Adv. Appl. Prob. (SGSA) **35,** 56–69 (2003)
Printed in Northern Ireland

DECOMPOSITION OF GAMMA-DISTRIBUTED DOMAINS CONSTRUCTED FROM POISSON POINT PROCESSES

RICHARD COWAN * AND
MALCOLM QUINE,* *University of Sydney*
SERGEI ZUYEV,** *University of Strathclyde*

Abstract

A known gamma-type result for the Poisson process states that certain domains defined through configuration of the points (or 'particles') of the process have volumes which are gamma distributed. By proving the corresponding sequential gamma-type result, we show that in some cases such a domain allows for decomposition into subdomains each having independent exponentially distributed volumes. We consider other examples—based on the Voronoi and Delaunay tessellations—where a natural decomposition does not produce subdomains with exponentially distributed volumes. A simple algorithm for the construction of a typical Voronoi flower arises in this work. In our theoretical development, we generalize the classical theorem of Slivnyak, relating it to the strong Markov property of the Poisson process and to a result of Mecke and Muche (1995). This new theorem has interest beyond the specific problems being considered here.

Keywords: Stopping set; gamma-type results; Slivnyak theorem; set-indexed martingale; Voronoi tessellation; Delaunay triangulation; Poisson process

AMS 2000 Subject Classification: Primary 60G48; 60D05
Secondary 62M30; 60G55

1. Introduction

Many figures constructed with respect to points of a homogeneous point process have a volume which is gamma distributed. We commence with a simple example.

Example 1. (*Volumes of annuli.*) Consider a homogeneous Poisson point (particle) process with intensity λ in $\mathbb{R}^d$ and let $x_1, x_2, \dots$ denote the first, second, $\dots$ closest particles to the origin. It is easily verified that the volume of the random ball B_n centred at the origin with radius $\|x_n\|$ has a gamma $\Gamma(n, \lambda)$ distribution in any dimension. Since this distribution coincides with the distribution of the sum of n independent exponential $\operatorname{Exp}(\lambda)$ variables with parameter λ, it is natural to ask: *can we decompose the ball into a disjoint union of domains with independent exponentially distributed volumes*? In this particular case, the answer is easy: yes, these domains are B_1 and the sequence of rings $B_k \setminus B_{k-1},\ k = 2, \dots, n$.

Situations where a random domain constructed with respect to Poisson particles has a gamma-distributed volume can be far more complex than Example 1. The first general result of this kind is the 'complementary theorem' of Miles (1970), developed further in a Palm-measure context by Møller and Zuyev (1996) and Zuyev (1999). Roughly speaking, any 'scale-invariant'

Received 6 August 2002; revision received 28 October 2002.
* Postal address: School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia.
** Postal address: Department of Statistics and Modelling Science, University of Strathclyde, Glasgow G1 1XH, UK. Email address: sergei@stams.strath.ac.uk

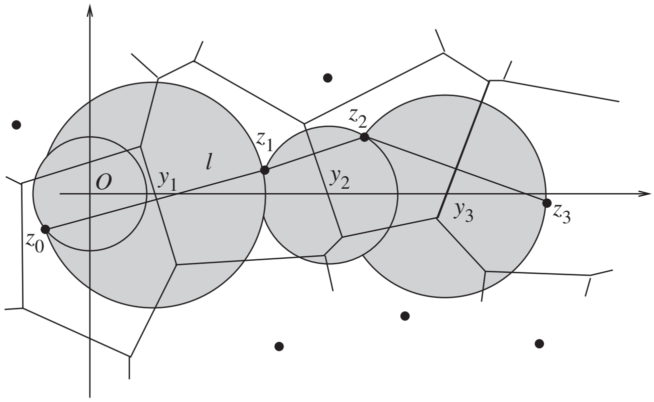

FIGURE 1: Voronoi sausage in $\mathbb{R}^2$.

random domain constructed with respect to $n+1$ Poisson particles and containing $m \geq 0$ other particles (or 'm-filled' in Miles' terminology) has $\Gamma(n+m, \lambda)$-distributed volume, as has an m-filled domain constructed from n particles and a given reference point such as the origin O *or* a domain constructed from $n+1$ particles containing $m-1$ other particles *and* also known to include (or 'cover') the given reference point.

The aim of this paper is to address the above decomposition question for such domains. The following are nontrivial examples of domains which conform to the complementary theorem. All use homogeneous Poisson particle processes of intensity λ.

Example 2. (*Voronoi sausage.*) Consider the Voronoi tessellation in $\mathbb{R}^d$. Denote by y_k the points of intersection of the Voronoi cells' hyper-faces with the positive first semi-axis. For each such y_k, there is, by definition, a ball B_k centred at it which contains on its boundary exactly two particles, with no particle inside. A Voronoi sausage is a finite union of these consecutive balls and the ball B_0 centred at the origin O with z_0, the closest particle to O (and the centre of the polygon covering O), on its boundary (see Figure 1). Each added ball B_{k+1} has exactly one particle, call it z_k, in common with the previous ball B_k and one 'new' particle z_{k+1}. If we connect these particles, we obtain a path l on a Delaunay graph which does not deviate 'very much' from the first axis. This corresponds to connecting the centres of the cells in the order the axis crosses them. The geometrical properties of this path are completely determined by the sausage. For instance, its average length in $\mathbb{R}^2$ is just $4/\pi \approx 1.273$ times larger than the Euclidean distance; see Baccelli *et al.* (2000) for details. A Voronoi sausage formed by n balls is constructed from the reference point O and n sequentially determined Poisson particles. Therefore, according to the complementary theorem, its area is $\Gamma(n, \lambda)$ distributed, the domain being empty of other particles. The area of B_0 conforms to $\operatorname{Exp}(\lambda)$, but are the areas of the lunes $B_k \setminus B_{k-1},\ k = 1, \ldots, n$, independent $\operatorname{Exp}(\lambda)$ random variables?

Example 3. (*Delaunay lunes.*) Consider a 'typical' Delaunay circumdisc in $\mathbb{R}^2$, that is, a disc which contains a 'typical' triple of Poisson particles on its boundary and is further constrained to have no other particle inside. The disc's centre is a vertex of the corresponding Voronoi tessellation in which three edges meet. According to the complementary theorem, its area has $\Gamma(2, \lambda)$ distribution. There have been no suggestions in the literature about how one might partition the disc into two $\operatorname{Exp}(\lambda)$ parts, but Figure 2 is immediately suggestive: choose

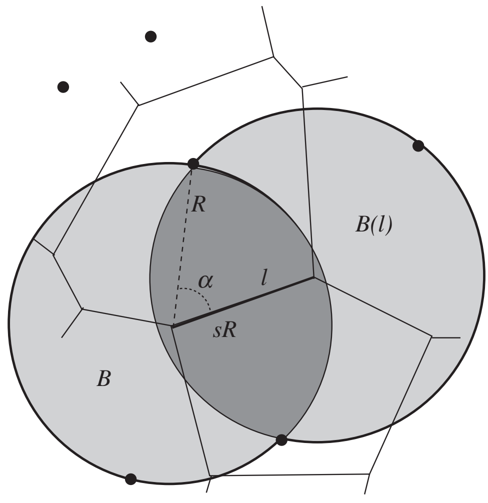

FIGURE 2: Two lunes and a lens corresponding to a Voronoi edge.

uniformly one of the three Voronoi edges emanating from the centre; call it l, and consider another Delaunay circumdisc at its other end. The two discs, say B and $B(l)$, define two 'lunes' $B \setminus B(l)$ and $B(l) \setminus B$ as well as one 'lens' $B \cap B(l)$. Note that the chosen edge l is a typical edge of the Voronoi tessellation, so the disc $B(l)$ is a typical Delaunay circumdisc too and hence its area is also $\Gamma(2, \lambda)$. It is therefore a natural hypothesis that both lunes and the lens all have $\text{Exp}(\lambda)$ distribution. We will address this hypothesis in Section 5.

Example 4. (*Voronoi flower or fundamental region.*) Consider a Voronoi polygon V in $\mathbb{R}^2$, centred at a typical particle labelled O. Its geometry is determined by the union of balls $B_1, \ldots, B_n$, n being the number of its sides, each ball centred at a vertex of V and having the particle O and exactly two other Poisson particles—three in all—on its boundary *and* no other particle inside (see Figure 3). We call this union the *Voronoi flower* F associated with the cell V. Conditioned on the number of sides n, the flower is constructed from $n + 1$ particles with no other particle inside; hence it has a $\Gamma(n, \lambda)$-distributed area. This result is first mentioned in Miles and Maillardet (1982), though without detailed proof; a proof is in Møller and Zuyev (1996) (see also Zuyev (1999)). Hayen and Quine (2002) consider a natural decomposition of F into n 'petals' formed by cutting F by the radii vectors of the flower's particles, and show that petal areas are neither exponentially distributed nor independent. Here we try another approach; we decompose F into the lunes $B_1 \setminus B_n, B_2 \setminus B_1, \ldots, B_n \setminus B_{n-1}$ and ask ourselves if their areas are $\text{Exp}(\lambda)$ distributed.

The trivial decomposition result in Example 1 requires little more than the definition of the Poisson process for its proof. The other examples, however, require more sophisticated machinery. Our technique, described briefly in Section 2, is based on *set-indexed martingales* developed for the case of point processes in Zuyev (1999); it uses the idea of a *stopping set*, extending the notion of a stopping time, well known for temporal stochastic processes.

In Section 3, we prove a generalization of the Slivnyak theorem, a result more appropriately called the 'Slivnyak–Mecke theorem' because Mecke (1967) extended Slivnyak's theorem for Poisson processes on the line to a general state space, also providing a characterization result

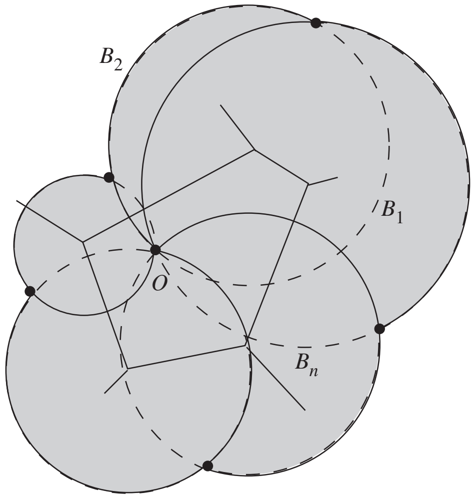

FIGURE 3: Voronoi flower in $\mathbb{R}^2$ decomposed into lunes.

for Poisson processes. Their theorems say, roughly, that the Poisson process Π has the same distribution when viewed from one of the Poisson particles as from a given reference point (provided that the viewing location is not part of the view). We prove a similar result in Section 3; when viewing from within a random compact domain S, any view of that part of Π which lies outside S is unaffected by the process structure in S, so long as this domain is a stopping set with respect to Π. A similar conditional independence was established in Miles (1970), and more formally in Mecke and Muche (1995), for the case when S is a typical Delaunay circumdisc with respect to Π.

In Section 4, we develop two *sequential* gamma-type theorems extending the corresponding result of Zuyev (1999); their application, presented in Section 5, proves the required decomposition in some cases.

Although these results are formulated for the case of Euclidean space $\mathbb{R}^d$, a generalization to a locally compact separable topological group (LCS-space) is straightforward.

Richard Cowan writes: I corresponded with Joseph Mecke long before I finally met him at one of the Oberwolfach meetings. These letters were warm and stimulating, and so, from early times more than 25 years ago, I have held him in high personal regard. Best wishes, Joseph.

Sergei Zuyev writes: I belong to the generation next to Professor Mecke's and surely have a deep respect for him. Although I have never had the chance to work with him, I also have a rather personal feeling towards Joseph as my teacher, as in almost every subject where my interests have led me I soon discovered that it was marked by his fundamental results, be it in stochastic geometry, point processes or Palm theory. I should have started by reading all his publications first. Long live dear Professor Mecke!

2. Preliminaries and notation

Let $X \subseteq \mathbb{R}^d$ be a subset of d-dimensional Euclidean space $\mathbb{R}^d$ with its Borel σ-algebra $\mathcal{B}$, $\mathbb{F}$ and $\mathbb{K}$ being the system of its closed and compact subsets respectively in the induced topology.

In this paper we deal with point processes on X that are treated by us as random counting measures. More specifically, a point process Φ is a measurable mapping from some abstract probability space into $(\mathcal{N}, \mathcal{F})$, where $\mathcal{N}$ is a set of all counting measures without accumulation points: for any $\phi \in \mathcal{N}$, we have $\phi(K) \in \mathbb{Z}_+$ for all $K \in \mathbb{K}$; and $\mathcal{F}$ is the σ-algebra generated by events $\{\phi \in \mathcal{N} : \phi(K) = k\}$ for $k \in \mathbb{Z}_+$, $K \in \mathbb{K}$. A point process Φ is *canonically defined* if the probability space is $(\mathcal{N}, \mathcal{F}, \mathrm{P})$ itself and Φ is the identity mapping. The process Φ is *simple* if, almost surely, $\Phi(\{x\})$ is 0 or 1 for any singleton $\{x\}$.

Together with the σ-algebra $\mathcal{F}$, we will also consider a family of σ-algebras

$$\mathcal{F}_K = \sigma(\{\phi \in \mathcal{N} : \phi(K' \cap K) = k\},\ k \in \mathbb{Z}_+, K' \in \mathbb{K})$$

indexed by compact sets $K \in \mathbb{K}$ which are generated by the truncated counting measures $\phi|_K(\cdot) = \phi(K \cap \cdot)$.

The following two properties of the ensemble $\{\mathcal{F}_K\}$ allow us to call it a *filtration*:

(i) monotonicity: $\mathcal{F}_{K_1} \subseteq \mathcal{F}_{K_2}$ for any two compact sets $K_1 \subseteq K_2$;

(ii) continuity from above: $\mathcal{F}_K = \bigcap_{n=1}^{\infty} \mathcal{F}_{K_n}$ if $K_n \downarrow K$.

We also put $\mathcal{F}_\varnothing$ to be the trivial σ-algebra. We will use the term *natural filtration* for this minimal filtration generated by simple counting measures.

When $X = \mathbb{R}^d$, any $x \in X$ generates a measurable bijection $\theta_x : X \to X$ defined by $\theta_x(y) = y + x$. We will also write

$$\theta_x B = B + x = \{y + x : y \in B\}.$$

The family $\theta_\cdot$, called a *flow*, gives rise to the compatible flow in $\mathcal{N}$: $\theta_x\phi$ is the measure taking value $(\theta_x\phi)(B) = \phi(\theta_x B)$ on a Borel set B. In particular, $\theta_x\delta_y = \delta_{y-x}$, where δ_y is the unit measure concentrated on $\{y\}$.

The probability measure P is *stationary* if $\mathrm{P}(\Sigma) = \mathrm{P}(\theta_x\Sigma)$ for all $\Sigma \in \mathcal{F}$ and $x \in X$. Any process Φ *compatible* with the flow, i.e. such that $\theta_x\Phi(\phi) = \Phi(\theta_x\phi)$ for all ϕ and x, is then automatically stationary in the sense that the distributions of Φ and $\theta_x\Phi$ coincide.

The *intensity measure* of a point process defined as $\lambda(B) = \mathrm{E}\Phi(B)$ in the stationary case is proportional to the Lebesgue measure: $\lambda_\Phi\,\mathrm{d}x$; the term λ_Φ is called the *intensity* of the point process.

The *Palm distribution* Q^Φ on $[\mathcal{N}, \mathcal{F}]$ corresponding to a stationary point process Φ is defined on any $\Sigma \in \mathcal{F}$ by means of

$$\lambda_\Phi \mathrm{Q}^\Phi(\Sigma) = \mathrm{E}\int \mathbf{1}_\Sigma(\theta_x\Phi)h(x)\Phi(\mathrm{d}x), \tag{1}$$

where $\mathbf{1}_\Sigma(\phi)$ is the indicator function of these events Σ and $h : X \to \mathbb{R}_+$ is any function such that $\int h(x)\,\mathrm{d}x = 1$. Here and afterwards, the integration domain is the whole of X unless explicitly written. The Palm distribution Q^Φ for a simple point process Φ is concentrated on configurations ϕ such that $\phi(\{0\}) = 1$ and can be regarded as a distribution of a random configuration viewed from a typical point of Φ.

A *random closed set* Ξ is a measurable mapping $\Xi : (\mathcal{N}, \{\mathcal{F}_K\}, \mathrm{P}) \to (\mathbb{F}, \sigma_f)$, where σ_f is the σ-algebra generated by the system $\{F \in \mathbb{F} : F \cap K \neq \varnothing\},\ K \in \mathbb{K}$.

A random compact set S is called a *stopping set* (more precisely, $\{\mathcal{F}_K\}$-stopping set) if the event $\{S \subseteq K\}$ is $\mathcal{F}_K$-measurable for all $K \in \mathbb{K}$. It is a natural generalization of the notion of a stopping time: knowing the configuration inside a compact set K is sufficient to conclude whether S is a subset of K or not, very much in the same way as, for a stopping time τ, occurrence of the event $\{\tau \leq t\}$ is determined by the history up to a moment t only.

Let $\mathcal{F} = \bigvee_{K\in\mathbb{K}} \mathcal{F}_K$. The *stopping σ-algebra* is the following collection:

$$\mathcal{F}_S = \{A \in \mathcal{F} : A \cap \{S \subseteq K\} \in \mathcal{F}_K \ \forall K \in \mathbb{K}\}.$$

It is straightforward to verify that $\mathcal{F}_S$ is indeed a σ-algebra and that a random compact set S is a stopping set if and only if it is $\mathcal{F}_S$-measurable.

A *Poisson process* with intensity measure $\lambda(\mathrm{d}x)$ is a point process Π with the following two properties: the variables $\Pi(B_1), \dots, \Pi(B_k)$ are mutually independent for disjoint $B_1, \dots, B_k$ for any k; and $\Pi(B)$ follows a Poisson distribution with parameter $\lambda(B)$. As a result, for any Borel set B and any functional $F(\cdot)$,

$$\begin{aligned}\int F(\phi)\,\mathrm{P}(\mathrm{d}\phi) &= \int F(\phi|_B + \phi|_{B^c})\,\mathrm{P}(\mathrm{d}\phi)\\ &= \iint F(\phi_1|_B + \phi_2|_{B^c})\,\mathrm{P}(\mathrm{d}\phi_1)\,\mathrm{P}(\mathrm{d}\phi_2)\\ &= \iint F(\phi_1 + \phi_2)\,\mathrm{P}_B(\mathrm{d}\phi_1)\,\mathrm{P}_{B^c}(\mathrm{d}\phi_2). \end{aligned} \tag{2}$$

Here and afterwards, P_B denotes restriction of P onto $\mathcal{F}_B$. The above property implies complete independence of the Poisson process distribution, due to which $\mathrm{P} = \mathrm{P}_B \otimes \mathrm{P}_{B^c}$. In particular, a Poisson process is a Markov process. Therefore, it also possesses the *strong Markov property*:

$$\int F(\phi)\,\mathrm{P}(\mathrm{d}\phi) = \iint F(\phi_1|_{S(\phi_1)} + \phi_2|_{S^c(\phi_1)})\,\mathrm{P}(\mathrm{d}\phi_1)\,\mathrm{P}(\mathrm{d}\phi_2) \tag{3}$$

for every compact stopping set S; see Rozanov (1982, Theorem 4).

The property (3) can also be expressed as

$$\mathrm{E}[F(\Pi) \mid \mathcal{F}_S](\phi|_{S(\phi)}) = \mathrm{E}_{S^c(\phi)} F(\phi|_{S(\phi)} + \Pi) \tag{4}$$

(more exactly, this is one version of the conditional expectation). Since nonrandom compact sets are also stopping sets, (4) also covers (2).

A set-indexed process $X_K = X_K(\phi)$ for $K \in \mathbb{K}$ is $\{\mathcal{F}_K\}$-adapted if, for any $K \in \mathbb{K}$, the process X_M for $M \subseteq K$, $M \in \mathbb{K}$ is $\mathcal{F}_K$-measurable. Such a process is called a *martingale* (more precisely, a $(\{\mathcal{F}_K\}, \mathrm{P})$-martingale) if, for any two compact sets $K \subseteq L$,

$$\mathrm{E}[X_L \mid \mathcal{F}_K] = X_K \qquad \mathrm{P}\text{-a.s.}$$

The following statement is an analogue of Doob's optional sampling theorem for set-indexed martingales.

Theorem 1. *Let S_1, S_2 be two almost surely compact stopping sets such that $S_1 \subseteq S_2$ almost surely. Let X_K be a uniformly integrable martingale. Then*

$$\mathrm{E}[X_{S_2} \mid \mathcal{F}_{S_1}] = X_{S_1} \qquad a.s. \tag{5}$$

provided that $\mathrm{E}|X_{S_2}| < \infty$.

The proof is based on the result of Kurtz (1980) for directed processes and can be found in Zuyev (1999, Theorem 1).

An example of a uniform integrable martingale is provided by the *likelihood ratio*. If P and P′ are two probability distributions on $(\mathcal{N}, \{\mathcal{F}_K\})$ such that the restriction P'_K of P′ onto $\mathcal{F}_K$

is absolutely continuous with respect to the restriction P_K of P onto the same σ-algebra for any $K \in \mathbb{K}$, then the Radon–Nikodym derivative $L_K = \mathrm{d}\,\mathrm{P}'_K / \mathrm{d}\,\mathrm{P}_K$ is a uniformly integrable P-martingale. In particular, if P_λ and P_ρ are the distributions of stationary Poisson processes with intensities λ and ρ respectively, then

$$L_K(\phi) = \frac{\mathrm{d}\,\mathrm{P}_\rho}{\mathrm{d}\,\mathrm{P}_\lambda}(\phi) = \left(\frac{\rho}{\lambda}\right)^{\phi(K)} \mathrm{e}^{(\lambda-\rho)|K|}, \tag{6}$$

where $|K|$ stands for the d-volume of K.

3. Generalization of Slivnyak's theorem

In this section, we establish a result generalizing Slivnyak's famous theorem for a Poisson process Π, stating that its distribution is still the same when viewed from one of the Poisson particles. We show that even when viewed from one of the points of a new process $\Phi(\Pi)$ effectively dependent on realizations of Π and given that only a compact domain is used for construction of each Φ, the distribution of Π outside of this compact domain is still Poisson. A similar fact was established in Mecke and Muche (1995) in the particular case where Φ is the vertices of the Voronoi tessellation and the domain is the Delaunay empty disc.

Let $(\mathcal{N}, \{\mathcal{F}_K\}, \mathrm{P})$ be a filtered probability space, where P is invariant under the flow $\theta.$. Here and below, Π is a homogeneous Poisson process in X compatible with the flow: $\Pi(\theta_x\phi) = \theta_x\Pi(\phi)$. Assume that on the same space there is defined a process of random compact sets $\mathcal{G} = \{G_i\}$ compatible with the flow: $\mathcal{G}(\theta_x\phi) = \theta_x\mathcal{G}(\phi)$ for all $x \in X$. Let us also assume that each compact set G_i is supplied with a *centroid* $z_i = z(G_i)$ so that there is a bijection between the point process $\Phi = \{z(G_i)\}$ and $\mathcal{G}$. Finally, we assume that the centroids are chosen in a compatible manner: $z(\theta_x G_i) = \theta_x z(G_i)$. Under the above conditions, the point process Φ is stationary and $G(z_i, \phi) = G(0, \theta_{z_i}\phi)$, where obviously $G(z_i, \phi)$ is the random compact set with centroid z_i. Let λ_Φ be the intensity of Φ.

Theorem 2. *If $G(\phi) = G(0, \phi)$ is an $\{\mathcal{F}_K\}$-stopping set, then, for any* P*-integrable functional F,*

$$\int F(\phi)\mathrm{Q}^\Phi(\mathrm{d}\phi) = \iint F(\phi|_{G(\phi)} + \phi'|_{G^{\mathrm{c}}(\phi)})\mathrm{Q}^\Phi(\mathrm{d}\phi)\,\mathrm{P}(\mathrm{d}\phi'). \tag{7}$$

Proof. By the definition (1) of the Palm distribution and by the strong Markov property (3), we have

$$\begin{aligned}
\lambda_\Phi \int F(\phi)\mathrm{Q}^\Phi(\mathrm{d}\phi) &= \iint F(\theta_x\phi)h(x)\Phi(\mathrm{d}x)\,\mathrm{P}(\mathrm{d}\phi) \\
&= \iint F(\theta_x(\phi|_{G(\phi)} + \phi|_{G^{\mathrm{c}}(\phi)}))h(x)\Phi(\mathrm{d}x)\,\mathrm{P}(\mathrm{d}\phi) \\
&= \iiint F(\theta_x(\phi|_{G(\phi)} + \phi''|_{G^{\mathrm{c}}(\phi)}))h(x)\Phi(\mathrm{d}x)\,\mathrm{P}(\mathrm{d}\phi)\,\mathrm{P}(\mathrm{d}\phi'').
\end{aligned}$$

Now put $\phi' = \theta_{-x}\phi''$ and use the stationarity of the distribution due to which $\mathrm{P}(\mathrm{d}\phi'') = \mathrm{P}(\mathrm{d}\phi')$. Thus, the above expression equals

$$\begin{aligned}
&\iiint F((\theta_x\phi)|_{G(\theta_x\phi)} + \phi'|_{G^{\mathrm{c}}(\theta_x\phi)})h(x)\Phi(\mathrm{d}x)\,\mathrm{P}(\mathrm{d}\phi)\,\mathrm{P}(\mathrm{d}\phi') \\
&\qquad = \lambda_\Phi \iint F(\phi|_{G(\phi)} + \phi'|_{G^{\mathrm{c}}(\phi)})\mathrm{Q}^\Phi(\mathrm{d}\phi)\,\mathrm{P}(\mathrm{d}\phi'),
\end{aligned}$$

which ends the proof.

Corollary 1. *If G_i are just the points of Π and hence $\Phi = \Pi$, then (7) transforms into Slivnyak's theorem:*

$$\int F(\phi)\,\mathrm{P}^0(\mathrm{d}\phi) = \int F(\delta_0 + \phi')\,\mathrm{P}(\mathrm{d}\phi'),$$

where P^0 is the Palm distribution of Π; see Slivnyak (1962) and Mecke (1967).

4. Sequential gamma-type results

In this part, we consider a homogeneous Poisson process Π of intensity λ in $X = \mathbb{R}^d$, $\{\mathcal{F}_K\}$ being its natural filtration and P_λ its distribution. Let $S_0 \subseteq S_1 \subseteq S_2 \subseteq \cdots$ be a (possibly finite) sequence of $\{\mathcal{F}_K\}$-stopping sets. Define $\mathcal{F}_k = \mathcal{F}_{S_k}$, $\Delta_k = S_k \setminus S_{k-1}$, $A_k = |\Delta_k|$ and $\pi_k = \Pi(\Delta_k)$ (by convention, $S_{-1} = \varnothing$ and $\mathcal{F}_{-1}$ is the trivial σ-algebra).

Theorem 3. *Assume that $\{\pi_k\}$ is a predictable sequence, i.e. π_k is $\mathcal{F}_{k-1}$-measurable for all $k = 0, 1, \dots$. Then*

(i) *the conditional distribution of A_k given $\mathcal{F}_{k-1}$ is $\Gamma(\pi_k, \lambda)$;*

(ii) *the Laplace transforms $\phi_{A_k}(z) = \mathrm{E}_\lambda \mathrm{e}^{zA_k}$ of A_k and $\phi_{\pi_k}(z)$ of π_k are related through*

$$\phi_{A_k}(z) = \phi_{\pi_k}\left(-\log\left(1 - \frac{\lambda}{z}\right)\right);$$

(iii) *if, in addition, there exists an $m \geq 0$ such that $\pi_0, \dots, \pi_m$ are constants, then A_k for $k = 0, \dots, m$ are independent $\Gamma(\pi_k, \lambda)$-distributed random variables.*

Proof. Let $\xi_k = L_{S_k}$, where L is the likelihood ratio (6) with a parameter ρ to be specified later. By Theorem 1, the sequence ξ_k is a positive $(\mathcal{F}_k, \mathrm{P}_\lambda)$-martingale (ordinary). Thus (5) can also be written as

$$\mathrm{E}_\lambda\left[\frac{\xi_k}{\xi_{k-1}} \,\middle|\, \mathcal{F}_{k-1}\right] = 1 \qquad \mathrm{P}_\lambda\text{ -a.s.}$$

Using the expression (6) this can be written as

$$\mathrm{E}_\lambda\left[\left(\frac{\rho}{\lambda}\right)^{\pi_k} \mathrm{e}^{(\lambda-\rho)A_k} \,\middle|\, \mathcal{F}_{k-1}\right] = 1 \qquad \text{a.s.}$$

The function $(\rho/\lambda)^{\pi_k}$ is $\mathcal{F}_{k-1}$-measurable, and therefore

$$\mathrm{E}_\lambda[\mathrm{e}^{(\lambda-\rho)A_k} \mid \mathcal{F}_{k-1}] = \left(\frac{\lambda}{\rho}\right)^{\pi_k} \qquad \text{a.s.}$$

Putting now $z = \lambda - \rho$, we see that

$$\mathrm{E}_\lambda[\mathrm{e}^{zA_k} \mid \mathcal{F}_{k-1}] = \left(1 - \frac{\lambda}{z}\right)^{-\pi_k} \qquad \text{a.s.,} \tag{8}$$

which proves part (i) of the theorem. Part (ii) follows immediately after taking the expectation of both sides in (8).

In the case when π_k are constants, the conditional distributions of A_k are seen not to depend on $\mathcal{F}_k$. So (iii) is proved by the following chain:

$$\begin{aligned}
\mathrm{E}_\lambda \exp\left\{\sum_{k=0}^{m} z_k A_k\right\} &= \mathrm{E}_\lambda \exp\left\{\sum_{k=0}^{m-1} z_k A_k\right\}\mathrm{E}_\lambda[\mathrm{e}^{z_m A_m} \mid \mathcal{F}_{m-1}] \\
&= \mathrm{E}_\lambda \exp\left\{\sum_{k=0}^{m-1} z_k A_k\right\}\left(1-\frac{\lambda}{z_m}\right)^{-\pi_m} \\
&= \cdots \\
&= \prod_{k=0}^{m}\left(1-\frac{\lambda}{z_k}\right)^{-\pi_k}.
\end{aligned}$$

The condition used above that π_k be $\mathcal{F}_{k-1}$-measurable for all $k = 0, 1, \dots$ means that the number of particles in Δ_k must be known before the construction of Δ_k from information available after S_{k-1} has been constructed. Predictability of π_k is, however, not always needed for a gamma-type result to hold, as the next theorem shows.

Theorem 4. *If, for some n and $k \geq 1$, $\mathrm{P}_\lambda\{\pi_k = n \mid \mathcal{F}_{k-1}\}$ is positive and not dependent on λ, then the $\mathrm{P}_\lambda[\cdot \mid \pi_k = n; \mathcal{F}_{k-1}]$ distribution of A_k is $\Gamma(n, \lambda)$.*

Proof. From Theorem 2, the point process on $\mathbb{R}^d \setminus S_{k-1}$ is Poisson, independent of the events in $\mathcal{F}_{k-1}$. So the current theorem follows from Theorem 2 of Zuyev (1999). It should only be noted that this theorem deals with compact stopping sets while $\Delta_k = S_k \setminus S_{k-1}$ is not compact in $\mathbb{R}^d$. However, it *is* compact in the topology *induced* on $X = \mathbb{R}^d \setminus S_{k-1}$, so the machinery of Zuyev (1999) is applicable to the stopping set Δ_k and the restricted process $\Pi|_X$, which is Poisson on X under the distribution $\mathrm{P}_\lambda(\cdot \mid \mathcal{F}_{S_{k-1}})$ by Theorem 2.

Theorem 2 considered earlier has further important consequences in the context of these two 'sequential' results: since the distribution of points under Q^Φ outside of $G = G(0, \phi)$ is still Poisson, we may apply the same machinery embodied in Theorems 3 and 4 to arrive at the following Palm version of the sequential gamma-type results.

Theorem 5. (i) *Assume that the conditions of Theorem 3 hold with $S_0(\phi) = G(0, \phi)$. Then the statements of Theorem 3 are valid for $k \geq 1$ (not for $k = 0$ in general!) almost surely with respect to a conditional Palm distribution Q^Φ_G (i.e. with $\mathrm{E}_\lambda[\cdot]$ replaced by $\mathrm{E}^\Phi[\cdot \mid \mathcal{F}_G]$ everywhere).*

(ii) *Theorem 4 is valid with $\mathrm{P}_\lambda[\cdot]$ replaced by $\mathrm{P}^\Phi[\cdot \mid \mathcal{F}_G]$.*

In particular, both (i) and (ii) are true for the Palm distribution P^0 of Π.

5. Applications

Here we address the decomposition issues posed in Section 1; we refer to the corresponding examples described there.

Example 2. (*Voronoi sausage.*) Theorem 3(iii) applies directly here with $S_k = \bigcup_{i=0}^{k} B_i$ and $\pi_i = 1$ for all i. So the areas of B_0 and the lunes are indeed independent $\mathrm{Exp}(\lambda)$, implying a gamma distribution for the volume of the whole sausage. The whole sausage is, of course, a stopping set.

Example 3. (*Delaunay lunes.*) Consider the process Φ of the Voronoi vertices. On the Palm space of Φ, the Delaunay circumdisc centred at the origin O is a stopping set (disc B in Figure 2); we denote its area by C. Let l_1, l_2, l_3 be the Delaunay edges emanating from O and ordered according to the size of their angle with the positive x-semi-axis. Let $B(l_i)$ be the Delaunay empty disc centred at the other endpoint of l_i for $i = 1, 2, 3$. Then $S_0 = B$, $S_1 = S_0 \cup B(l_1)$, $S_2 = S_1 \cup B(l_2)$ and $S_3 = S_2 \cup B(l_3)$ are stopping sets and, according to Theorem 5 applied to Theorem 3(iii), the areas of the lunes $S_i \setminus S_{i-1}$ for $i = 1, 2, 3$ given $\mathcal{F}_{S_0}$ are conditionally independent and $\mathrm{Exp}(\lambda)$ distributed. Thus, also for a uniformly randomly chosen edge l, the Q^Φ-conditional distribution of the area A_l of the lune $B(l) \setminus B$ given $\mathcal{F}_B$ is $\mathrm{Exp}(\lambda)$, independent of $\mathcal{F}_B$ (and in particular of C). So the unconditional Q^Φ distribution of A_l is also $\mathrm{Exp}(\lambda)$.

Recall from Section 1 that both B and $B(l)$ are typical Delaunay circumdiscs, each with $\Gamma(2, \lambda)$-distributed areas. Thus, by rôle reversal, the area A of the lune $B \setminus B(l)$ is also $\mathrm{Exp}(\lambda)$ independently of the area of $B(l)$. So we have the intriguing fact that C, the area of the disc B, is $\Gamma(2, \lambda)$ and that a well-defined part of it, namely the lune $B \setminus B(l)$, has an area A distributed as $\mathrm{Exp}(\lambda)$. Knowing that the sum of two independent $\mathrm{Exp}(\lambda)$ variables has the $\Gamma(2, \lambda)$ distribution, we may ask: is the area of the remaining part of B, the lens $B \cap B(l)$ (whose area we denote by $A_\cap$), distributed as $\mathrm{Exp}(\lambda)$ too? The answer turns out to be no!

Let R be the radius of B and α be the angle between the edge l and the radius vector of the intersection point of the two discs. Also, let sR be the length of l. Given C and α, we can show that the area of $B(l)$ is, for $s \geq 0$,

$$A_\cap + A_l = C(s^2 - 2s\cos\alpha + 1) \tag{9}$$

while we can write $A_\cap$ as

$$A_\cap = \frac{C}{\pi}\left[\alpha - s\sin\alpha + (s^2 - 2s\cos\alpha + 1)\arccos\frac{s - \cos\alpha}{\sqrt{s^2 - 2s\cos\alpha + 1}}\right]. \tag{10}$$

For any given C and α, $A_\cap$ is a monotone increasing function of s, starting at C and declining to an infimum of $C(\alpha - \sin\alpha\cos\alpha)/\pi$. On the other hand, from (9) and (10), A_l increases monotonely with s from 0 to infinity. Note that, if $\alpha < \pi/2$, then $A_\cap + A_l$ attains a minimum value of $C\sin^2\alpha$ at $s = \cos\alpha$.

Denote the monotone increasing function $s \mapsto A_l$ by $Cg_\alpha(\cdot)$. Then

$$g_\alpha(s) = \frac{1}{\pi}(s^2 - 2s\cos\alpha + 1)\arccos\frac{\cos\alpha - s}{\sqrt{s^2 - 2s\cos\alpha + 1}} - \frac{1}{\pi}(\alpha - s\sin\alpha).$$

Since A_l has probability density function $\lambda\exp\{-\lambda x\}$ independently of (C, α), we can immediately write the conditional distribution function of s given (C, α) as

$$F_s(C, \alpha) = 1 - \exp\{-\lambda C g_\alpha(s)\}, \qquad s \geq 0.$$

The probability density function of the angle α is known to be

$$f(\alpha) = \frac{4\sin\alpha}{3\pi}[(\pi - \alpha)\cos\alpha + \sin\alpha], \qquad \alpha \in (0, \pi),$$

proved in Mecke and Muche (1996) and, by a different method, in Cowan (2002) (where further detail of the above calculations is provided).

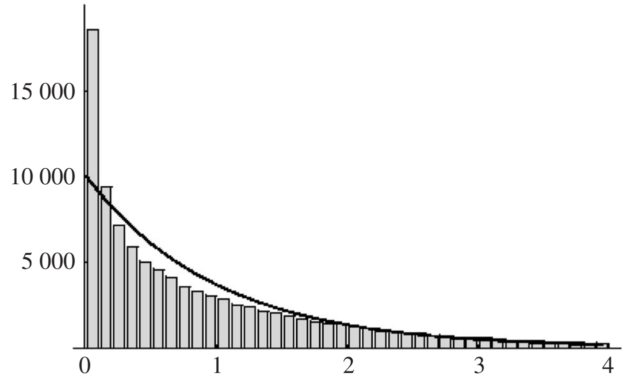

FIGURE 4: Frequency distribution of $A_\cap$ for a sample of 100 000 simulated cases with $\lambda = 1$. Overlaid on the empirical results is the probability density function of the exponential distribution of mean 1.

It is simple to generate the independent random variables C, α and A_l. Given these, we can calculate the random variable s as the unique root of $A_l = Cg_\alpha(s)$. Then $A_\cap$ can be calculated using (10). This has been done 100 000 times (using $\lambda = 1$); the histogram for $A_\cap$ is plotted in Figure 4.

Under the null hypothesis that the true distribution of $A_\cap$ is Exp(1), the data were classified into 10 equi-probable classes, leading to class frequencies of

19 158, 10 532, 8819, 8071, 7886, 7631, 8015, 8486, 9272 and 12 128.

The resulting Pearson χ^2 statistic was 11 064; obviously the hypothesis is rejected. The distribution is longer tailed, with greater weight at very small and very large values. While the sample mean was close to 1 at 0.9951, the second sample moment, 2.3278, was clearly larger than the hypothesized value of 2.

As a check, the simulation was repeated for $A_\cap + A_l$ and A; results were highly consistent with the known $\Gamma(2, 1)$ and Exp(1) laws respectively.

Thus, we have demonstrated (though not exactly proved mathematically) that $A_\cap$ does *not* conform to the Exp(λ) distribution, even though:

(i) both A and A_l have Exp(λ) laws;

(ii) both $A + A_\cap$ and $A_l + A_\cap$ are $\Gamma(2, \lambda)$ distributed;

(iii) $A + A_\cap + A_l$ is $\Gamma(3, \lambda)$ distributed;

(iv) A is independent of $A_\cap + A_l$; and

(v) A_l is independent of $A + A_\cap$.

Example 4. (*Voronoi flower.*) When trying to apply Theorem 3 to this example, we face the problem of the choice of the first stopping set S_0 as the lunes thus defined are *not* stopping sets. The geometry of a lune $B_{k+1} \setminus B_k$ depends on information outside it, namely, on B_k. We can circumvent this by defining a domain S_0 as a starting set for a sequence of stopping sets. Let S_0 be the largest disc centred on the positive x-axis passing through the origin and one of the Poisson particles (call it x_1), *and* not having any Poisson particles in its interior (see Figure 5). Clearly S_0 is a stopping set.

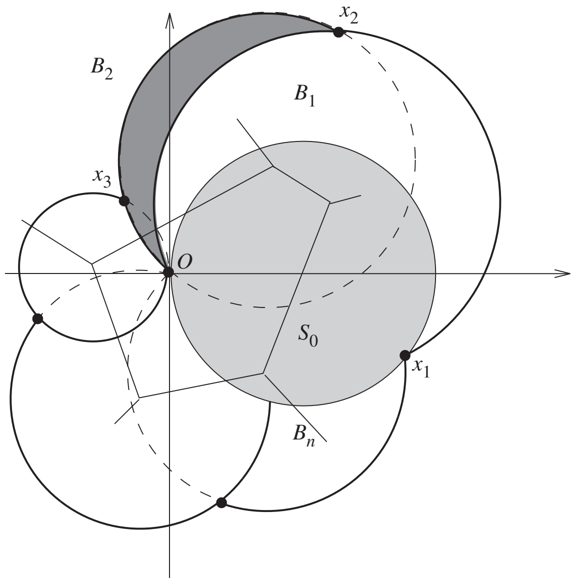

FIGURE 5: Stopping set S_0 and the second lune L_2.

The right bisector of O and x_1 can be seen in Figure 5; it is the side of the Voronoi polygon cut by the positive x-axis. To form the next stopping set in the sequence, consider the continuum of discs passing through O and x_1, with centre moving upward along this right bisector. Stop when this 'growing' disc first hits another particle (which is labelled x_2). This disc is B_1 and our second stopping set in the sequence is now defined as $S_1 = S_0 \cup B_1$. In a similar fashion, we move a circle-centre along the next right bisector, stopping the growing disc (which passes through O and x_2) when it hits another particle, x_3. Now $S_2 = S_1 \cup B_2$. The last of these constructions stops when x_1 is encountered by a growing disc. This algorithm successfully constructs the Voronoi flower.

So, with B_1 being the circumcircle of O, x_1 and x_2, B_2 the circumcircle of O, x_2 and x_3, and so on, we have the sequence of stopping sets S_0, $S_1 = S_0 \cup B_1$, $S_2 = S_1 \cup B_2, \dots, S_n = S_{n-1} \cup B_n = F$. Note that not all of the 'added' domains $S_i \setminus S_{i-1}$ are lunes. Some are discs with more than one 'bite' taken out.

In Section 1, we defined the lunes $L_1 = B_1 \setminus B_n$, $L_2 = B_2 \setminus B_1, \dots, L_n = B_n \setminus B_{n-1}$. While pictures may sometimes suggest that these lunes correspond to the 'added' domains—see the most darkly shaded L_2 in Figure 5, which does coincide with $S_2 \setminus S_1$—this is not general.

By Theorem 3, the area of S_0 has $\mathrm{Exp}(\lambda)$ distribution. The same is true for the areas of $S_1 \setminus S_0$ and $S_2 \setminus S_1$. The theorem is not applicable to $S_k \setminus S_{k-1}$ for $k \geq 3$ as $\pi_k = \Pi(S_k \setminus S_{k-1})$ is not predictable any more. Indeed, for $k < 3$, we can predict that $\pi_k = 1$; for larger k, we cannot predict from $\mathcal{F}_{k-1}$ whether it equals 1 or 0. Theorem 4 does not help when $k \geq 3$ because $\mathrm{P}\{\pi_k = 1 \mid \mathcal{F}_{k-1}\}$ depends on λ. Although it is disappointing not to find application of the theorems for $k \geq 3$, we consider that we have benefited from clarification of the precise conditions necessary for the complementary theorem to work in a sequential way. For example, $\mathrm{P}\{\pi_k = 1\}$, without the conditioning, does not depend on λ; it is just the probability that the typical Voronoi polygon has more than k sides, a scale-independent quantity. Yet, with the conditioning on $\mathcal{F}_{k-1}$, there *is* dependence.

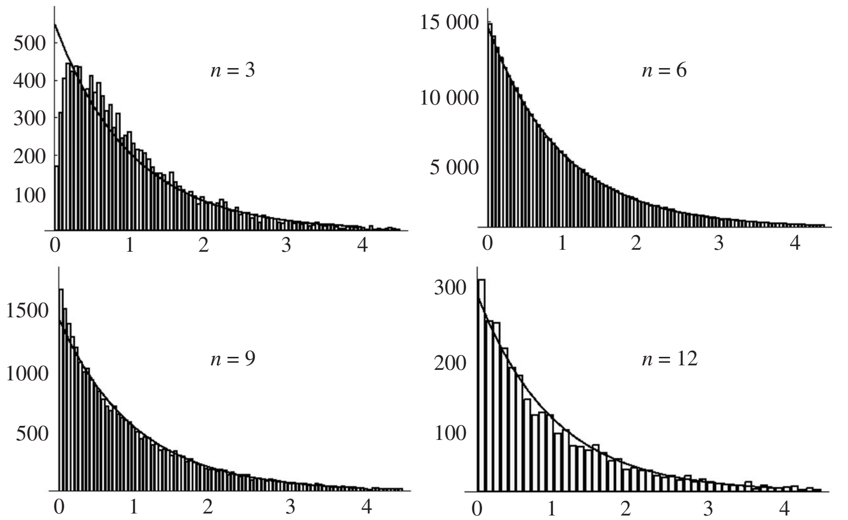

FIGURE 6: Frequency distributions for lune areas for $n = 3, 6, 9, 12$. For the first three cases, one lune was sampled from each independently generated Voronoi flower to build the histogram. Since we had only 239 independent flowers with $n = 12$, that histogram includes all 12 lunes, albeit slightly correlated in area, from each flower.

We note that the domain $S_1 \setminus S_0$ (the white part of B_1 in Figure 5) is always a lune, but not one of the lunes in our decomposition of F. Indeed, it is always smaller than L_1. Thus, we may conclude that the area of L_1 is *larger* than an Exp(λ)-distributed area of $S_1 \setminus S_0$; so its mean is greater than $1/\lambda$. Quite a special case is the lune L_2. Although in Figure 5 it coincides with the 'added domain' $S_2 \setminus S_1$, this is not the case if the centre of B_2 lies below the x-axis. Then $S_2 \setminus S_1 = L_2 \setminus S_0$. Thus, even the second lune has an area which is stochastically larger than the exponential distribution.

The observations noted above on theoretical grounds were reinforced by a large simulation study. We simulated one million typical Poisson–Voronoi cells with intensity 1 and then separately analysed polygons with n sides for each observed value of n. The simulation algorithm used was based on a radially symmetric method proposed in Quine and Watson (1984) and coded by Andrew Hayen.

The upward bias in the area of L_1, anticipated from the theory, was clearly evident in the simulations. For instance, for the 106 413 quadrilaterals, the means of the areas of $L_1, \dots, L_4$ were 1.13, 1.00, 1.02 and 0.85, while for the 89 929 octogons the means were 1.41, 1.02, 1.01, 1.02, 1.02, 0.99, 0.87 and 0.66. Of course, this is a labelling artefact. The side of the Voronoi cell which is hit by an axis passing through its centre is biased upward, so discs B_1 and B_n have an upwardly biased separation between centres. That makes both discs bigger than average, with an upward bias in the area of L_1 and a downward bias in the area of L_n.

The simulations showed that the difference between the areas of $S_2 \setminus S_1$ and L_2 is very small; the mean value of the area of $L_2 \cap S_0$ calculated on the one million flowers was approximately 0.00098. Thus, L_2 has an area which is only marginally larger than the exponential distribution. Although the hypothesis regarding exponentially distributed lunes is false, there is one lune in the Voronoi flower which has an 'almost exponential' area, and there is a domain which is 'almost a lune' that has an Exp(λ) area.

Of interest are the simulation results when, for each polygon, the lunes are numbered from a random starting point s. Figure 6 shows the histograms of areas of L_s for 3-, 6-, 9- and 12-sided polygons, with the Exp(1) density superimposed. Means of L_s were very close to 1 while variances ranged from about 0.78 at $n = 3$ to 1.32 at $n = 12$.

The 6-sided case is of special interest. The χ^2 goodness-of-fit test for Exp(1), for the area of L_s, for the 295 199 flowers could not reject the hypothesis. Although a formal test rejected the hypothesis of independence of the lunes' areas, the observed correlations were mostly negative, but less than 0.01 in absolute value.

Neighbouring lune areas for other values of n are clearly not independent; estimated correlations between randomly labelled neighbouring lunes L_s and L_{s+1} were about 0.1 for $n = 3$, 0.04 for $n = 4$, 0.01 for $n = 5$ and mostly negative and smaller than -0.01 for $n > 6$. Details can be found on the authors' Web pages, http://www.maths.usyd.edu.au/u/richardc/, http://www.maths.usyd.edu.au/u/malcolmq/, http://www.stams.strath.ac.uk/∼sergei/.

Acknowledgement

This research was supported by an Australian Research Council grant.

References

Baccelli, F., Tchoumatchenko, K. and Zuyev, S. (2000). Markov paths on the Poisson–Delaunay graph with applications to routeing in mobile networks. *Adv. Appl. Prob.* **32,** 1–18.

Cowan, R. (2000). Some Delaunay circumdisks and related lunes. Tech. Rep., University of Sydney. Available at http://www.maths.usyd.edu.au/u/richardc/.

Hayen, A. and Quine, M. P. (2002). Areas of components of a Voronoi polygon in a homogeneous Poisson process in the plane. *Adv. Appl. Prob.* **34,** 281–291.

Kurtz, T. (1980). The optional sampling theorem for martingales indexed by directed sets. *Ann. Prob.* **8,** 675–681.

Mecke, J. (1967). Stationäre zufällige Masse auf localcompakten Abelischen Gruppen. *Z. Wahrscheinlichkeitsth.* **9,** 36–58.

Mecke, J. and Muche, L. (1995). The Poisson Voronoi tessellation I. A basic identity. *Math. Nachr.* **176,** 199–208.

Mecke, J. and Muche, L. (1996). The Poisson Voronoi tessellation II. Edge length distribution functions. *Math. Nachr.* **178,** 271–283.

Miles, R. (1970). On the homogeneous planar Poisson process. *Math. Biosci.* **6,** 85–127.

Miles, R. and Maillardet, R. (1982). The basic structures of Voronoi and generalized Voronoi polygons. In *Essays in Statistical Science* (J. Appl. Prob. Spec. Vol. **19A**), eds J. Gani and E. J. Hannan, Applied Probability Trust, Sheffield, pp. 97–111.

Møller, J. and Zuyev, S. (1996). Gamma-type results and other related properties of Poisson processes. *Adv. Appl. Prob.* **28,** 662–673.

Quine, M. P. and Watson, D. (1984). Radial generation of n-dimensional Poisson processes. *J. Appl. Prob.* **21,** 548–557.

Rozanov, Y. A. (1982). *Markov Random Fields*. Springer, New York.

Slivnyak, I. (1962). Some properties of stationary flows of homogeneous random events. *Teor. Veroyat. Primen.* **7,** 347–352 (in Russian). English translation: *Theory Prob. Appl.* **7,** 336–341.

Zuyev, S. (1999). Stopping sets: gamma-type results and hitting properties. *Adv. Appl. Prob.* **31,** 355–366.