Wireless queues in Poisson interference fields: the continuum between zero and infinite mobility

Nithin S. Ramesan
Department of ECE
University of Texas at Austin
Austin, USA
nithinseyon@utexas.edu

François Baccelli
Departments of Mathematics & ECE
University of Texas at Austin
Austin, USA
francois.baccelli@austin.utexas.edu

***Abstract*—This paper considers the time evolution of a queue that is placed in a Poisson point process of moving wireless interferers. The queue evolves according to an external arrival process and a time-varying service process that is a function of the SINR that it experiences. Static configurations of interferers result in infinite queue workload with positive probability. In contrast, a velocity-independent stability condition for the queue is established in the case where interferers possess any non-zero mobility that results in displacements that are both independent across interferers and with a distribution which is invariant to interferer positions. The proof leverages mixing properties of point processes under non-zero mobility. The effects of increasing mobility on statistical averages of queueing metrics are studied, and convex ordering tools are used to establish that faster moving interferers result in a queue workload that is dominated in the increasing convex sense by a queue workload resulting from slower interferers. As a corollary, it is shown that mean queue workload and mean delay improve as network mobility increases. It is shown that there is positive correlation between SINR level-crossing events at finitely separated time points. The notion of correlation between interference over time is made precise via an explicit correlation function. System behaviour is empirically analyzed using discrete-event simulation and the impact of the mobility model on system level performance is evaluated.**

***Index Terms*—Time-varying queues, queues in random environments, mobility, interference correlation, ad hoc networks, stochastic geometry, network dynamics, stochastic ordering, convex ordering.**

I. Introduction

With the advent of self-driving cars or drones, for example, and more generally autonomous vehicles, the analysis of the effects of motion on wireless networks is increasingly relevant. Current wireless networks inherently involve a certain degree of motion - users in cellular networks are very often mobile, and there exist multiple use cases for ad hoc networks that involve mobile nodes. The effect of mobility on the performance of wireless networks has been extensively investigated (see [1] & [2] and references therein). This paper, however, takes a different viewpoint on performance - that of queueing delay, or latency. To study this metric, we model our system by relaxing the full buffer assumption that is widely used in wireless network analysis (for example, in [3] and throughout [4]). This rather ubiquitous assumption assumes that wireless transmitters always have packets to transmit and ignores queueing issues. While this assumption may be an attractive one due to its simplification of the analysis of wireless networks, most buffers behave in reality as queues. Further, the full-buffer assumption abstracts out the aforementioned metric of queueing delay - sacrificing knowledge of packet-level effects for tractability in network-level ones.

Our model assumes that a receiver-transmitter pair of interest has a queue associated with it - a queue whose departure process is affected by the interference it experiences, and hence by the positions of interfering nodes around it, which are modelled as a Poisson point process. We assume that these points are moving according to a mobility model with unconstrained but *finite* velocity. This motion must be uncorrelated across interferers - the displacements that two different interferers see in any interval of time must be independent - and these displacements should not be a function of the locations of interferers in space. Our model studies the continuum of velocities in the interval $[0, \infty)$ and leads to scenarios where the point processes seen at two different time instants are correlated. We will note at this juncture that significant novelty lies in this correlation - most analytical work that considers mobility in networks uses uncorrelated mobility models for tractablility - for instance, an i.i.d. redistribution of nodes in time slots (for example, see [5] and [6]). The correlation we see in interferer configurations will induce time-correlations in the service of the queue, correlations that make analysis of the behaviour of the queue highly non-trivial - as we will see in the sequel.

The analysis of this problem hence involves dealing with two types of network dynamics - the motion associated with nodes which triggers some evolution of the interference seen by the queue, and the queueing dynamics that induces changes in the lengths of buffers and hence has implications on metrics like delay. The two types of dynamics that we mention are coupled - the time evolution of the queue considered depends directly on the time evolution of the positions of interferers around it. In a broader sense, the interplay of stochastic geometry and queueing is a challenging problem - the former deals with space averages, and the latter with time averages. We will see here however that mobility unifies these disjoint paradigms.

Contributions. The two main questions that we answer are queueing-theoretic in nature - when, if at all, is this system stable, and how do queue statistics change with varying degrees of mobility? The first question considers a notion of capacity that arises from queueing theory, i.e., the maximum input data rate that a buffer can support whilst remaining stable. In answering this question, we show that in the absence of mobility, universal stability guarantees do not exist. This result stems from the observation that a static network is exactly that - static and unchanging, and hence a queue will see the same interference at all times - dooming it to stability or instability forever, subject to the caprice of the interferer configuration around it. But all

This work is supported in part by the National Science Foundation under Grant No. NSF-CCF-1514275 and an award from the Simons Foundation (#197982), both to the University of Texas at Austin.

ISBN 978-3-903176-29-4 © 2020 IFIP

is not lost - the introduction of *any* non-zero mobility results in enough change in the system to guarantee stability of the queues - independently of the degree of mobility in the system. The second question considers the queueing delay experienced by packets in this queue, and the effects of mobility on this metric. Considering a continuum of velocities, we find that increasing motion in the network results in an accelerated rate of variation in the service rate of the queue - which in turns induces an averaging effect that provides increasingly reliable, or smooth service. The overall effect is that of decreasing mean queue workload and mean delay. The continuum of velocities hence is shown to induce a continuum of ordered queue workloads. In addition to these main results, we also provide results on the correlations that exist in mobile wireless networks, and connect them to the phenomena we observe. We supplement our theoretical results with simulations that study the effects of different mobility models.

A. Related Work

1) Mobility and wireless network performance

The effects of mobility on the performance of wireless networks have been studied extensively, beginning with [5], and we do not provide a comprehensive survey here. The vast majority of work in this area, however, is concerned with the concepts of transport capacity or throughput capacity, as first proposed in [3] or of multihop delay (as in [7]) rather than queueing capacity and delay, as is our focus. These works also do not consider the traffic dynamics of the systems they study by virtue of assuming a backlogged/full buffer. Nonetheless, a common thread between this body of work and our work is the observation that mobility improves performance.

2) Time-varying and Markov modulated queues

Our model involves a queue that sees time-varying service rates. There is an existing line of work on time-varying queues modelled using non-homogenous Poisson processes as arrival and service processes (e.g. [8], [9]) - our service process is not amenable to this framework because it is not a Poisson process. Modelling more general arrival and service processes has been accomplished by considering Markov-modulated queues - where arrival and service rates are functions of the state of a continuous-time Markov chain with finite state space (the classical Gilbert model of queueing theory is a special case where the CTMC has two states). Under this framework, computational methods to calculate the steady state probabilities have been proposed ([10], [11]). To attempt to use such an approach, our model will have to be made Markovian by assuming knowledge of the interferer point process instead of SINR. Its evolution will then have to be modelled as a CTMC with infinite dimensional state space (one coordinate for every interferer) - rendering this method intractable.

3) Local Delay And Mobility

The concept of local delay, introduced in [12], is the random number of time slots required by a node to transmit a packet successfully to its intended (single-hop) destination. This can instead be thought of as the number of time slots required for the packet located at the head of a queue such as the one we consider to be successfully transmitted. The local delay hence implicitly ignores queueing-level effects. The mean local delay is computed as a spatial average of the local delays of all nodes in a network, where by contrast, our queueing system involves temporal evolution and averaging to obtain queue statistics. The rather unrealistic infinite mobility model is applied to local delay in [12] and [6]. Another work that considers finite mobility ([13]) uses a constrained mobility model where locations of nodes are independent excursions from a fixed home-location point process, allowing for tractability. In contrast, we consider an unconstrained mobility model.

4) Queueing in Wireless Networks

There has been a recent body of work that seeks to unify queueing and wireless networks ([14]). A number of these works seek to establish necessary and sufficient conditions for stability using stochastic dominance ideas (see e.g. [15]), but these conditions are often loose. Other works make independence assumptions that decorrelate interference across time slots ([16]), which we do not do. The work closest to ours in terms of model, [17], assumes an infinite mobility model. Another line of thought treats the entire wireless network as a queue, dealing with the birth and death of transmitters and receivers ([18], [19]).

II. Model

We assume a continuous time model. We assume that interferers (which are for all intents and purposes, transmitters) are initially distributed as a homogenous Poisson point process Φ_0 of intensity Λ on $\mathbb{R}^2$. Denote by Φ_t the point process of interferers at time t. We assume that the receiver of interest is placed at the origin. Its transmitter transmits at unit power. The location of the transmitter is a point chosen uniformly at random on the circle centered at the origin, with radius R. We can also analyze the case where the receiver is moving using the same approaches - we will revisit this later in Section IV. We assume that all interferers transmit with unit power at all times. But it is straightforward to generalize this model to one that incorporates spatial ALOHA, simply by independently thinning the point process of interferers. A path-loss function $l(.) : \mathbb{R}^+ \to \mathbb{R}^+$ is assumed, with the following properties: 1) $\lim_{r\to\infty} l(r) = 0$ (in other words, we assume that an interferer that is infinitely far away does not have any effect at the origin), and 2) $\int rl(r)dr < \infty$ (which ensures that the mean interference shot-noise at the origin is finite). Denote by $F_j^0(t)$ the fading variable between interferer j and the origin at time t. We assume that $F_j^0(t)$ is stationary. Denote by $\sigma^2(t)$ the thermal noise power at time t. Let $S_0(t)$ be the received signal power at the origin, and $I_0(t)$ be the interference seen at the origin. The SINR seen by the receiver at the origin at time t can then be described by:

$$\text{SINR}_0(t) = \frac{S_0(t)}{I_0(t) + \sigma^2(t)} = \frac{l(R)F_0^0(t)}{\sum_{x_j \in \Phi_t} l(|x_j|)F_j^0(t) + \sigma^2(t)}. \quad (1)$$

Note that the primary object of interest is the behaviour of $I_0(t)$, which is the value at the origin of a time-space interference shot noise field (see [4], Section 2.3.4), with the positions of interferers changing over time. This time variation is governed by our mobility model. We will now present three examples of mobility models that satisfy the conditions mentioned in the introduction.

Random Direction Model (RD): Assume that at the beginning of time, every interferer i samples an angle θ_i uniformly at random from the interval $[0, 2\pi]$ and independently. Each interferer will then move with constant velocity $v \in [0, \infty)$ along this angle - hence, in a time interval Δt, interferer i will be displaced by $(v\Delta t \cos\theta_i, v\Delta t \sin\theta_i)$.

Random Waypoint Model (RWP): We consider a version of the random waypoint model where every interferer i moves with constant velocity v and angle θ_i sampled uniformly and independently from $[0, 2\pi]$ for a deterministic amount of time Δp. At the end of every such time interval, every node resamples its angle of motion and continues moving.

Brownian Motion (BW): We assume that every interferer moves according to an independent 2-D Wiener process, with the variance of each one-dimensional Wiener process being γ^2. The magnitude of displacement is hence given by a Rayleigh distribution with parameter γ, whose mean is $\gamma\sqrt{\frac{\pi}{2}}$. In a small time interval δs, an interferer moving with velocity v will be displaced on average by $v\delta s = \gamma\sqrt{\frac{\pi}{2}}$. This can easily be extended to more general diffusion processes.

By the displacement theorem for PPPs ([4], Theorem 1.3.9), it follows that, for all three models, Φ_t is a homogenous PPP for all t. Note that all queue comparison results presented in Section IV hold for these models and for more general mobility models that satisfy the aforementioned independence conditions.

A. Queueing model

We assume that a queue is associated with the transmitter-receiver pair of interest. Data arriving to the transmitter enter the queue and must be transmitted to the receiver, upon which the data leave the queue. To describe the queueing dynamics, we assume that a discrete time scale is overlaid on our continuous time scale, with time slots of duration δ. We consider a single-server infinite buffer queue. We assume that packets, each of unit size, arrive to the queue according to an external discrete time stationary and ergodic arrival process, $A(n)$, which is independent of the spatial PPP and the interferer motion processes. In other words, the workload W_n (amount of data waiting in the queue for transmission) at time $n\delta$ of the queue increases by $A(n)$. Note that we can approximate the case of a continuous time arrival process by making δ arbitrarily small. The departure process associated with this queue depends on the time-series of interference that the receiver sees. We assume a continuous-time random service process, $s(t)$, upon which we place the dual restrictions that $s(t)$ must be a measurable function of $I_0(t)$, and that $s(t)$ must be integrable. Examples of possible service processes include: a) Shannon rate: $s(t) = \log_2(1 + \text{SINR}_0(t))$ and b) truncated Shannon rate: $s(t) = \log_2(1 + \text{SINR}_0(t))\mathbb{1}[\text{SINR}_0(t) > T]$, where the former assumes adaptive modulation and coding based on SINR, and the latter assumes the same as long as the SINR is above a certain threshold T. It is easy to see that $s(t)$ is a stationary process. Such a model is a fluid queueing model - we are implicitly assuming that while data arrive as packets, they are transmitted continuously as a data stream. The cumulative service that the queue sees in the time interval $(t_1, t_2]$ is then denoted by $S(t_1, t_2) = \int_{t_1}^{t_2} s(t)dt$. Since the arrival process is in discrete time, the queue also evolves in discrete time according to the Lindley recursion:

$$W_{n+1} = (W_n + A(n) - V(n))^+ , \tag{2}$$

where W_n is the workload of the queue at time slot n or at time $n\delta$, $V(n) = S(n\delta, (n+1)\delta)$ and $x^+ = \max(0, x)$.

This queue is a G/G/1/∞ queue - where we emphasize that service to the queue is stationary, generally distributed, and *dependent*, or *not independent over time slots* - seeing a large service rate in one time slot will increase the likelihood of a large service rate in the next time slot, and similarly for small service rates as well. This fact stems from the finite mobility we consider in our model, which induces positive time-correlations in interferer configurations and hence in $s(t)$. We will make the idea of this correlation more precise in Section V. This lack of independence makes an already non-trivial problem even more complex.

III. Mobility Guarantees Stability

In this section, we will contrast the effects of a static configuration of inteferers on the evolution of our queue with the effects of a configuration of interferers that possess a non-zero degree of mobility, modelled using the random direction model of the last section. The mixing results that we show in this section will extend to the other two models we presented, but we do not include these proofs here. For this section, denote by $(\Omega, \mathcal{F}, \mathbb{P})$ the probability space upon which the initial PPP of interferers is defined. Consider a static version of our model, with $v = 0$. The initial configuration of interferers will be fixed for all times, i.e., $\Phi_0 = \Phi_t\ \forall t$. The service process $V(n)$ is clearly stationary, but it is not ergodic - time averages of $V(n)$ will not be the same as ensemble averages over point process realizations. Assume that the arrival process $A(n)$ has rate $\lambda = \mathbb{E}[A(n)]$.

Remark III.1. *Since the queue workload is driven by a non-ergodic sequence $V(n)$, Loynes' Theorem does not apply for this queue (see Section 2.1, [20] for the necessary conditions to apply Loynes' Theorem) and stability is not guaranteed for any positive λ.*

Indeed, we will now see that instances of *instability* are present in the system for all positive λ. Note that there exists a subset Ω_U of Ω that has non-zero measure and that corresponds to point process realizations that have interferers clustered near the origin, resulting in the instability of the queue. The other side of this coin is also true - there will be interferer configurations (for $\omega \in \Omega_U^C$) where all interferers are far enough away that the queue will experience high rates of service. The stationary queue workload for $\omega \in \Omega_U^C$ will be finite, and the queue workload for $\omega \in \Omega_U$ will be infinite. Characterizing Ω_U is straightforward - $\Omega_U = \{\omega \in \Omega : \lambda > \mathbb{E}[s(t)|\Phi(\omega)]\}$, with $\Phi(\omega) = \Phi_0$. As an example, let $s(t)$ be the Shannon rate, $s(t) = \log_2[1 + \text{SINR}_0(t)]$. Then, $\Omega_U = \left\{\omega \in \Omega : \lambda > \mathbb{E}\left[\log_2\left(1 + \frac{l(R)F_0^0(t)}{\sum_{x_j \in \Phi(\omega)} l(|x_j|)F_j^0(t) + \sigma^2(t)}\right)\middle|\Phi(\omega)\right]\right\}$.

Proposition III.2. *The mean queue workload for a static field of interferers (averaged over the point process of interferers, fading and the arrival process) will be infinite for all positive λ. Equivalently, for all positive λ, the network will have a positive fraction of unstable queues.*

Proof. Follows from noting that the stationary queue workload is infinite on Ω_U, a set with non-zero measure on the space $(\Omega, \mathcal{F}, \mathbb{P})$. □

It is clear that we cannot guarantee sample path stability, nor stability in a mean sense. From a system design perspective, this is not ideal - it is preferable that we can guarantee that for λ small enough, all our wireless buffers are *always* stable, with finite workload. The solution to this need, as we will show now, is to introduce mobility into the network.

Definition III.3. [21] Strong mixing: We say that a time-indexed stationary stochastic process $\{\Psi_t\}$ is strongly mixing if $\mathbf{P}\{\Psi_t \in \Gamma, \Psi_{t+s} \in \Delta\} \to \mathbf{P}\{\Psi_t \in \Gamma\}\mathbf{P}\{\Psi_{t+s} \in \Delta\}$ as $s \to$

∞, for all configuration sets Δ, Γ.

We then have the following theorem:

Theorem III.4. *If $v > 0$, the process $\{\Phi_t\}$ is strongly mixing.*

Proof. Denote by s the random displacement that transforms the point process Φ_t to Φ_{t+1} for any time t. This random displacement is characterized by a law $G(.)$, which in our case uniformly samples angles for every point at time $t = 0$ from the interval $[0, 2\pi]$, while holding speed constant. Hence, s takes values in the boundary of the 2-D ball centered at the origin with radius v, $\partial B(0, v)$. From the definition of our mobility model, we know that the displacement that transforms the point process Φ_t to Φ_{t+n} is given by ns. To prove the theorem, we must show that in the limit $n \to \infty$, the point processes Φ_t and Φ_{t+n} are asymptotically independent. We will accomplish this using the joint Laplace functional of the two point processes.

We can express the joint Laplace functional of two arbitrary Poisson point processes, Φ (homogenous with intensity Λ) and Φ' (where Φ' is obtained from Φ via transformation by a probability kernel $p(x, .)$, as:

$$\mathbb{E}\left(e^{-\int_{\mathbb{R}^2} f(x)\Phi(dx)} e^{-\int_{\mathbb{R}^2} g(y)\Phi'(dy)}\right) = \mathbb{E}\left(e^{-\sum_i f(X_i) - \sum_i g(Y_i)}\right)$$

$$= \mathbb{E}\left(\left(\exp\left[\sum_j \log e^{-f(x_j)}\right]\right)\right.$$

$$\left.\left(\exp\left[\sum_j \log\left(\int_{\mathbb{R}^2} e^{-g(y_j)} p(x_j, dy_j)\right)\right]\right)\right)$$

$$= \mathcal{L}_\Phi(h), \text{ where } h(x) = -\log\left[\int_{\mathbb{R}^2} e^{-f(x)-g(y)} p(x, dy)\right]$$

$$= \exp\left[-\int_{\mathbb{R}^2}\left(\left(1 - \int_{\mathbb{R}^2} e^{-f(x)-g(y)}\right) p(x, dy)\right)\Lambda(dx)\right]$$

$$= \mathcal{L}_{\Phi'}(g) \exp\left[-\int_{\mathbb{R}^2}\left(1 - e^{-f(x)}\right)\right.$$

$$\left.\left(\int_{\mathbb{R}^2} e^{-g(y)} p(x, dy))\right)\Lambda(dx)\right].$$

Setting $\Phi = \Phi_t$, $\Phi' = \Phi_{t+n}$, we get:

$$\mathcal{L}_{\Phi,\Phi'}(f, g) = \mathcal{L}_{\Phi'}(g) \exp\left[-\int_{\mathbb{R}^2}\left(1 - e^{-f(x)}\right)\right.$$

$$\left.\left(\int_{\partial B(0,v)} e^{-g(x+ns)} G(ds)\right)\Lambda dx\right].$$

As long as $g(.)$ is such that $g(x) \to 0$ in all directions as $x \to \infty$, we have that $\left(\int_{\mathbb{R}^2} e^{-g(x+ns)} G(ds)\right) \to 1$ as $n \to \infty$. This follows directly from the Dominated Convergence Theorem, where we will upper-bound the exponential by 1. This in turn implies that:

$$\lim_{n\to\infty} \mathcal{L}_{\Phi,\Phi'}(f, g) = \mathcal{L}_\Phi(f)\mathcal{L}_{\Phi'}(g).$$

□

Corollary III.4.1. *In the absence of fading, the interference shot-noise process, $I_0(t)$, is strongly mixing.*

Proof. The proof follows from Theorem III.4 and the observation that measurable transformations preserve strong mixing (Thm 5.2, [21]). □

While considering the presence of fading, it is clear that the shot-noise will be strongly mixing when the fading processes at two points sufficiently separated in space are independent of each other.

Since strong mixing implies ergodicity, we then have that $s(t)$ and $V(n) = S(n\delta t, (n+1)\delta t)$ are ergodic processes.

We have now established that the driving sequences of the queue, $A(n)$ and $V(n)$ are ergodic and stationary. Since we also assumed that they are independent of each other, these are sufficient conditions to apply the theory of Loynes ([22]) to our queue, and to obtain the main result of this section.

Theorem III.5. *When $v > 0$, all network queues are stable (i.e., their workloads have a finite limiting distribution) if and only if $\lambda = \mathbb{E}[A(n)] < \delta\mathbb{E}[s(0)]$.*

Note that here $\mathbb{E}[s(0)]$ is a spatial average over the PPP. As an example, if $s(t) = \log_2[1 + \mathrm{SINR}_0(t)]$ and assuming that $\sigma^2(t) = 0$, fading is Rayleigh and $l(r) = (Ar)^{-\beta}, A > 0, \beta > 2$, then $\mathbb{E}[s(0)] = \int_0^\infty \exp(-2\pi^2\Lambda R^2 v^{\frac{2}{\beta}}\beta^{-1}(\sin(2\pi/\beta))^{-1})/(v+1))dv$ (see Section 16.2.3, [23]). While the structure of Theorem III.5 may seem obvious, we would like to emphasize its significance. What we have shown is that introducing any non-zero mobility, however small, into the system causes the non-ergodic and unchanging nature of the queue to vanish - leaving us with a queue that can be stabilized and cope with *all Poisson interferer configurations*. Indeed, under the assumption of non-zero mobility, it follows from mixing that the queue will eventually see all possible Poisson interferer configurations. Motion in the network has in a sense unified the spatial averages that are traditionally associated with stochastic geometry and the temporal averages that are traditionally associated with queueing theory. Finally, we note the parallel to the infinite mobility models often considered in the analysis of mean local delay - in those models, the point process decorrelates in a single time slot. In ours, the point process decorrelates after a sufficiently long amount of time - which, as we show, is sufficient for a well-behaved (i.e., stable) queue.

IV. Mobility Orders Queue Workloads

Now that we are assured of the queue's stability, we will investigate the effect of increasing degrees of mobility on the statistics of the queue workload. To begin, we define useful partial orderings.

Definition IV.1. Convex order ($\leq_{cx}$): Consider two random variables X and Y. X is said to be smaller than Y in the convex order, denoted by $X \leq_{cx} Y$, if $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$ for all convex functions f, provided expectations exist.

Definition IV.2. Increasing convex order ($\leq_{icx}$): Consider two random variables X and Y. X is said to be smaller than Y in the increasing convex order, denoted by $X \leq_{icx} Y$, if $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$ for all increasing convex functions f, provided expectations exist.

For further results on convex orderings, see [24]. A useful equivalent definition for these orderings is presented in the lemma below (taken from [20], Chapter 4).

Lemma IV.3. *Given two random variables X and Y, $X \leq_{cx}$ (resp. $\leq_{icx}$)Y if and only if there exist two random variables A and B, identically distributed as X and Y respectively and defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$, such that $A =$ (resp. $\leq$) $\mathbb{E}[B|\mathcal{G}]$ a.s., for some sub σ-field $\mathcal{G}$ of $\mathcal{F}$.*

Now, we will consider 3 environments for a queue and its moving interferers. These environments differ only in the velocities of the interferers - they have the same arrival process, $A^{(1)}(n) = A^{(2)}(n) = A^{(3)}(n) = A(n)$. Let the velocities be v_1, v_2 and v_3 respectively, with $v_2 = mv_1$ and $v_3 = nv_1$ and $n > m$ ($n, m \in \mathbb{N}$, $n > 1$, $m > 1$), so that $v_3 > v_2 > v_1$. Let $s^{(i)}(t)$ be the instantaneous service rate that the queue can provide in Environment i. Let the cumulative service that the queue can provide in the time interval $(t_1, t_2]$ for Environments 1, 2 and 3 be $S^{(1)}(t_1, t_2)$, $S^{(2)}(t_1, t_2)$ and $S^{(3)}(t_1, t_2)$ respectively. Let the workloads of the queue at time slot n be $W_n^{(1)}$, $W_n^{(2)}$ and $W_n^{(3)}$ respectively. The key observation now is that an increase in the velocity of interferers is equivalent to an acceleration of time while keeping velocity fixed. Consider a simple example - the distance an interferer moving with velocity $2v$ will cover in time Δt is the same as the distance an interferer moving with velocity v will cover in time $2\Delta t$. In the context of the environments described above, scaling velocities by a factor of $m \in \mathbb{N}$ is equivalent to accelerating time by a factor of m. Hence, we have

$$\begin{aligned} S^{(2)}(t_1, t_2) &= \int_{t_1}^{t_2} s^{(1)}(mt)dt = \frac{1}{m}\int_{mt_1}^{mt_2} s^{(1)}(t)dt \\ &= \frac{1}{m}\sum_{i=1}^{m}\int_{mt_1+(i-1)(t_2-t_1)}^{mt_1+(i)(t_2-t_1)} s^{(1)}(t)dt \\ &= \frac{1}{m}\sum_{i=1}^{m} S^{(1)}(mt_1 + (i-1)(t_2-t_1), mt_1 + i(t_2-t_1)) \\ &= \frac{1}{m}\sum_{i=1}^{m} Y_i^{(2)}, \end{aligned}$$

where the variables $Y_i^{(2)} = S^{(1)}(mt_1 + (i-1)(t_2 - t_1), mt_1 + i(t_2 - t_1))$ are identically distributed (since $s(t)$ is stationary) but not independent. Similarly, we have

$$S^{(3)}(t_1, t_2) = \frac{1}{n}\sum_{j=1}^{n} Y_j^{(3)},$$

for $Y_j^{(3)} = S^{(1)}(nt_1 + (i-1)(t_2-t_1), nt_1 + (i)(t_2-t_1))$, where $\{Y_j^{(3)}\}$ is distributed identically to the variables $\{Y_i^{(2)}\}$. Note also that $S^{(1)}(t_1, t_2) = Y_j^{(1)}$ is identically distributed to $\{Y_j^{(2)}\}$ and $\{Y_j^{(3)}\}$. Henceforth, we will drop the superscript for the variables Y_i, and write

$$S^{(1)}(t_1, t_2) = Y_i, S^{(2)}(t_1, t_2) = \frac{1}{m}\sum_{i=1}^{m} Y_i,$$

$$S^{(3)}(t_1, t_2) = \frac{1}{n}\sum_{i=1}^{n} Y_i. \tag{3}$$

Now, let $M_d^{(i)} \in \mathbb{R}^d$ be a vector whose j^{th} element is $S^{(i)}(j\delta t, (j+1)\delta t)$, for $d > j \geq 0$. Then, we have the following lemma.

Lemma IV.4. *$M_d^{(3)} \leq_{cx} M_d^{(2)} \leq_{cx} M_d^{(1)}$ for all $d > 0$.*

Proof. From (3), we see that we can write $M_d^{(1)} = Z_i, M_d^{(2)} = \frac{1}{m}\sum_{i=1}^{m} Z_i$, and $M_d^{(3)} = \frac{1}{n}\sum_{i=1}^{n} Z_i$, where Z_i is an $d-$dimensional vector, with each element of the vector distributed identically to Y_0 (but not necessarily independent of the other elements in the vector). Further define $X_n = \sum_{i=1}^{n} Z_i$ and $\overline{X}_n = \frac{X_n}{n} \in \mathbb{R}^d$. We will now show that $\overline{X}_{n-1} \geq_{cx} \overline{X}_n$. The result will then follow by the transitivity of the convex ordering.

First, note that $\mathbb{E}[Z_i|X_n]$ is a function of X_n, and is independent of i due to the stationarity of $\{Z_i\}$. Let $\mathbb{E}[Z_i|X_n] = \Gamma(X_n)$. Then,

$$\Gamma(X_n) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_i|X_n] = \mathbb{E}\left[\frac{X_n}{n}|X_n\right] = \overline{X}_n.$$

Since $\mathbb{E}[Z_i|X_n] = \overline{X}_n$, we have that

$$\mathbb{E}[Z_1 + Z_2 + ... + Z_{n-1}|X_n] = (n-1)\overline{X}_n,$$

which implies $\mathbb{E}[\overline{X}_{n-1}|\overline{X}_n] = \overline{X}_n$. From Lemma IV.3, we have that $\overline{X}_{n-1} \geq_{cx} \overline{X}_n$. □

The intuition behind Lemma IV.4 is straightforward - sample means calculated using a larger number of samples will possess lower variability, resulting in a convex ordering. Hence, an accelerated service process will result in the queue "seeing" a larger number of samples of instantaneous service rate, which in turn will achieve an averaging effect that results in more reliable service being provided to the queue. It follows that this increasing reliability of service will reflect in the performance of the queue, a qualitative result that we will present in Theorem IV.6.

Now, consider two sequences of variables $\{\widetilde{y}_i\}$ and $\{y_i\}$, $i \geq 0$, whose evolution is governed by the same stochastic recurrence function h, and two sets of driving sequences $\{\widetilde{\beta}_n\}$ & $\{\beta_n\}$ and initial conditions $\widetilde{y}_0$ & y_0 respectively:

$$y_{n+1} = h(y_n, \beta_n)$$
$$\widetilde{y}_{n+1} = h(\widetilde{y}_n, \widetilde{\beta}_n).$$

These sequences have the same dynamics, characterized by h, and differ only in their driving sequences and initial conditions. Then, we have the following lemma:

Lemma IV.5. *([20], Property 4.2.5): Assume that the driving sequences and initial conditions are integrable, and that the function $y \to h(y, \beta)$ is non-decreasing and that the function $(y, \beta) \to h(y, \beta)$ is convex. Then, $(y_0, \beta_0, \beta_1, ...) \leq_{cx} (\widetilde{y}_0, \widetilde{\beta}_0, \widetilde{\beta}_1, ...)$ implies $(y_0, y_1, y_2, ...) \leq_{icx} (\widetilde{y}_0, \widetilde{y}_1, \widetilde{y}_2, ...)$.*

We now present the main theorem of this section.

Theorem IV.6. *For queues that start with initial workload $W_0^i = 0$ and evolve according to the Lindley recursion (2), and for environments as defined previously,*

$$W_n^{(3)} \leq_{icx} W_n^{(2)} \leq_{icx} W_n^{(1)}, \forall n. \tag{4}$$

Further, the steady state workloads that the queues converge to in the limit are also similarly ordered:

$$W_\infty^{(3)} \leq_{icx} W_\infty^{(2)} \leq_{icx} W_\infty^{(1)}. \tag{5}$$

Proof. The convex ordering is preserved by the operation of multiplication by -1 and by the addition of an independent random variable. Let $A_n \in \mathbb{R}^n$ be a vector whose j^{th} element is $A(j)$, for $n > j \geq 0$. From Lemma IV.4,

$$A(n) - M_n^{(3)} \leq_{cx} A(n) - M_n^{(2)} \leq_{cx} A(n) - M_n^{(1)}.$$

Now, we define $h(y, \beta) = (y + \beta)^+$, and set y_n to be queue workload W_n, initial conditions $y_0 = \widetilde{y}_0 = 0$ and $(\beta_0, \beta_1, ..., \beta_{n-1}) = A_n - M_n$. Then the proof of the first part of the theorem follows from Lemma IV.5. Since we are operating in the regime where the queue is stable, we know that the queue workload will converge to a stationary steady state variable $W_\infty^{(i)}$. The second part of the theorem then follows from a direct application of the Monotone Convergence Theorem (see [20], Section 4.2.6 for more details). □

We emphasize again here that this result holds for queues that are driven by *correlated* service processes, and is a significantly more general result than results of the Pollaczek-Khinchine type, which hold only for independent service processes.

A first consequence of this theorem is the following corollary.

Corollary IV.6.1. $\mathbb{E}[W_\infty^{(3)}] \leq \mathbb{E}[W_\infty^{(2)}] \leq \mathbb{E}[W_\infty^{(1)}]$.

It also follows that the mean delay that packets see are similarly ordered. Let us assume that a packet is successfully transmitted once all the data associated with it are transmitted (assume a packet is of size k bits), and the delay a packet sees is the length of the time interval from its arrival to its succesful transmission. Then, at time slot n, the (fractional) number of packets in the queue will be $N_\infty = \frac{W_\infty}{k}$. N_∞ is hence ordered in expectation in a way similar to W_∞. It follows from Little's Law that the mean delay, $\mathbb{E}[D] = \mathbb{E}[N_\infty]/\mathbb{E}[A(0)]$ is also ordered in the same way, provided that it exists.

Corollary IV.6.2. $\mathbb{E}[D^{(3)}] \leq \mathbb{E}[D^{(2)}] \leq \mathbb{E}[D^{(1)}]$.

A. Discussion

These results state that queues that see the same arrival process see decreasing workloads and delays as they are placed amongst interferers that move with increasing mobility. Hence, the performance of queues in wireless networks improve as the degree of mobility in the network increases. This performance improvement can be intuitively thought of as a consequence of mobility introducing diversity into the network. The diversity that we refer to is with respect to interferer configurations - faster moving interferers allow a larger number of configurations to be seen at a higher rate. In other words, the network mixes faster. An intuitive explanation that is a dual to the faster averaging/mixing argument is the following scenario. Assume the queue is surrounded by a "bad" configuration of interferers - namely, a large number of interferers are located close to the origin. The interference seen at the origin will hence be large. A slowly moving environment will cause the interference at the origin to be large for a long period of time, which will result in a build-up in the workload of the queue. On the flip side, interferer configurations that are "good" will also persist for long periods of time, allowing the queue to drain effectively during those periods. A fast-moving environment, on the other hand, will result in shorter build-ups and drainings of the queue. Fig. 1 shows the difference in the lengths of these cycles and hence in the average queue workload for environments with different degrees of mobility. From an expectation point-of-view,

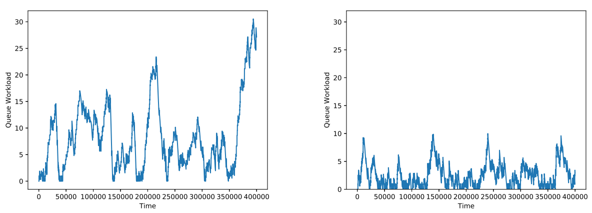

(a) Slow environment ($v = 1$)

(b) Fast environment ($v = 1000$)

Fig. 1: Comparison of queue build-up-and-drain cycles for different degrees of motion in the random direction model

it seems clear that longer build-up-and-drain cycles will result in larger mean workloads and mean delays, which is what we see in Corollary IV.6.1 and IV.6.2. Interestingly, the more general result of Theorem IV.6 also holds - increasing convex functions of workloads are also similarly ordered.

We would like to draw attention to the generality of results presented in this section. First, the results presented are such that in order to compare the workloads of two environments, the velocity of one environment must be a rational multiple of the other. Since the rationals are dense in the reals, we do not lose any generality here. In terms of modelling, note that we make no assumptions on the service process $s(t)$ save that it is integrable and a measurable function of $I_0(t)$. We only need a stationarity assumption for the fading processes. These results hold for all the mobility models described in Section II. Indeed, these results will hold for any general mobility model that admits the interpretation of an increase in velocity as an acceleration in time. The case of the receiver at the origin and its corresponding transmitter being mobile as opposed to static is subsumed under this set of mobility models - the relative velocities of interferers with respect to the tagged receiver will no longer be independent of each other, but will still be amenable to the framework used in this section. We briefly discuss the limiting case of infinite velocity. Under this mobility assumption, the point process of interferers will be resampled at every time instant. This will decorrelate the service process - the variables Y_i will now be i.i.d. and the service process seen in a time interval (t_1, t_2) will be a constant - the infinite averaging will result in the expected service rate being seen. This results in a G/D/1 queue. Making further assumptions on the arrival process and network model could lead to analytical expressions for queue statistics, but we will not pursue this here.

We have hence shown the existence of a continuum of queues corresponding to the continuum of velocities between ∞ and 0 - which correspond to a stationary and ergodic G/D/1 queue and a non-ergodic G/G/1 queue that sees infinite queue workload in a subset of the sample space that has non-zero measure respectively. We move from the former extreme to the latter via a sequence of queue workloads that are increasingly convex ordered - the expectation of any increasing convex function of the queue workload decreases along this sequence.

V. Correlation structures in mobile networks

In this section, we formalize the correlations that we discussed while presenting intuition for the results in the previous section. First, we consider SINR level set events - L_t is defined as the event that at time t, $\text{SINR}_0(t) > T$ for some fixed threshold T. We will show that if we see a level set event at time t, there is an increased likelihood of seeing another level set event in some neighbourhood of that time. Assume Rayleigh fading between the transmitter and receiver of interest (modelled by exponential random variables of mean $1/\mu$), and let all other fading variables be generally distributed but with unit mean. We then have the following theorem.

Theorem V.1. $\mathbb{P}(L_{t+1}|L_t) \geq \mathbb{P}(L_{t+1}) = \mathbb{P}(L_t)$.

Proof. Presented in the appendix. □

Theorem V.1 formalizes the intuition we presented in Section IV. It implies that if we observe a high (resp. low) SINR, we are likely to continue to observe a high (resp. low) SINR for a while. The notion of buildup-and-drain cycles of the queue is a direct consequence of this effect. This positive temporal correlation of level-crossing events will decrease as the correlation between the point processes Φ_t and Φ_{t+1} decreases - or in other words, as mobility increases.

Finally, for completeness, we will derive the correlation coefficient between the interference shot-noise observed at times t and $t+1$, $\gamma(t, t+1)$, defined as:

$$\gamma(t, t+1) = \frac{\mathbb{E}[I_0(t)I_0(t+1)] - \mathbb{E}[I_0(t)]^2}{\mathbb{E}[I_0(t)^2] - \mathbb{E}[I_0(t)]^2}.$$

Lemma V.2.

$$\gamma(t, t+1) = \frac{\int l(x) \int l(y) p(x, dy) dx}{\mathbb{E}[h^2] \int l(x)^2 dx}. \quad (6)$$

Proof. We use Campbell's formula (Thm 1.4.3, [4]) to calculate the first and second moments of interference, and then use the fact that $\mathbb{E}[I_0(t) I_0(t+1)] = \frac{\partial^2 (L_{t,t+1}(s_1, s_2))}{\partial s_1 \partial s_2}\Big|_{s_1=0, s_2=0}$. □

We see that the correlation is highest when the probability kernel that defines mobility is such that $p(x, y) = 1$ for $x = y$, and is lowest in the limit of infinite mobility.

VI. Simulation Studies

Our simulation setup is as follows. We consider a 100×100 square whose edges are wrapped around to avoid edge effects. In order to approximate continuous time, we let our overall system evolve in very small intervals of discrete time, of length $\Delta = 10^{-3}$ seconds. The queue associated with the bipole placed at the origin evolves in time slots of length δ. We set $\delta = \Delta$ in simulations, which can be interpreted as approximating a continuous time queue. The arrival process $A(n)$ to the queue is a Bernoulli process of rate λ, where at every time slot a packet arrives with probability $\lambda\delta$, with λ chosen such that $\lambda\delta < 1$. Interfering nodes are distributed as a PPP, and move according to the mobility models described before. We assume that the service process for the queue is the truncated Shannon rate process described in Section II. The reduction in workload that the queue sees at time t in the interval Δ is hence $\log_2(1 + \text{SINR}_0(t)) \mathbb{1}[\text{SINR}_0(t) > T]\Delta$. The queue evolves according to (2). From Theorem III.5, to guarantee stability of the queue, we must ensure that $\mathbb{E}[A(n)] = \lambda < \mathbb{E}[s(0)]$ (since $\Delta = \delta$). We empirically estimate $\mathbb{E}[s(0)]$ for our system and set λ accordingly. For our simulations, we set $\Lambda = 0.1$, noise power $\sigma = 0$, $R = 0.3$, $T = 8$ and $l(r) = (1 + r)^{-4}$. We consider Rayleigh fades with mean 1.

To begin, we verify the ordering of mean queue workload that is indicated by Corollary IV.6.1. We estimate the mean workload, and 95% confidence intervals using the batch mean method detailed in Chapter IV.5 of [25]. Fig. 2 shows the anticipated trend in mean workload, $\mathbb{E}[W]$, for all mobility models. Mean delay, $\mathbb{E}[D] = \mathbb{E}[W]/\lambda$, will show the same trend. We see an order of magnitude difference between the mean workload for $v = 1$ and the mean workload for $v = 1000$. Next, we compare the effects of different mobility models on mean queue workload. This comparison can be thought of as a measure of how fast each mobility model causes the PPP of interferers to mix - the faster the network mixes, the lower the mean workload. Fig. 2 shows the comparison of the mobility models presented in Section II. We see the expected ordering of workloads. Further, we see that workloads associated with Brownian motion (BM) dominate those associated with the random waypoint model (RWP), which in turn dominate those associated with the random direction (RD) model. The faster mixing of the RD model is explained by noting that all interferers travel with a constant direction over time, while the directions of motion change over time in the RWP and BM models, which will lead to slower mixing. Further, the time-scale of this change in direction is smaller for BM than RWP, which results in the slower mixing of BM. Note that the RD model is used in the literature to model vehicular networks (see [26]). In contrast, the BM and RWP models have been used to model more erratic movement, such as that of pedestrians. Our simulation studies suggest that the former motion leads to better performance of wireless queues, due to its more directed nature.

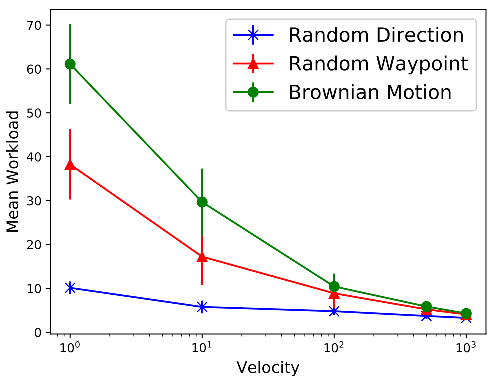

Fig. 2: Comparison of the effects of different mobility models. $\lambda = 1.2$, $\mathbb{E}[s(t)] \approx 1.37$, $\Lambda = 0.1$.

We also evaluate and plot the empirical CDF of latency (as described before Corollary IV.6.2) for the random waypoint model. We observe, via comparison with the CDF of the exponential distribution, how the distributions of latency become heavier with decreasing velocities. This observation has engineering implications for communication networks - in the context of our system, for example, this indicates that packets will often see large latencies at low mobility levels. We hence conclude that higher velocities will be advantageous in scenarios where reliable service is required, such as in URLLC systems.

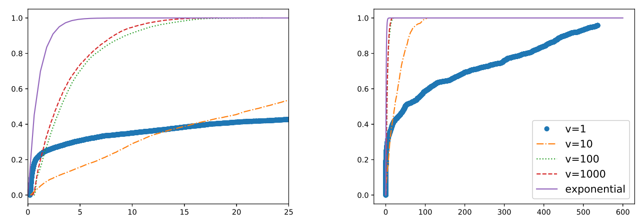

(a) X-axis = [0, 25] (b) X-axis = [0, 600]

Fig. 3: Empirical CDF of latency, observed at two different scales. Legends are the same for both plots.

VII. Conclusion

We have shown that introducing mobility in a large wireless network introduces mixing diversity that breaks non-ergodicity and allows us to state universal stability guarantees for a queue driven by the resulting interference process. We next investigated the continuum of mobility scenarios that arise from considering all velocities in the interval $(0, \infty)$. Accelerating the motion of interferers accelerates the mixing of the interference process, and the resulting averaging effect causes less variable and more reliable service that improves mean queue workload and mean delay. This effect can also be explained by considering the structure of correlations of the SINR and interference processes, which we have formalized. An interesting future line of thought is generalizing this model to an interacting queue model, where all interferers are themselves wireless queues. We conjecture that a similar stability condition and a similar ordering of workloads will exist in that model as well, for similar reasons. Another problem that remains to be pursued is to obtain (possibly approximate) closed-form expressions for mean queue workload and delay, so that the explicit dependence on the degree of mobility in the system can be studied. Finally, ongoing work involves statistically testing the distributions of latency for the occurrence of heavy tails, and investigating whether increasing velocity alleviates this phenomenon.

REFERENCES

[1] L. Ying, S. Yang, and R. Srikant, "Optimal Delay-Throughput Tradeoffs in Mobile Ad Hoc Networks," *IEEE Trans. on Info. Theory*, vol. 54, no. 9, pp. 4119–4143, 2008.

[2] W. Bao and B. Liang, "Stochastic geometric analysis of user mobility in heterogeneous wireless networks," *IEEE JSAC*, vol. 33, no. 10, pp. 2212–2225, 2015.

[3] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. on Info. Theory*, vol. 46, no. 2, pp. 388–404, 2000.

[4] F. Baccelli and B. Błaszczyszyn, *Stochastic Geometry and Wireless Networks, Volume I—Theory*. NoW Publishers, 2009.

[5] M. Grossglauser and D. N. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Trans. on Networking (ToN)*, vol. 10, no. 4, pp. 477–486, 2002.

[6] M. Haenggi, "Local delay in static and highly mobile Poisson networks with ALOHA," in *ICC*. IEEE, 2010.

[7] X. Lin, G. Sharma, R. R. Mazumdar, and N. B. Shroff, "Degenerate delay-capacity tradeoffs in ad-hoc networks with brownian mobility," *IEEE/ACM Trans. on Networking (TON)*, vol. 14, no. SI, pp. 2777–2784, 2006.

[8] W. A. Massey, "The analysis of queues with time-varying rates for telecommunication models," *Telecommunication Systems*, vol. 21, no. 2-4, pp. 173–204, 2002.

[9] Z. Feldman, A. Mandelbaum, W. A. Massey, and W. Whitt, "Staffing of time-varying queues to achieve time-stable performance," *Management Science*, vol. 54, no. 2, pp. 324–338, 2008.

[10] M. F. Neuts, "The M/M/1 Queue with Randomly Varying Arrival and Service Rates." Delaware Univ., Newark, Tech. Rep., 1977.

[11] E. P. Kao and C. Lin, "The M/M/1 queue with randomly varying arrival and service rates: a phase substitution solution," *Management science*, vol. 35, no. 5, pp. 561–570, 1989.

[12] F. Baccelli and B. Blaszczyszyn, "A new phase transitions for local delays in MANETs," in *IEEE INFOCOM*, 2010, pp. 1–9.

[13] Z. Gong and M. Haenggi, "The local delay in mobile Poisson networks," *IEEE Trans. on Wireless Comm.*, vol. 12, no. 9, pp. 4766–4777, 2013.

[14] Y. Zhong, M. Haenggi, F.-C. Zheng, W. Zhang, T. Q. Quek, and W. Nie, "Towards a tractable delay analysis in large wireless networks," *arXiv preprint arXiv:1612.01276*, 2016.

[15] Y. Zhong, M. Haenggi, T. Q. Quek, and W. Zhang, "On the stability of static poisson networks under random access," *IEEE Trans. on Commun.*, vol. 64, no. 7, pp. 2985–2998, 2016.

[16] H. H. Yang and T. Q. Quek, "SIR coverage analysis in cellular networks with temporal traffic: a stochastic geometry approach," *arXiv preprint arXiv:1801.09888*, 2018.

[17] K. Stamatiou and M. Haenggi, "Random-access Poisson networks: stability and delay," *IEEE Comm. Letters*, vol. 14, no. 11, pp. 1035–1037, 2010.

[18] A. Sankararaman and F. Baccelli, "Spatial Birth-Death Wireless Networks," *IEEE Trans. on Info. Theory*, vol. 63, no. 6, pp. 3964–3982, 2017.

[19] A. AlAmmouri, J. G. Andrews, and F. Baccelli, "Stability and Metastability of Traffic Dynamics in Uplink Random Access Networks," *arXiv preprint arXiv:1906.04683*, 2019.

[20] F. Baccelli and P. Brémaud, *Elements of queueing theory: Palm Martingale calculus and stochastic recurrences*. Springer Science & Business Media, 2013, vol. 26.

[21] R. C. Bradley, "Basic properties of strong mixing conditions. a survey and some open questions," *Probability Surveys*, vol. 2, pp. 107–144, 2005.

[22] R. M. Loynes, "The stability of a queue with non-independent inter-arrival and service times," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 58, no. 3. Cambridge University Press, 1962, pp. 497–520.

[23] F. Baccelli and B. Błaszczyszyn, *Stochastic Geometry and Wireless Networks, Volume II - Applications, volume 4, No 1–2 of Foundations and Trends in Networking*. NoW Publishers, 2009.

[24] A. K. Gupta and M. A. Aziz, "Convex ordering of random variables and its applications in econometrics and actuarial science," *European Journal of Pure and Applied Mathematics*, vol. 3, no. 5, pp. 779–785, 2010.

[25] S. Asmussen and P. W. Glynn, *Stochastic simulation: algorithms and analysis*. Springer Science & Business Media, 2007, vol. 57.

[26] C.-S. Choi and F. Baccelli, "Poisson cox point processes for vehicular networks," *IEEE Trans. on Vehicular Technology*, vol. 67, no. 10, pp. 10 160–10 165, 2018.

[27] A. AlAmmouri, J. G. Andrews, and F. Baccelli, "A unified asymptotic analysis of area spectral efficiency in ultradense cellular networks," *IEEE Trans. on Info. Theory*, vol. 65, no. 2, pp. 1236–1248, 2018.

[28] F. Baccelli, B. Blaszczyszyn, and P. Muhlethaler, "An Aloha protocol for multihop mobile wireless networks," *IEEE Trans. on Info. Theory*, vol. 52, no. 2, pp. 421–436, 2006.

APPENDIX

A. Proof of Theorem V.1

We consider level-crossings spaced unit time apart for ease of exposition, but the result will hold for any interval of time. That $\mathbb{P}(L_{t+1}) = \mathbb{P}(L_t)$ is clear from the stationarity of $\text{SINR}_0(t)$. Let h_1 and h_2 be independent fading variables. The results in this section hold for any motion that preserves the homogeneity of Φ_t. We have:

$$\begin{aligned}\mathbb{P}[L_{t+1}, L_t] &= \mathbb{P}[l(R)h_1 > TI(t), l(R)h_2 > TI(t+1)] \\ &= \mathbb{E}[\exp(-\mu TI(t)/l(R)) \exp(-\mu TI(t+1)/l(R))] \\ &= \mathcal{L}_{t,t+1}(s,s),\end{aligned}$$

where $\mathcal{L}_{t,t+1}(.,.)$ is the joint Laplace transform of $I(t)$ and $I(t+1)$, and $s = \frac{\mu T}{l(R)}$. We assume that the motion that transforms Φ_t to Φ_{t+1} can be represented by a probability kernel p, where $p(x, B)$ is the probability that a point at location x at time t moves to region B at time $t+1$. Denote by $\mathcal{L}_h$ the Laplace transform of the general fading variables h that govern the channel between inteferers and the receiver of interest.

$$\mathcal{L}_{t,t+1}(s_1, s_2) = \mathbb{E}\left[\prod_{X_i \in \Phi_t} \mathcal{L}_h(s_1 l(X_i)) \prod_{Y_j \in \Phi_{t+1}} \mathcal{L}_h(s_2 l(Y_j))\right]$$

$$= \mathbb{E}\left[\prod_{X_i \in \Phi_t} \mathcal{L}_h(s_1 l(X_i)) \int \mathcal{L}_h(s_2 l(y_i)) p(X_i, dy_i)\right]$$

$$= \exp\left(-\Lambda \int (1 - v(x)) dx\right),$$

where $v(x) = \mathcal{L}_h(s_1 l(x)) \int \mathcal{L}_h(s_2 l(y)) p(x, dy)$

$$= \mathcal{L}_h(s_1 l(x)) + \mathcal{L}_h(s_1 l(x)) \left[\int \left(\mathcal{L}_h(s_2 l(y)) - 1\right) p(x, dy)\right].$$

Setting $s_1 = s_2 = s = \frac{\mu T}{l(R)}$, we have that $\mathcal{L}_{t,t+1}(s,s) = \mathbb{P}(L_t, L_{t+1})$.

Noting also that $\mathbb{P}(L_t) = \exp\left(-\Lambda \int 1 - \mathcal{L}_h(sl(x)) dx\right)$,

$$\begin{aligned}\mathbb{P}(L_{t+1}|L_t) &= \frac{\mathbb{P}(L_t, L_{t+1})}{\mathbb{P}(L_t)} = \frac{\mathcal{L}_{t,t+1}(s,s)}{\mathbb{P}(L_t)} \\ &= \exp\left(\Lambda \int\int \mathcal{L}_h(sl(x)) \mathcal{L}_h(sl(y)) p(x, dy) dx\right) \\ &\quad \times \exp\left(-\Lambda \int \mathcal{L}_h(sl(x)) dx\right).\end{aligned}$$

$$\begin{aligned}\implies \frac{\mathbb{P}(L_{t+1}|L_t)}{\mathbb{P}(L_t)} &= \exp\Bigg(\Lambda \int \Big[1 - 2\mathcal{L}_h(sl(x)) \\ &\quad + \int \mathcal{L}_h(sl(x)) \mathcal{L}_h(sl(y)) p(x, dy)\Big] dx\Bigg). \qquad (7)\end{aligned}$$

$\mathbb{P}(L_{t+1}) = \mathbb{P}(L_t)$, which yields $\int \left(\int \mathcal{L}_h(sl(y)) p(x, dy)\right) dx = \int \mathcal{L}_h(sl(x)) dx$.

Substituting this in (7),

$$\begin{aligned}\frac{\mathbb{P}(L_{t+1}|L_t)}{\mathbb{P}(L_t)} &= \exp\Bigg(\Lambda \int [1 - \mathcal{L}_h(sl(x))] \\ &\quad \times \left[1 - \left(\int \mathcal{L}_h(sl(y)) p(x, dy)\right)\right] dx\Bigg) > 1. \qquad \square\end{aligned}$$