

ON THE MODIFIED PALM VERSION

By HERMANN THORISSON

Abstract

The interpretation of the ‘standard’ Palm version of a stationary random measure ξ is that it behaves like ξ conditioned on containing the origin in its mass. The interpretation of the ‘modified’ Palm version is that it behaves like ξ seen from a typical location in its mass. In this paper we shall focus on the modified Palm version, comparing it with the standard version in the transparent case of mixed biased coin tosses, and then establishing a limit theorem that motivates the above interpretation in the case of random measures on locally compact second countable Abelian groups possessing Følner averaging sets.

Keywords: Random measure; Palm version; shift-coupling; mass-stationarity

2010 Mathematics Subject Classification: Primary 60G57; 60G55
Secondary 60G60

1. Introduction

The Palm version of a *stationary random measure* ξ is an important tool in probability. It is however not well known that there are in fact two Palm versions, with related but different interpretations. For want of other terminology, let us here call the well-known version *standard* and the less-known version *modified*. Denote them ξ^s and ξ^m , respectively.

The informal interpretation of the standard Palm version ξ^s is that it behaves like the stationary ξ conditioned on the origin lying in the mass of ξ (or, in the case of a simple point process, conditioned on there being a point of ξ at the origin).

The informal interpretation of the modified Palm version ξ^m is that it behaves like the stationary ξ seen from a typical location in the mass of ξ (or, in the case of a simple point process, seen from a typical point of ξ). Now, this is also a common interpretation of the standard Palm version ξ^s , but such an interpretation is flawed unless ξ^s and ξ^m have the same distribution, which happens for instance in the ergodic case (see Section 4).

In this paper we focus on the modified Palm version and its interpretation, using coin tosses as a transparent example. The concepts of shift-coupling and mass-stationarity play key roles in proving the main result, Theorem 7.1. That result motivates the above interpretation of ξ^m in the case of random measures on locally compact second countable Abelian groups possessing *Følner averaging sets* (see Section 3), for instance on \mathbb{R}^d .

The plan of the paper is as follows. In Section 2 we compare the two Palm versions in the case of mixed (possibly) biased coin tosses. Section 3 sets up the general framework for the formal definition of the two Palm versions in Section 4. Shift-coupling is considered in Section 5 and mass-stationarity in Section 6. In Section 7 we prove the main result and Section 8 concludes with some historical remarks.

2. Example: coin tosses (possibly biased)

For each $0 < p < 1$, let $(\xi_p(i) : i \in \mathbb{Z})$ be a doubly infinite sequence of independent and identically distributed (i.i.d.) Bernoulli p random variables representing tosses of a coin having probability p of coming up with a head. Let ξ_p be the associated Bernoulli random measure on the integers \mathbb{Z} defined by

$$\xi_p(B) = \sum_{i \in B} \xi_p(i), \quad B \subseteq \mathbb{Z}.$$

Let η_p denote the restriction of ξ_p to the nonzero integers $\mathbb{Z} \setminus \{0\}$, that is,

$$\eta_p(B) = \sum_{i \in B} \xi_p(i), \quad B \subseteq \mathbb{Z} \setminus \{0\}.$$

Conditioning on a head at the origin, i.e. $\xi_p(0) = 1$, does not affect the other coin tosses; they remain i.i.d. with probability p of coming up with a head. Thus, the standard Palm version ξ_p^s of ξ_p has a head at the origin, $\xi_p^s(0) = 1$, and the restriction of ξ_p^s to $\mathbb{Z} \setminus \{0\}$ behaves like η_p .

In order to deduce the structure of the modified Palm version ξ_p^m of ξ_p , note that the inter-head distances of ξ_p are independent and all geometric p except the distance between the last head before the origin 0 and the first head at or after the origin (that distance is the sum of two independent geometric p variables minus 1, this is the waiting time paradox). Thus, if we shift the origin of ξ_p to a typical head (a head chosen uniformly at random among all the infinitely many heads so the exceptional interval disappears to infinity in the limit, stretch your imagination) then we should obtain a random measure with a head at the (new) origin and inter-head distances that are all i.i.d. geometric p . This is exactly the distributional structure of the standard Palm version ξ_p^s . Thus, the two Palm versions ξ_p^s and ξ_p^m have the same distribution.

Now take $q \neq p$, and let M be a (nontrivial) ‘mixing’ variable that takes only the values p and q , and is independent of ξ_p and ξ_q (and of their Palm versions). Then the mixed Bernoulli random measure ξ_M is still stationary, but the coin tosses are no longer independent.

In this mixed setting, conditioning on a head at the origin, $\xi_M(0) = 1$, affects the other coin tosses η_M . Note that $\mathbb{P}(\xi_M(0) = 1) = p\mathbb{P}(M = p) + q\mathbb{P}(M = q)$, that is,

$$\mathbb{P}(\xi_M(0) = 1) = \mathbb{E}[M]. \tag{2.1}$$

Note also that, by independence,

$$\begin{aligned} \mathbb{P}(\eta_M \in \cdot, \xi_M(0) = 1, M = x) &= \mathbb{P}(\eta_x \in \cdot, \xi_x(0) = 1, M = x) \\ &= \mathbb{P}(\eta_x \in \cdot) x \mathbb{P}(M = x), \quad x = p, q. \end{aligned}$$

Divide by (2.1) to obtain

$$\mathbb{P}(\eta_M \in \cdot, M = x \mid \xi_M(0) = 1) = \mathbb{P}(\eta_x \in \cdot) \frac{x\mathbb{P}(M = x)}{\mathbb{E}[M]}, \quad x = p, q.$$

Thus, the standard Palm version $(\xi_M)^s$ of ξ_M behaves as ξ_p^s and ξ_q^s mixed by a *size-biased* version of M ,

$$\mathbb{P}(M = x \mid \xi_M(0) = 1) = \frac{x\mathbb{P}(M = x)}{\mathbb{E}[M]}, \quad x = p, q.$$

On the other hand, if we consider ξ_M as seen from a typical head then, conditionally on $M = x$ (for $x = p, q$), we would see i.i.d. geometric x inter-point distances. This is exactly the distributional structure of the modified Palm version ξ_x^m of ξ_x (which also happens to be the standard Palm version ξ_x^s of ξ_x). Thus, the modified Palm version $(\xi_M)^m$ of ξ_M behaves as ξ_p^m and ξ_q^m mixed by the *original* M ,

$$(\xi_M)^m = \xi_M^m.$$

Thus, the two Palm versions $(\xi_M)^s$ and $(\xi_M)^m$ do *not* have the same distribution.

Remark 2.1. The random measure ξ_p has mass points of size 0 and 1. Let us briefly consider random measures ξ_μ on \mathbb{Z} with mass points of i.i.d. random size $\xi_\mu(i)$ having distribution μ . Let the intensities $c_\mu = \mathbb{E}[\xi_\mu(i)]$ be positive and finite (if μ is Bernoulli p then $c_\mu = p$). Again, $\xi_\mu^s = \xi_\mu^m$; in both cases the restriction to $\mathbb{Z} \setminus \{0\}$ is as the restriction of ξ_μ to $\mathbb{Z} \setminus \{0\}$, and the distribution of $\xi_\mu^s(0) = \xi_\mu^m(0)$ is now a size-biased version of μ ,

$$\mathbb{P}(\xi_\mu^s(0) \in dy) = \mathbb{P}(\xi_\mu^m(0) \in dy) = \frac{y\mu(dy)}{\mathbb{E}[\xi_\mu(0)]}.$$

This size biasing is easy to guess from the interpretation of the modified version: placing a new origin at a typical location in the mass of ξ_μ suggests that the new location should be chosen with probability proportional to the mass y at that location. The results of the size biasing in three particular cases are as follows. If $\xi_\mu(0)$ is Poisson then $\xi_\mu^s(0) = \xi_\mu^m(0)$ is distributed as $\xi_\mu(0) + 1$. If $\xi_\mu(0)$ is exponential then $\xi_\mu^s(0) = \xi_\mu^m(0)$ is the sum of two independent copies of $\xi_\mu(0)$. If $\xi_\mu(0)$ is geometric minus 1 then $\xi_\mu^s(0) = \xi_\mu^m(0)$ is the sum of two independent copies of $\xi_\mu(0)$ plus 1.

Now take $\nu \neq \mu$, and let M be a (nontrivial) mixing variable that takes only the values μ and ν , and is independent of ξ_μ and ξ_ν (and of their Palm versions). As before, although $\xi_\mu^s = \xi_\mu^m$, the two Palm versions $(\xi_M)^s$ and $(\xi_M)^m$ do *not* have the same distribution: the standard Palm version $(\xi_M)^s$ turns out (see Section 4) to behave as ξ_μ^s and ξ_ν^s mixed by an *intensity-biased* version of M ,

$$\mathbb{P}((\xi_M)^s \in \cdot) = \mathbb{P}(\xi_\mu^s \in \cdot) \frac{c_\mu \mathbb{P}(M = \mu)}{\mathbb{E}[\xi_M(0)]} + \mathbb{P}(\xi_\nu^s \in \cdot) \frac{c_\nu \mathbb{P}(M = \nu)}{\mathbb{E}[\xi_M(0)]},$$

while the modified Palm version behaves as ξ_μ^m and ξ_ν^m mixed by the *original* M .

3. The stationary setting

The discussion in Section 2 is based on the informal interpretation of the two Palm versions. We now set up the framework for their formal definition in Section 4. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the probability space supporting the random elements below.

Let G be a locally compact second countable topological Abelian group with Borel sets \mathcal{G} and Haar measure λ (in Section 2, $G = \mathbb{Z}$ under addition and $\mathcal{G} = 2\mathbb{Z}$). For $t, s \in G$, denote t applied to s by ts , and denote the group inverse of t by t^{-1} . In Corollary 5.1 and Theorem 7.1 below we further assume the existence of *Følner sets*, that is, sets $B_r \in \mathcal{G}$, $r > 0$, of finite positive λ -measure such that (with ‘ Δ ’ denoting symmetric difference), for $r \rightarrow \infty$,

$$B_r \nearrow G \quad \text{and} \quad \frac{\lambda(B_r \Delta t^{-1} B_r)}{\lambda(B_r)} \rightarrow 0 \quad \text{for all } t \in G.$$

For example, if $G = \mathbb{R}^d$ (under addition, with λ the Lebesgue measure) and B is a convex set of positive finite volume containing 0 in its interior, then $B_r = rB$, $r > 0$, are Følner sets.

Let ξ be a random measure on (G, \mathcal{G}) . In association with ξ we consider a random element X in a space (H, \mathcal{H}) on which the group G acts measurably. This X could be another random measure on (G, \mathcal{G}) or a random field $X = (X_s)_{s \in G}$ indexed by G . Write θ_t for the shift map placing a new origin at $t \in G$. For example, $\theta_t \xi$ is defined by

$$\theta_t \xi(B) = \xi(tB), \quad B \in \mathcal{G},$$

while, if $X = (X_s)_{s \in G}$ is a random field, the definition is

$$\theta_t X = (X_{ts})_{s \in G}.$$

Also, the joint shift is defined by

$$\theta_t(X, \xi) = (\theta_t X, \theta_t \xi).$$

Let (X, ξ) be *stationary*, that is, with ‘ $\stackrel{D}{=}$ ’ denoting identity in distribution,

$$\theta_t(X, \xi) \stackrel{D}{=} (X, \xi), \quad t \in G.$$

Let P be the *stationary distribution*

$$P(A) = \mathbb{P}((X, \xi) \in A), \quad A \in \mathcal{H} \otimes \mathcal{G}.$$

4. The Palm versions

We now give a formal definition of the two Palm versions in the setting of Section 3. The *Palm measure* P^ξ of ξ with respect to the stationary P is defined by

$$P^\xi(A) = \frac{\mathbb{E}[\int_B \mathbf{1}_{\{\theta_s(X, \xi) \in A\}} \xi(ds)]}{\lambda(B)}, \quad A \in \mathcal{H} \otimes \mathcal{G},$$

where $B \in \mathcal{G}$ and $0 < \lambda(B) < \infty$. Note that P^ξ does not depend on B because the measure $B \mapsto \mathbb{E}[\int_B \mathbf{1}_{\{\theta_s(X, \xi) \in A\}} \xi(ds)]$ (with A fixed) is shift-invariant, owing to the stationarity of (X, ξ) , and thus is proportional to λ .

The Palm measure P^ξ is not necessarily a probability measure because it need not have total mass 1. The total mass of P^ξ is the *intensity* c of ξ ,

$$c = \frac{\mathbb{E}[\xi(B)]}{\lambda(B)}.$$

If $c < \infty$, define the *standard Palm distribution* P^s by

$$P^s = \frac{P^\xi}{c},$$

and call a pair (X^s, ξ^s) with distribution P^s the *standard Palm version* of (X, ξ) .

Let \mathcal{J} be the *invariant σ -algebra* on $(H, \mathcal{H}) \otimes (G, \mathcal{G})$ defined by

$$\mathcal{J} = \{A \in \mathcal{H} \otimes \mathcal{G} : \theta_t^{-1} A = A, t \in G\}.$$

Let \mathcal{J} be the *invariant σ -algebra* of (X, ξ) , that is, the sub- σ -algebra of \mathcal{F} defined by

$$\mathcal{J} = (X, \xi)^{-1} \mathcal{J}.$$

Define the *sample intensity* C of ξ by

$$C = \frac{\mathbb{E}[\xi(B) \mid \mathcal{F}]}{\lambda(B)}.$$

If $C < \infty$ almost surely (a.s.), define the *modified Palm distribution* P^m by

$$P^m(A) = \frac{\mathbb{E}[\int_B \mathbf{1}_{\{\theta_s(X, \xi) \in A\}} \xi(ds) / C]}{\lambda(B)}, \quad A \in \mathcal{G} \otimes \mathcal{H}, \tag{4.1}$$

and call a pair (X^m, ξ^m) with distribution P^m the *modified Palm version* of (X, ξ) . Note that $\mathbb{E}[C] = c$ and that $P^s = P^m$ if and only if $C = c$ a.s.

For notational convenience, let both (X^s, ξ^s) and (X^m, ξ^m) be defined on $(\Omega, \mathcal{F}, \mathbb{P})$.

Remark 4.1. The Palm versions (X^s, ξ^s) and (X^m, ξ^m) can be constructed in a similar two-step (pathwise, measurewise) way as applied to random measures on \mathbb{R}^d in [9] (see also the notes in [18]), as follows.

Fix some $B \in \mathcal{G}$ such that $0 < \lambda(B) < \infty$. Let S_B be a *typical location* in the mass of ξ on the set B conditionally on (X, ξ) , that is, let S_B be a random element in G with conditional distribution $\xi(\cdot \mid B)$ given (X, ξ) ,

$$\mathbb{P}(S_B \in \cdot \mid (X, \xi)) = \xi(\cdot \mid B) := \frac{\xi(\cdot \cap B)}{\xi(B)} \quad (\text{any distribution when } \xi(B) = 0).$$

If $c < \infty$ then $\theta_{S_B}(X, \xi)$ has the standard Palm distribution P^s under the measure \mathbb{P}^s defined on (Ω, \mathcal{F}) by

$$d\mathbb{P}^s = \frac{\xi(B)}{c} d\mathbb{P}.$$

If $C < \infty$ \mathbb{P} -a.e. then $\theta_{S_B}(X, \xi)$ has the modified Palm distribution P^m under the measure \mathbb{P}^m defined on (Ω, \mathcal{F}) by

$$d\mathbb{P}^m = \frac{\xi(B)}{C} d\mathbb{P}.$$

In [9] the constructions could be reversed to obtain the stationary version from the Palm versions. Here the constructions can be reversed only when $\mathbb{P}(\xi(B) = 0) = 0$.

5. Shift-coupling

In this section we present a shift-coupling result for the stationary (X, ξ) and its modified Palm version (X^m, ξ^m) . The proof is based on theory from [16] and [18].

Let Y and Z be random elements in a measurable space (K, \mathcal{K}) on which the group G acts measurably. Say that Y and Z admit *shift-coupling* if there exists (possibly after extension of the underlying probability space) a random element T in G such that $\theta_T Y$ and Z have the same distribution, that is,

$$\theta_T Y \stackrel{D}{=} Z \quad (\text{shift-coupling}).$$

The following theorem is from [16] (and holds without the restriction that G is Abelian); see also Section 7.4 of [18, Chapter 7].

Theorem 5.1. *The random elements Y and Z admit shift-coupling if and only if their distributions agree on the invariant σ -algebra of (K, \mathcal{K}) .*

Let $\|\cdot\|$ denote the total variation norm and let ‘ Δ ’ denote symmetric difference. From $\theta_T Y \stackrel{D}{=} Z$ it follows (see [18, Chapter 7, Section 7.3]) that, for $B \in \mathcal{G}$ with $0 < \lambda(B) < \infty$,

$$\|\mathbb{P}(\theta_{U_B} Y \in \cdot) - \mathbb{P}(\theta_{U_B} Z \in \cdot)\| \leq \mathbb{E} \left[\frac{\lambda(B \Delta T^{-1} B)}{\lambda(B)} \right] \quad (\text{shift-coupling inequality}),$$

where U_B is distributed according to $\lambda(\cdot \mid B)$ and is independent of Y and Z (and T^{-1} denotes the group inverse of T and not the inverse of T as a function of $\omega \in \Omega$).

Furthermore, suppose that there exist Følner sets $B_r, r > 0$; see Section 3. Then the shift-coupling inequality yields (by bounded convergence)

$$\|\mathbb{P}(\theta_{U_{B_r}} Y \in \cdot) - \mathbb{P}(\theta_{U_{B_r}} Z \in \cdot)\| \rightarrow 0 \quad \text{as } r \rightarrow \infty.$$

In particular, if Z is stationary then $\theta_{U_{B_r}} Z \stackrel{D}{=} Z$ and this limit result can be written as

$$\theta_{U_{B_r}} Y \xrightarrow{\text{TV}} Z \quad \text{as } r \rightarrow \infty,$$

where ‘ $\xrightarrow{\text{TV}}$ ’ denotes convergence in total variation.

We are now at a key observation on our way towards motivating the interpretation of the modified Palm version. (Here the condition that G is Abelian is not needed.)

Proposition 5.1. *The distributions of the stationary (X, ξ) and its modified Palm version (X^m, ξ^m) agree on the invariant σ -algebra of $(H, \mathcal{H}) \otimes (G, \mathcal{G})$:*

$$P(A) = P^m(A), \quad A \in \mathcal{J}.$$

Proof. Take $A \in \mathcal{J}$. Then $\{\theta_s(X, \xi) \in A\} = \{(X, \xi) \in A\}$, and thus (see (4.1))

$$P^m(A) = \frac{\mathbb{E}[\mathbf{1}_{\{(X,\xi) \in A\}} \xi(B)/C]}{\lambda(B)}. \tag{5.1}$$

Both $\{(X, \xi) \in A\}$ and $C := \mathbb{E}[\xi(B) \mid \mathcal{J}]/\lambda(B)$ are measurable with respect to \mathcal{J} , and thus

$$\mathbb{E} \left[\frac{\mathbf{1}_{\{(X,\xi) \in A\}} \xi(B)}{C} \mid \mathcal{J} \right] = \frac{\mathbf{1}_{\{(X,\xi) \in A\}} \mathbb{E}[\xi(B) \mid \mathcal{J}]}{C} = \mathbf{1}_{\{(X,\xi) \in A\}} \lambda(B).$$

Take the expectation, divide by $\lambda(B)$, and consult (5.1) to obtain

$$P^m(A) = \mathbb{E}[\mathbf{1}_{\{(X,\xi) \in A\}}] = P(A). \quad \square$$

The above (with $Y = (X^m, \xi^m), Z = (X, \xi)$ and with $Y = (X, \xi), Z = (X^m, \xi^m)$) yields the following result.

Corollary 5.1. *There exist (possibly after extension of $(\Omega, \mathcal{F}, \mathbb{P})$) random elements T^m and T in G such that*

$$\theta_{T^m}(X^m, \xi^m) \stackrel{D}{=} (X, \xi) \quad \text{and} \quad \theta_T(X, \xi) \stackrel{D}{=} (X^m, \xi^m).$$

Furthermore, suppose that there exist Følner sets $B_r, r > 0$, and let U_{B_r} be a typical location in B_r and independent of (X^m, ξ^m) , that is,

$$\mathbb{P}(U_{B_r} \in \cdot \mid (X^m, \xi^m)) = \lambda(\cdot \mid B_r).$$

Then

$$\theta_{U_{B_r}}(X^m, \xi^m) \xrightarrow{\text{TV}} (X, \xi) \quad \text{as } r \rightarrow \infty.$$

6. Mass-stationarity

Corollary 5.1 can be interpreted as saying that the stationary (X, ξ) behaves as (X^m, ξ^m) seen from a typical location in the mass of λ . Our aim however is to establish the reverse interpretation, namely, that the modified Palm version (X^m, ξ^m) behaves as (X, ξ) seen from a typical location in the mass of ξ . For that purpose, we need mass-stationarity, a characterising property of Palm measures introduced in [7].

Let Y be a random element in (H, \mathcal{H}) , and let η be a random measure on G having the origin in its support. A set $B \in \mathcal{G}$ is called a λ -continuity set if its boundary ∂B has λ -mass zero, $\lambda(\partial B) = 0$. The pair (Y, η) is *mass-stationary* if, for all bounded λ -continuity sets $B \in \mathcal{G}$ of positive λ -measure, it holds that

$$(\theta_{V_B}(Y, \eta), U_B V_B) \stackrel{D}{=} ((Y, \eta), U_B),$$

where U_B and V_B are random elements in G such that U_B is a typical location in (the mass of λ on) the set B and independent of (Y, η) , that is,

$$\mathbb{P}(U_B \in \cdot \mid (Y, \eta)) = \lambda(\cdot \mid B),$$

and V_B is a typical location in the mass of η on the set $U_B^{-1}B$ given $((Y, \eta), U_B)$ (with U_B^{-1} the group inverse of U_B), that is,

$$\mathbb{P}(V_B \in \cdot \mid (Y, \eta), U_B) = \eta(\cdot \mid U_B^{-1}B) := \frac{\eta(\cdot \cap U_B^{-1}B)}{\eta(U_B^{-1}B)}.$$

We still say that (Y, η) is mass-stationary when the above holds even if the distribution Q of (Y, η) is only σ -finite.

The following theorem was established in [7] for Abelian G ; the restriction to Abelian G was removed in [6].

Theorem 6.1. *A pair (Y, η) is mass-stationary if and only if its (possibly only σ -finite) distribution Q is the Palm measure P^ξ of the (possibly only σ -finite) distribution P of some stationary (X, ξ) .*

Note that, although the distribution P of (X, ξ) and the Palm measure P^ξ can both be only σ -finite, the distributions P^s and P^m of the two Palm versions are always probability measures (when they exist).

Corollary 6.1. *Both Palm versions (X^s, ξ^s) and (X^m, ξ^m) are mass-stationary.*

7. Interpretation of the modified Palm version

We are now ready for the limit theorem motivating the interpretation of the modified Palm version in the case that the Abelian group G possesses Følner sets, namely,

$$(X^m, \xi^m) \text{ behaves like } (X, \xi) \text{ seen from a typical location in the mass of } \xi.$$

In this result we demand (as throughout the paper except in Theorem 6.1) that the distribution P of the stationary (X, ξ) has mass 1.

Theorem 7.1. *Let (X, ξ) be stationary with sample intensity $C < \infty$ a.s. Let (X^m, ξ^m) be the modified Palm version of (X, ξ) . Suppose that G is Abelian and that there exist bounded λ -continuity Følner sets $B_r, r > 0$. Then*

$$\theta_{S_{B_r}}(X, \xi) \xrightarrow{\text{TV}} (X^m, \xi^m) \text{ as } r \rightarrow \infty,$$

where S_{B_r} is a typical location in the mass of ξ on the set B_r given (X, ξ) , that is,

$$\mathbb{P}(S_{B_r} \in \cdot \mid (X, \xi)) = \xi(\cdot \mid B_r), \quad r > 0.$$

Proof. For convenience, write X for (X, ξ) and X^m for (X^m, ξ^m) . Let U_{B_r} and V_{B_r} be as in the definition of mass-stationarity, that is,

$$\mathbb{P}(U_{B_r} \in \cdot \mid X^m) = \lambda(\cdot \mid B_r)$$

and

$$\mathbb{P}(V_{B_r} \in \cdot \mid X^m, U_{B_r}) = \xi^m(\cdot \mid U_{B_r}^{-1} B_r), \tag{7.1}$$

where $U_{B_r}^{-1}$ is the group inverse of U_{B_r} . By the mass-stationarity of X^m (Corollary 6.1),

$$\|\mathbb{P}(\theta_{S_{B_r}} X \in \cdot) - \mathbb{P}(X^m \in \cdot)\| = \|\mathbb{P}(\theta_{S_{B_r}} X \in \cdot) - \mathbb{P}(\theta_{V_{B_r}} X^m \in \cdot)\|.$$

Since G is Abelian, this can be rewritten as

$$\|\mathbb{P}(\theta_{S_{B_r}} X \in \cdot) - \mathbb{P}(X^m \in \cdot)\| = \|\mathbb{P}(\theta_{S_{B_r}} X \in \cdot) - \mathbb{P}(\theta_{U_{B_r} V_{B_r}} \theta_{U_{B_r}^{-1}} X^m \in \cdot)\|.$$

Since $\theta_{S_{B_r}} X$ is obtained from (X, S_{B_r}) by the same measurable map as $\theta_{U_{B_r} V_{B_r}} \theta_{U_{B_r}^{-1}} X^m$ from $(\theta_{U_{B_r}^{-1}} X^m, U_{B_r} V_{B_r})$, we deduce from this that

$$\|\mathbb{P}(\theta_{S_{B_r}} X \in \cdot) - \mathbb{P}(X^m \in \cdot)\| \leq \|\mathbb{P}((X, S_{B_r}) \in \cdot) - \mathbb{P}((\theta_{U_{B_r}^{-1}} X^m, U_{B_r} V_{B_r}) \in \cdot)\|. \tag{7.2}$$

Now (7.1) implies that $\mathbb{P}(U_{B_r} V_{B_r} \in \cdot \mid X^m, U_{B_r}) = (\theta_{U_{B_r}^{-1}} \xi^m)(\cdot \mid B_r)$, and thus

$$\mathbb{P}(U_{B_r} V_{B_r} \in \cdot \mid \theta_{U_{B_r}^{-1}} X^m) = (\theta_{U_{B_r}^{-1}} \xi^m)(\cdot \mid B_r).$$

Thus, the conditional distribution of $U_{B_r} V_{B_r}$ given the value of $\theta_{U_{B_r}^{-1}} X^m$ is the same as that of S_{B_r} given the value of X , that is,

$$\mathbb{P}(U_{B_r} V_{B_r} \in \cdot \mid \theta_{U_{B_r}^{-1}} X^m = \cdot) = \mathbb{P}(S_{B_r} \in \cdot \mid X = \cdot).$$

This implies (see Lemma 3.1 of [18, Chapter 6]) that

$$\|\mathbb{P}((X, S_{B_r}) \in \cdot) - \mathbb{P}((\theta_{U_{B_r}^{-1}} X^m, U_{B_r} V_{B_r}) \in \cdot)\| = \|\mathbb{P}(X \in \cdot) - \mathbb{P}(\theta_{U_{B_r}^{-1}} X^m \in \cdot)\|. \tag{7.3}$$

Since G is Abelian, the sets $B_r^{-1}, r > 0$, are λ -continuity Følner sets and

$$\mathbb{P}(U_{B_r}^{-1} \in \cdot \mid X^m) = \lambda(\cdot \mid B_r^{-1}).$$

Thus, Corollary 5.1 yields $\|\mathbb{P}(X \in \cdot) - \mathbb{P}(\theta_{U_{B_r}^{-1}} X^m \in \cdot)\| \rightarrow 0$ as $r \rightarrow \infty$, and combining this with (7.2) and (7.3) yields the desired result:

$$\|\mathbb{P}(\theta_{S_{B_r}} X \in \cdot) - \mathbb{P}(X^m \in \cdot)\| \rightarrow 0 \text{ as } r \rightarrow \infty. \quad \square$$

8. Notes

When Conny Palm [13] introduced his theory in 1943, he was focussed on the randomised-origin interpretation of what we now call Palm versions (see his Chapter 2). Thus, it is a pity to use the term ‘modified’ for the Palm version with this interpretation. Also, strictly speaking, the ‘modified’ Palm version is not really a modification of the ‘standard’ version. It can exist even when the ‘standard’ version does not; we can have $C < \infty$ a.s. when $c = \mathbb{E}[C] = \infty$. But the ‘modified’ Palm distribution is a modification of the (formally simpler and more general) Palm measure. So, what can be said?

The relationship between the two Palm versions and the role of the sample intensity C is implicit in the 1962 paper [14] by Slivnyak, dealing with point processes on the line; see also the 1978 book [11] by Matthes *et al.* In the early 1990s the modified Palm version and its interpretation was rediscovered independently by Glynn and Sigman [3], Nieuwenhuis [12], and this author [15], in the context of point processes on the line. In [17] the view was extended to point processes on \mathbb{R}^d .

The path of this author to the modified Palm version was through work on shift-coupling. Shift-coupling, in the one-sided context of nonnegative integers, dates back to the amazing 1979 monograph by Berbee [2], where the link to the Cesaro total variation convergence (such as in Corollary 5.1) was established. The term ‘shift-coupling’ was coined in the 1993 paper [1], where the link to the invariant σ -algebra was established. The extension to the group setting of the present paper is from the 1996 paper [16], where the shift-coupling inequality can also be found (embedded in the proof of Theorem 2). See [18] for more background.

Mass-stationarity of random measures, an intrinsic characterisation of Palm measures, was introduced in the 2009 paper [7]. It is a generalisation of point-stationarity of simple point processes which was introduced in [17] (the necessity part of the point-stationarity characterisation is in fact implicit in Satz 4.3 of the 1975 paper [10]). Point-stationarity is in turn a generalisation of cycle-stationarity, the well-known characterising property of Palm measures of simple point processes on the line that the inter-point distances (and the paths between the points) form a stationary sequence. See [8] for a survey of developments concerning point- and mass-stationarity and for some clarification in the compact-group case of the (mysterious) double-randomisation property (see Section 6) used to define mass-stationarity.

The proof above of Theorem 7.1 is an example of how mass-stationarity works. It is adapted from the proof of Theorem 9.3(b) of [17] (the double-randomisation property had been established in Theorem 5.1 in that paper as a characterisation of point-stationarity, and therefore of Palm versions, in the case of simple point processes on \mathbb{R}^d).

For a comprehensive treatment of random measures and Palm theory, see Olav Kallenberg’s new book [5]. See also his *Foundations of Modern Probability* [4].

References

- [1] ALDOUS, D. J. AND THORISSON, H. (1993). Shift-coupling. *Stoch. Process. Appl.* **44**, 1–14.
- [2] BERBEE, H. C. P. (1979). *Random Walks with Stationary Increments and Renewal Theory* (Math. Centre Tracts **112**). Mathematisch Centrum, Amsterdam.
- [3] GLYNN, P. AND SIGMAN, K. (1992). Uniform Cesàro limit theorems for synchronous processes with applications to queues. *Stoch. Process. Appl.* **40**, 29–43.
- [4] KALLENBERG, O. (2002). *Foundations of Modern Probability*, 2nd edn. Springer, New York.
- [5] KALLENBERG, O. (2017). *Random Measures, Theory and Applications*. Springer, Cham.
- [6] LAST, G. (2010). Modern random measures: Palm theory and related models. In *New Perspectives in Stochastic Geometry*, eds W. Kendall and I. Molchanov, Oxford University Press, pp. 77–110.
- [7] LAST, G. AND THORISSON, H. (2009). Invariant transports of stationary random measures and mass-stationarity. *Ann. Prob.* **37**, 790–813.

- [8] LAST, G. AND THORISSON, H. (2011). What is typical? In *New Frontiers in Applied Probability* (Adv. Appl. Prob. Spec. Vol. **48A**), eds C. M. Goldie and A. Mijatović, Applied Probability Trust, Sheffield, pp. 379–389.
- [9] LAST, G. AND THORISSON, H. (2015). Construction and characterisation of stationary and mass-stationary random measures on \mathbb{R}^d . *Stoch. Process. Appl.* **125**, 4473–4488.
- [10] MECKE, J. (1975). Invarianzeigenschaften allgemeiner Palmischer Maße. *Math. Nachr.* **65**, 335–344.
- [11] MATTHES, K., KERSTAN, J. AND MECKE, J. (1978). *Infinitely Divisible Point Processes*. John Wiley, Chichester.
- [12] NIEUWENHUIS, G. (1994). Bridging the gap between a stationary point process and its Palm distribution. *Statist. Neerlandica* **48**, 37–62.
- [13] PALM, C. (1943). *Intensitätsschwankungen im Fernsprecherkehr* (Ericsson Techniks **44**). Ericsson, Stockholm (in German). English translation: *Intensity Variations in Telephone Traffic*, North-Holland, Amsterdam, 1988.
- [14] SLINYAK, I. M. (1962). Some properties of stationary flows of homogeneous random events. *Theory Prob. Appl.* **7**, 336–341.
- [15] THORISSON, H. (1995). On time- and cycle-stationarity. *Stoch. Process. Appl.* **55**, 183–209.
- [16] THORISSON, H. (1996). Transforming random elements and shifting random fields. *Ann. Prob.* **24**, 2057–2064.
- [17] THORISSON, H. (1999). Point-stationarity in d dimensions and Palm theory. *Bernoulli* **5**, 797–831.
- [18] THORISSON, H. (2000). *Coupling, Stationarity, and Regeneration*. Springer, New York.

HERMANN THORISSON, *University of Iceland*

Science Institute, University of Iceland, Dunhaga 5, 107 Reykjavik, Iceland. Email address: hermann@hi.is