Electronic Journal of Probability

Vol. 14 (2009), Paper no. 85, pages 2438–2462.

Journal URL
http://www.math.washington.edu/~ejpecp/

Variational characterisation of Gibbs measures with Delaunay triangle interaction

David Dereudre
LAMAV
Université de Valenciennes et du Hainaut-Cambrésis, Le Mont Houy
59313 Valenciennes Cedex 09
France
email: david.dereudre@univ-valenciennes.fr

Hans-Otto Georgii
Mathematisches Institut
Ludwig-Maximilians-Universität München
Theresienstr. 39
80333 München
Germany
email: georgii@math.lmu.de

Abstract

This paper deals with stationary Gibbsian point processes on the plane with an interaction that depends on the tiles of the Delaunay triangulation of points via a bounded triangle potential. It is shown that the class of these Gibbs processes includes all minimisers of the associated free energy density and is therefore nonempty. Conversely, each such Gibbs process minimises the free energy density, provided the potential satisfies a weak long-range assumption.

Key words: Delaunay triangulation, Voronoi tessellation, Gibbs measure, variational principle, free energy, pressure, large deviations.

AMS 2000 Subject Classification: Primary 60K35; Secondary: 60D05, 60G55, 82B21.

Submitted to EJP on June 11, 2009, final version accepted November 4, 2009.

1 Introduction

It is well-known that stationary renewal processes with a reasonable spacing distribution can be characterised as Gibbs processes for an interaction between nearest-neighbour pairs of points [16, Section 6]. Here we consider an analogue in two dimensions, viz. Gibbsian point processes on $\mathbb{R}^2$ with an interaction depending on nearest-neighbour triples of points, where the nearest-neighbour triples are defined in terms of the Delaunay triangulation. Recall that the Delaunay triangulation is dual to the Voronoi tessellation, in the sense that two points are connected by a Delaunay edge if and only if their Voronoi cells have a common edge. Since the Voronoi cell of a point consists of the part of space that is closer to this point than to any other point, this means that the Delaunay graph defines a natural nearest-neighbour structure between the points. (Of course, the analogy with renewal processes does not reach too far because the independence of spacings under the Palm distribution, which is characteristic of one-dimensional renewal processes, is lost in two dimensions due to the geometric constraints.)

There is a principal difference between the Delaunay interactions considered here and the pair interactions that are common in Statistical Physics. Namely, suppose a point configuration ω is augmented by a new particle at x. In the case of pair interactions, x is subject to some additional interaction with the particles in ω, but the interaction between the particles of ω is not affected by x. In the Delaunay case, however, the particle at x not only gives rise to some new tiles of the Delaunay triangulation, but also destroys some other tiles that were present in the triangulation of ω. This so-called non-hereditary nature of the Delaunay triangulation blurs the usual distinction between attractive and repulsive interactions and makes it difficult to use a local characterisation of Gibbs measures in terms of their Campbell measures and Papangelou intensities. Such a local approach to the existence of Gibbs measures for Delaunay interactions was used in the previous work [2; 3; 5; 7] and made it necessary to impose geometric constraints on the interaction by removing triangles with small angles or large circumcircles.

In this paper we address the existence problem from a global point of view, which is based on stationarity and thermodynamic quantities such as pressure and free energy density. Specifically, we show that all minimisers of the free energy density are Gibbsian, which implies the existence of Delaunay-Gibbs measures because the entropy density has compact level sets. The converse part of this variational principle is harder and requires the comparison of different boundary conditions in the thermodynamic limit. It is here that the non-hereditary nature of the interaction shows up again, but it can be controlled with the help of stationarity and an additional condition which is much weaker than the geometric constraints mentioned above. In contrast to [7], however, we need to assume throughout that the interaction potential is bounded, and therefore do not cover hard-core interactions that forbid particular shapes of the tiles. We note, however, that some ideas of the present paper can be used to establish the existence of Delaunay-Gibbs measures in a more general setting that includes also the hard-core case, see [8]. The extension of the variational principle to such interactions is left to future work. As a final comment, let us emphasise that we make repeated use of Euler's polyhedral formula and the resulting linear complexity of the Delaunay triangulations, and are therefore limited to two dimensions, as was already the case in the previous papers mentioned above.

2 Preliminaries

2.1 Configurations and Delaunay triangulations

A subset ω of $\mathbb{R}^2$ is called locally finite if $\operatorname{card}(\omega \cap \Delta) < \infty$ for all bounded $\Delta \subset \mathbb{R}^2$; each such ω is called a *configuration*. We write Ω for the set of all configurations ω. The configuration space Ω is equipped with the σ-algebra $\mathscr{F}$ that is generated by the counting variables $N_\Delta : \omega \to \operatorname{card}(\omega \cap \Delta)$, with Δ an arbitrary bounded Borel subset of $\mathbb{R}^2$.

For each Borel set $\Lambda \subset \mathbb{R}^2$ we write $\Omega_\Lambda = \{\omega \in \Omega : \omega \subset \Lambda\}$ for the set of configurations in Λ, $\mathrm{pr}_\Lambda : \omega \to \omega_\Lambda := \omega \cap \Lambda$ for the projection from Ω to Ω_Λ, $\mathscr{F}'_\Lambda = \mathscr{F}|\Omega_\Lambda$ for the trace σ-algebra of $\mathscr{F}$ on Ω_Λ, and $\mathscr{F}_\Lambda = \mathrm{pr}_\Lambda^{-1}\mathscr{F}'_\Lambda$ for the σ-algebra of events in Ω that happen in Λ only.

For each configuration $\omega \in \Omega$ we consider the Delaunay triangulation $\mathsf{D}(\omega)$ associated to ω. By definition,

$$\mathsf{D}(\omega) = \{\tau \subset \omega : \operatorname{card}\tau = 3,\ \omega \cap B(\tau) = \emptyset\}, \tag{2.1}$$

where $B(\tau)$ is the unique open disc with $\tau \subset \partial B(\tau)$. $\mathsf{D}(\omega)$ is uniquely defined and determines a triangulation of the convex hull of ω whenever ω is in general circular position, in that no four points of ω lie on a circle that contains no further points of ω inside [17]. If this is not the case, one can apply some determistic rule to make the Delaunay triangulation unique. Indeed, let

$$\mathscr{T} := \{\tau \subset \mathbb{R}^2 : \operatorname{card}\tau = 3\} = \{(x,y,z) \in (\mathbb{R}^2)^3 : x \prec y \prec z\}$$

be the set of all triangles (or tiles) in $\mathbb{R}^2$ where '$\prec$' stands for the lexicographic order in $\mathbb{R}^2$. The triangles in $\mathscr{T}$ can be compared by the lexicographic order of $(\mathbb{R}^2)^3$, and this in turn induces a lexicographic order on finite collections of triangles. Now, if $n \geq 4$ points of ω lie on a circle with no points inside then the associated Delaunay cell is a convex polygon having these n points as vertices. To define a unique triangulation of this polygon one can then simply take the smallest among all possible triangulations. Conflicts with other possible polygons cannot arise because the tessellations inside and outside a fixed convex polygon of $n \geq 4$ points on a circle do not depend on each other.

Let us note that the prescription $\omega \to \mathsf{D}(\omega)$ is a mapping from Ω to the set $\Omega(\mathscr{T})$ of all locally finite subsets of $\mathscr{T}$. If $\Omega(\mathscr{T})$ is equipped with the σ-algebra $\mathscr{F}(\mathscr{T})$ that is defined in analogy to $\mathscr{F}$, one can easily check that this mapping is measurable.

Next we assign to each tile $\tau \in \mathscr{T}$ a centre and a radius. Specifically, for every $\tau \in \mathscr{T}$ we write $c(\tau)$ for the centre and $\varrho(\tau)$ for the radius of the circumscribed disc $B(\tau)$. The centres allow us to consider $\mathsf{D}(\omega)$ as a germ-grain system, i.e., as a marked point configuration of germs in $\mathbb{R}^2$ and marks in the space

$$\mathscr{T}_0 = \{\tau \in \mathscr{T} : c(\tau) = 0\}$$

of centred tiles, by considering the mapping

$$D : \omega \to \{(c(\tau), \tau - c(\tau)) : \tau \in \mathsf{D}(\omega)\} \tag{2.2}$$

from Ω to the point configurations on $\mathbb{R}^2 \times \mathscr{T}_0$. Here we write $\tau - c(\tau) := \{y - c(\tau) : y \in \tau\}$ for the shifted tile.

A crucial fact we need in the following is the linear complexity of Delaunay triangulations, which is expressed in the following lemma. This result follows directly from Euler's polyhedral formula, and is the main reason why we need to confine ourselves to two spatial dimensions; see [1], Chapter 11, and [17], Remark 2.1.4.

Lemma 2.1. *For a simple planar graph on $n \geq 3$ vertices, the number of edges is at most $3n - 6$, and the number of inner faces is at most $2n - 5$. In particular, every triangulation with $n \geq 3$ nodes consists of $2n - 2 - \partial$ triangles, where ∂ is the number of nodes (or: number of edges) along the outer boundary.*

2.2 Stationary point processes and their tile distribution

Let $\mathscr{P}_\Theta$ be the set of all probability measures P on $(\Omega, \mathscr{F})$ that satisfy the following two properties:

(S) P is *stationary*, that is, P is invariant under the shift group $\Theta = (\vartheta_x)_{x \in \mathbb{R}^2}$ on Ω, which is defined by $\vartheta_x : \omega \to \omega - x := \{y - x : y \in \omega\}$.

(I) P has a *finite intensity* $z(P) = |\Delta|^{-1} \int N_\Delta \, dP < \infty$. Here, $\Delta \subset \mathbb{R}^2$ is any bounded Borel set of non-zero Lebesgue measure $|\Delta|$. (Δ can be arbitrarily chosen due to stationarity.)

Each $P \in \mathscr{P}_\Theta$ is called a *stationary point process on* $\mathbb{R}^2$ *with finite intensity*. For $\Lambda \subset \mathbb{R}^2$, we write $P_\Lambda := P \circ \mathrm{pr}_\Lambda^{-1}$ for the projection image of P on $(\Omega_\Lambda, \mathscr{F}'_\Lambda)$, which can of course be identified with the restriction $P|\mathscr{F}_\Lambda$ of P to the events in Λ.

Every $P \in \mathscr{P}_\Theta$ defines a germ-grain model $\bar{P}$, namely the distribution of P under the mapping D defined in (2.2). That is, $\bar{P}$ is a stationary marked point process on $\mathbb{R}^2$ with mark space $\mathscr{T}_0$. Let $\bar{P}^0$ be the associated Palm measure on $\mathscr{T}_0 \times \Omega$ and $\mu_P = \bar{P}^0(\,\cdot \times \Omega)$ the associated mark distribution, or *centred tile distribution*, on $\mathscr{T}_0$. By definition,

$$\int dx \int \mu_P(d\tau) f(x, \tau) = \int P(d\omega) \sum_{\tau \in \mathrm{D}(\omega)} f(c(\tau), \tau - c(\tau)) \tag{2.3}$$

for all nonnegative measurable functions f on $\mathbb{R}^2 \times \mathscr{T}_0$. For each $P \in \mathscr{P}_\Theta$, μ_P has total mass $\|\mu_P\| = 2z(P)$, as follows from Euler's polyhedral formula; see, for example, [17, Eq. (3.2.11)] or [20, Theorem 10.6.1(b)].

Let us say a measure $P \in \mathscr{P}_\Theta$ is *tempered* if

$$\int |B(\tau)| \, \mu_P(d\tau) < \infty \tag{2.4}$$

We write $\mathscr{P}_\Theta^{\mathrm{tp}}$ for the set of all tempered $P \in \mathscr{P}_\Theta$. Of course, (2.4) is equivalent to the condition $\int \varrho(\tau)^2 \, \mu_P(d\tau) < \infty$. Moreover, (2.3) implies that

$$\begin{aligned} \int |B(\tau)| \, \mu_P(d\tau) &= \int dx \int \mu_P(d\tau) \, \mathbb{1}_{\{x \in B(\tau)\}} \qquad (2.5) \\ = \int \sum_{\tau \in \mathrm{D}(\omega)} \mathbb{1}_{\{c(\tau) \in B(\tau - c(\tau))\}} \, P(d\omega) &= \int \mathrm{card}\{\tau \in \mathrm{D}(\omega) : 0 \in B(\tau)\} \, P(d\omega). \end{aligned}$$

So, P is tempered if and only if the last expression is finite. A sufficient condition for temperedness will be given in Proposition 4.9.

The most prominent members of $\mathscr{P}_\Theta^{\mathrm{tp}}$ are the Poisson point processes, which will take the role of reference processes for the models we consider. Recall that the Poisson point process Π^z with intensity $z > 0$ is characterised by the following two properties:

(P1) For every bounded Borel set Δ, the counting variable N_Δ is Poisson distributed with parameter $z|\Delta|$.

(P2) Conditional on $N_\Delta = n$, the n points in Δ are independent with uniform distribution on Δ, for every bounded Borel set Δ and each integer n.

The temperedness of Π^z follows from Proposition 4.3.1 of [17], or Proposition 4.9 below.

Another type of measures in $\mathscr{P}_\Theta^{\mathrm{tp}}$ are the stationary empirical fields that are defined as follows. Let $\Lambda \subset \mathbb{R}^2$ be an open square of side length L, and for $\omega \in \Omega_\Lambda$ let $\omega_{\Lambda,\mathrm{per}} = \{x + Li : x \in \omega, i \in \mathbb{Z}^2\}$ be its periodic continuation. The associated *stationary empirical field* is then given by

$$R_{\Lambda,\omega} = \frac{1}{|\Lambda|} \int_\Lambda \delta_{\vartheta_x \omega_{\Lambda,\mathrm{per}}}\, dx\,. \tag{2.6}$$

It is clear that $R_{\Lambda,\omega}$ is stationary. In addition, it is tempered because $2\varrho(\tau) \le \operatorname{diam}\Lambda$ for each triangle $\tau \in \mathsf{D}(\omega_{\Lambda,\mathrm{per}})$. Finally, one easily finds that

$$\mu_{R_{\Lambda,\omega}} = |\Lambda|^{-1} \sum_{\tau \in \mathsf{D}(\omega_{\Lambda,\mathrm{per}}):\, c(\tau)\in\Lambda} \delta_{\tau - c(\tau)}\,; \tag{2.7}$$

see the proof of the similar result in [14, Remark 2.3(3)]. In the following we will often consider the probability kernel $R_\Lambda : \omega \to R_{\Lambda,\omega} = R_{\Lambda,\omega_\Lambda}$, where ω may be taken from either Ω or Ω_Λ depending on the context.

2.3 The topology of local convergence

In contrast to the traditional weak topology on the set $\mathscr{P}_\Theta$ of stationary point processes, we exploit here a finer topology, which is such that the intensity is a continuous function, but nonetheless the entropy density has compact level sets.

Let $\mathscr{L}$ denote the class of all measurable functions $f : \Omega \to \mathbb{R}$ which are *local* and *tame*, in that there exists some bounded Borel set $\Delta \subset \mathbb{R}^2$ such that $f = f \circ \mathrm{pr}_\Delta$ and $|f| \le b(1 + N_\Delta)$ for some constant $b = b(f) < \infty$. The *topology $\mathscr{T}_\mathscr{L}$ of local convergence* on $\mathscr{P}_\Theta$ is then defined as the weak* topology induced by $\mathscr{L}$, i.e., as the smallest topology for which the mappings $P \to \int f\, dP$ with $f \in \mathscr{L}$ are continuous. By the definition of the intensity, it is then clear that the mapping $P \to z(P)$ is continuous.

A further basic continuity property is the fact that the centred tile distribution μ_P depends continuously on P. Let $\mathscr{L}_0$ be the class of all *bounded* measurable functions on the space $\mathscr{T}_0$ of centred tiles, and $\mathscr{T}_0$ the associated weak* topology on the set $\mathscr{M}(\mathscr{T}_0)$ of all finite measures on $\mathscr{T}_0$. (This is sometimes called the τ-topology.)

Proposition 2.2. *Relative to the topologies $\mathscr{T}_\mathscr{L}$ and $\mathscr{T}_0$ introduced above, the mapping $P \to \mu_P$ from $\mathscr{P}_\Theta$ to $\mathscr{M}(\mathscr{T}_0)$ is continuous.*

This result will be proved in Section 4.1. It takes advantage of the linear complexity of finite Delaunay triangulations, and therefore relies on the planarity of our model.

2.4 The entropy density

Recall that the relative entropy (or Kullback-Leibler information) of two probability measures P, Q on any measurable space is defined by

$$I(P;Q) = \begin{cases} \int f \log f \, dQ & \text{if } P \ll Q \text{ with Radon-Nikodym density } f, \\ \infty & \text{otherwise.} \end{cases}$$

It is well-known that $I(P;Q) \geq 0$ with equality precisely for $P = Q$. For a point process $P \in \mathscr{P}_\Theta$ and a bounded Borel set Λ in $\mathbb{R}^2$, we write $I_\Lambda(P, \Pi^z) = I(P_\Lambda; \Pi^z_\Lambda)$ for the relative entropy of the restrictions P_Λ and Π^z_Λ. By the independence properties of Π^z, these quantities are superadditive in Λ, which implies that the limit

$$I_z(P) = \lim_{|\Lambda| \to \infty} I_\Lambda(P; \Pi^z)/|\Lambda| \in [0, \infty] \tag{2.8}$$

exists and is equal to the supremum of this sequence. For our purposes, it is sufficient to take this limit along a fixed sequence of squares; for example, one can take squares with vertex coordinates in $\mathbb{Z} + 1/2$. The claim (2.8) then follows from the well-known analogous result for lattice models [10, Chapter 15] by dividing $\mathbb{R}^2$ into unit squares. I_z is called the (negative) *entropy density with reference measure* Π^z.

We set $I = I_1$. Each I_z differs from I only by a constant and a multiple of the particle density. In fact, an easy computation shows that

$$I(P) = I_z(P) + 1 - z + z(P) \log z \qquad \text{for all } z > 0 \text{ and } P \in \mathscr{P}_\Theta . \tag{2.9}$$

A crucial fact we need later is the following result obtained in Proposition 2.6 of [14].

Lemma 2.3. *In the topology $\mathscr{T}_{\mathscr{L}}$, each I_z is lower semicontinuous with compact level sets $\{I_z \leq c\}$, $c > 0$.*

2.5 Triangle interactions

This paper is concerned with point processes with a particle interaction which is induced by the associated Delaunay triangulation. We stick here to the simplest kind of interaction, which depends only on the triangles that occur in each configuration. Specifically, let $\varphi : \mathscr{T}_0 \to \mathbb{R}$ be an arbitrary measurable function. It can be extended to a unique shift-invariant measurable function φ on $\mathscr{T}$ via $\varphi(\tau) := \varphi(\tau - c(\tau))$, $\tau \in \mathscr{T}$. Such a φ will be called a *triangle potential*. We will assume throughout that φ is bounded, in that

$$|\varphi| \leq c_\varphi \tag{2.10}$$

for some constant $c_\varphi < \infty$. In Theorem 3.4 we will need the following additional condition to prove the temperedness of Gibbs measures. Let us say that a triangle potential φ is *eventually increasing* if there exist a constant $r_\varphi < \infty$ and a measurable nondecreasing function $\psi : [r_\varphi, \infty[\to \mathbb{R}$ such that $\varphi(\tau) = \psi(\varrho(\tau))$ when $\varrho(\tau) \geq r_\varphi$. This condition is clearly satisfied when φ is constant for all triangles τ with sufficiently large radius $\varrho(\tau)$.

Example 2.4. Here are some examples of triangle potentials. For each triangle $\tau \in \mathscr{T}$ let $b(\tau) = \frac{1}{3}\sum_{x\in\tau} x$ be the barycentre and $A(\tau)$ the area of τ. Examples of bounded (and scale invariant) interactions that favour equilateral Delaunay triangles are

$$\varphi_1(\tau) = \beta\,|c(\tau) - b(\tau)|/\varrho(\tau) \quad \text{or} \quad \varphi_2(\tau) = -\beta A(\tau)/\varrho(\tau)^2$$

with $\beta > 0$. Of course, many variants are possible; e.g., one can replace the barycentre by the centre of the inscribed circle. By way of contrast, to penalise regular configurations one can replace the φ_i's by their negative.

The triangle potentials φ_i above are not eventually increasing. But each triangle potential φ can be modified to exhibit this property by setting

$$\tilde{\varphi}(\tau) = \begin{cases} \varphi(\tau) & \text{if } \varrho(\tau) < r\,, \\ K & \text{otherwise} \end{cases}$$

with $r > 0$ and K a suitable constant. When K is large, one has the additional effect of favouring small circumcircles.

Remark 2.5. The type of interaction introduced above is the simplest possible that is adapted to the Delaunay structure. In particular, we avoid here any explicit interaction ψ along the Delaunay edges. This has two reasons: First, we might add a term of the form $\frac{1}{2}\sum_{e\subset\tau:\,\mathrm{card}\,e=2}\varphi_{\mathrm{edge}}(e)$ to the triangle interaction φ. Such a term would take account for an edge interaction φ_{edge} whenever $\mathsf{D}(\omega)$ is a triangulation of the full plane. Secondly, we often need to control the interaction over large distances; the condition of φ being eventually increasing is tailored for this purpose. It is then essential to define the range in terms of triangles rather than edges. Namely, if a configuration ω is augmented by a particle at a large distance from ω, the circumcircles of all destroyed triangles must be large, but their edges can be arbitrarily short. So, a large-circumcircle assumption on the triangle potential allows to control this effect, but a long-edge asumption on an edge potential would be useless.

3 Results

Let φ be a fixed triangle potential. *We assume throughout that φ is bounded*, see (2.10), but do not require in general that φ is eventually increasing. To introduce the Hamiltonians and Gibbs distributions for φ we first need to introduce a class of 'reasonable' configurations.

Definition. Let us say a configuration $\omega \in \Omega$ is *admissible* if for every bounded Borel set Λ there exists a compact set $\bar{\Lambda}(\omega) \supset \Lambda$ such that $B(\tau) \subset \bar{\Lambda}(\omega)$ whenever $\zeta \in \Omega$ and $\tau \in \mathsf{D}(\zeta_\Lambda \cup \omega_{\Lambda^c})$ is such that $B(\tau) \cap \Lambda \neq \emptyset$. We write Ω^* for the set of all admissible configurations.

In Corollary 4.2 we will show that $P(\Omega^*) = 1$ for all $P \in \mathscr{P}_\Theta$ with $P(\{\emptyset\}) = 0$. Suppose now that $\omega \in \Omega^*$ and $\Lambda \subset \mathbb{R}^2$ is a bounded Borel set. The *Hamiltonian for φ in Λ with boundary condition ω* is then defined by

$$H_{\Lambda,\omega}(\zeta) = \sum_{\tau\in\mathsf{D}(\zeta_\Lambda\cup\omega_{\Lambda^c}):\,B(\tau)\cap\Lambda\neq\emptyset} \varphi(\tau)\,, \tag{3.1}$$

where $\zeta \in \Omega$ is arbitrary. Since $\omega \in \Omega^*$, the sum is finite, so that the Hamiltonian is well-defined. Note also that $H_{\Lambda,\omega}(\emptyset) \neq 0$ in general. Moreover, Lemma 2.1 shows that

$$H_{\Lambda,\omega} \geq -2c_\varphi N_\Lambda - 2c_\varphi N_{\bar{\Lambda}(\omega)\setminus\Lambda}(\omega),$$

where $\bar{\Lambda}(\omega)$ is as above and c_φ from (2.10). This in turn implies that for each $z > 0$ the associated *partition function*

$$Z_{\Lambda,z,\omega} = \int e^{-H_{\Lambda,\omega}} \, d\Pi_\Lambda^z \tag{3.2}$$

is finite. We can therefore define the *Gibbs distribution with activity* $z > 0$ by

$$G_{\Lambda,z,\omega}(A) = Z_{\Lambda,z,\omega}^{-1} \int \mathbb{1}_A(\zeta \cup \omega_{\Lambda^c}) \, e^{-H_{\Lambda,\omega}(\zeta)} \, \Pi_\Lambda^z(d\zeta), \quad A \in \mathscr{F}. \tag{3.3}$$

The measure $G_{\Lambda,z,\omega}$ depends measurably on ω and thus defines a probability kernel from $(\Omega^*, \mathscr{F}_{\Lambda^c})$ to $(\Omega, \mathscr{F})$.

Definition. Let P be a probability measure on $(\Omega, \mathscr{F})$ which is of first moment, in that $\int N_\Lambda \, dP < \infty$ for all bounded Borel sets Λ in $\mathbb{R}^2$. P is called a *Gibbs point process for the Delaunay triangle potential* φ *and the activity* $z > 0$, or a *Delaunay-Gibbs measure* for short, if $P(\Omega^*) = 1$ and $P = \int P(d\omega) \, G_{\Lambda,z,\omega}$ for all bounded Borel sets $\Lambda \subset \mathbb{R}^2$. We write $\mathscr{G}_\Theta(z, \varphi)$ for the set of all stationary Gibbs measures for φ and z, and $\mathscr{G}_\Theta^{\mathrm{tp}}(z, \varphi)$ for the set of all *tempered* stationary Gibbs measures; recall Eq. (2.4).

The above definition corresponds to the classical concept of a Gibbs measure, which is based on the location of points. We note that an alternative concept of Gibbs measure that considers the location of Delaunay triangles has been proposed and used by Zessin [21] and Dereudre [7].

Intuitively, the interaction of a configuration in Λ with its boundary condition ω reaches not farther than the set $\bar{\Lambda}(\omega)$ above, which guarantees some kind of quasilocality. So one can expect that a limit of suitable Gibbs distributions $G_{\Lambda,z,\omega}$ as $\Lambda \uparrow \mathbb{R}^2$ should be Gibbsian, and the existence problem reduces to the question of whether such limits exist. Our approach here is to take the necessary compactness property from Lemma 2.3, the compactness of level sets of the entropy density. In fact, we even go one step further and show that the stationary Gibbs measures are the minimisers of the free energy density. Since such minimisers exist by the compactness of level sets, this solves in particular the existence problem. The free energy density is defined as follows; recall the definition of the centred tile distribution μ_P before (2.3).

Definition. The *energy density* of a stationary point process $P \in \mathscr{P}_\Theta$ is defined by

$$\Phi(P) = \int \varphi \, d\mu_P = |\Delta|^{-1} \int P(d\omega) \sum_{\tau \in \mathrm{D}(\omega):\, c(\tau) \in \Delta} \varphi(\tau),$$

where Δ is an arbitrary bounded Borel set of positive Lebesgue measure. The *free energy density* of P relative to Π^z is given by $I_z(P) + \Phi(P)$.

The definition of Φ is justified by Proposition 3.6 below. Here are some crucial facts on the free energy density, which will be proved in Subsection 4.1.

Proposition 3.1. *Relative to the topology $\mathscr{T}_{\mathscr{L}}$ on $\mathscr{P}_\Theta$, Φ is continuous, and each $I_z + \Phi$ is lower semicontinuous with compact level sets. In particular, the set $\mathscr{M}_\Theta(z,\varphi)$ of all minimisers of $I_z + \Phi$ is a non-empty convex compact set, and in fact a face of the simplex $\mathscr{P}_\Theta$.*

Next we observe that the elements of $\mathscr{M}_\Theta(z,\varphi)$ are nontrivial, in that the empty configuration $\emptyset \in \Omega$ has zero probability; this result will also be proved in Subsection 4.1.

Proposition 3.2. *For all $z > 0$ we have $\delta_\emptyset \notin \mathscr{M}_\Theta(z,\varphi)$, and thus $P(\{\emptyset\}) = 0$ for all $P \in \mathscr{M}_\Theta(z,\varphi)$.*

Our main result is the following variational characterisation of Gibbs measures.

Theorem 3.3. *Let φ be a bounded triangle potential and let $z > 0$. Then every minimiser of the free energy density is a stationary Gibbs measure. That is, the inclusion $\mathscr{M}_\Theta(z,\varphi) \subset \mathscr{G}_\Theta(z,\varphi)$ holds. In particular, Gibbs measures exist. Conversely, every tempered stationary Gibbs measure is a minimiser of the free energy density, which means that $\mathscr{G}^{\mathrm{tp}}_\Theta(z,\varphi) \subset \mathscr{M}_\Theta(z,\varphi)$.*

The proof will be given in Subsections 4.2 and 4.5. Theorem 3.3 raises the problem of whether all stationary Gibbs measures are tempered. It is natural to expect that $\mathscr{G}_\Theta(z,\varphi) = \mathscr{G}^{\mathrm{tp}}_\Theta(z,\varphi)$, but we did not succeed to prove this in general. In fact, we even do not know whether $\mathscr{G}^{\mathrm{tp}}_\Theta(z,\varphi)$ is always non-empty. But we can offer the following sufficient condition, which will be proved in Subsection 4.6.

Theorem 3.4. *Suppose φ is eventually increasing and let $z > 0$. Then every stationary Gibbs measure is tempered, so that $\mathscr{G}^{\mathrm{tp}}_\Theta(z,\varphi) = \mathscr{G}_\Theta(z,\varphi)$.*

Combining Theorems 3.3 and 3.4 we arrive at the following result.

Corollary 3.5. *Suppose φ is bounded and eventually increasing, and let $z > 0$. Then the minimisers of the free energy density are precisely the stationary tempered Gibbs measures. That is, $\mathscr{M}_\Theta(z,\varphi) = \mathscr{G}_\Theta(z,\varphi) = \mathscr{G}^{\mathrm{tp}}_\Theta(z,\varphi)$ for all $z > 0$.*

The proof of Theorem 3.3 is based on an analysis of the mean energy and the pressure in the infinite volume limit when $\Lambda \uparrow \mathbb{R}^2$. For simplicity, we take this limit through a fixed reference sequence, namely the sequence

$$\Lambda_n = \left] -n - \tfrac{1}{2}, n + \tfrac{1}{2} \right[^2 \tag{3.4}$$

of open centred squares. We shall often write n when we refer to Λ_n. That is, we set $\omega_n = \omega_{\Lambda_n}$, $P_n = P_{\Lambda_n}$, $R_{n,\omega} = R_{\Lambda_n,\omega}$, $H_{n,\omega} = H_{\Lambda_n,\omega}$, and so on. We also write $v_n = |\Lambda_n| = (2n+1)^2$ for the Lebesgue measure of Λ_n. Our first result justifies the above definition of $\Phi(P)$. Besides the Hamiltonian (3.1) with configurational boundary condition ω, we will also consider the *Hamiltonian with periodic boundary condition*, namely

$$H_{n,\mathrm{per}}(\omega) := v_n \Phi(R_{n,\omega}) = \sum_{\tau \in \mathrm{D}(\omega_{n,\mathrm{per}}) :\, c(\tau) \in \Lambda_n} \varphi(\tau)\,. \tag{3.5}$$

(The last equation follows from (2.7).) By definition, we have $H_{n,\mathrm{per}}(\emptyset) = 0$. Applying Lemma 2.1 and using (2.10), we see that $|H_{n,\mathrm{per}}| \le v_n c_\varphi\, 2\, z(R_n) = 2c_\varphi\, N_n$. The following result will be proved in Subsection 4.1.

Proposition 3.6. *For every $P \in \mathscr{P}_\Theta$ we have*

$$\lim_{n\to\infty} v_n^{-1} \int H_{n,\mathrm{per}}\, dP = \Phi(P)\,.$$

Moreover, if P is tempered then

$$\lim_{n\to\infty} v_n^{-1} \int H_{n,\omega}(\omega)\, P(d\omega) = \Phi(P)\,.$$

Finally we turn to the pressure. Let

$$Z_{n,z,\mathrm{per}} = \int e^{-H_{n,\mathrm{per}}}\, d\Pi_n^z$$

be the partition function in Λ_n with periodic boundary condition.

Proposition 3.7. *For each $z > 0$, the pressure*

$$p(z,\varphi) := \lim_{n\to\infty} v_n^{-1} \log Z_{n,z,\mathrm{per}}$$

exists and satisfies

$$p(z,\varphi) = -\min_{P\in\mathscr{P}_\Theta} \left[I_z(P) + \Phi(P) \right]. \tag{3.6}$$

Proof: This is a direct consequence of Theorem 3.1 of [14] because Φ is continuous by Proposition 3.1. $\Diamond$

A counterpart for the partition functions with configurational boundary conditions follows later in Proposition 4.11. Let us conclude with some remarks on extensions and further results.

Remark 3.8. *Large deviations.* The following large deviation principle is valid. For every measurable $A \subset \mathscr{P}_\Theta$,

$$\limsup_{n\to\infty} v_n^{-1} \log G_{n,z,\mathrm{per}}(R_n \in A) \le -\inf I_{z,\varphi}(\mathrm{cl} A)$$

and

$$\liminf_{n\to\infty} v_n^{-1} \log G_{n,z,\mathrm{per}}(R_n \in A) \ge -\inf I_{z,\varphi}(\mathrm{int} A)\,,$$

where $G_{n,z,\mathrm{per}} = Z_{n,z,\mathrm{per}}^{-1} e^{-H_{n,\mathrm{per}}} \Pi_n^z$ is the Gibbs distribution in Λ_n with periodic boundary condition, $I_{z,\varphi} = I_z + \Phi + p(z,\varphi)$ is the excess free energy density, and the closure cl and the interior int are taken in the topology $\mathscr{T}_{\mathscr{L}}$. Since φ is bounded so that Φ is continuous, this is a direct consequence of Theorem 3.1 of [14].

Remark 3.9. *Marked particles.* Our results can be extended to the case of point particles with marks, that is, with internal degrees of freedom. Let E be any separable metric space, which is equipped with its Borel σ-algebra and a reference measure ν, and $\overline{\Omega}$ the set of all pairs $\bar\omega = (\omega, \sigma_\omega)$ with $\omega \in \Omega$ and $\sigma_\omega \in E^\omega$. In place of the reference Poisson point process Π^z, one takes the Poisson point process $\overline{\Pi}^z$ on $\overline{\Omega}$ with intensity measure $z\lambda \otimes \nu$, where λ is Lebesgue measure on $\mathbb{R}^2$. For $\bar\omega \in \overline{\Omega}$ let

$$\mathsf{D}(\bar\omega) = \left\{ \bar\tau = (\tau, \sigma_\tau) : \tau \in \mathsf{D}(\omega),\ \sigma_\tau = \sigma_\omega|_\tau \right\}.$$

Of course, the centre, radius and circumscribed disc of a marked triangle $\bar{\tau}$ are still defined in terms of the underlying τ. In the germ-grain representation, $\mathscr{T}_0$ is replaced by the set $\overline{\mathscr{T}_0}$ of all centred $\bar{\tau}$. The tile distribution $\mu_{\bar{P}}$ of a stationary point process $\bar{P}$ on $\overline{\Omega}$ is a finite measure on $\overline{\mathscr{T}_0}$ and is defined by placing bars in (2.3). A triangle potential is a bounded function φ on $\overline{\mathscr{T}_0}$. Such a φ is eventually increasing if $\varphi(\bar{\tau}) = \psi(\varrho(\tau))$ for some nondecreasing ψ when $\varrho(\tau)$ is large enough. It is then easily seen that all our arguments carry over to this setting without change.

Remark 3.10. *Particles with hard core.* There is some interest in the case when the particles are required to have at least some distance $r_0 > 0$. This is expressed by adding to the Hamiltonian (3.1) a hard-core pair interaction term $H^{\mathrm{hc}}_{\Lambda,\omega}(\zeta)$ which is equal to ∞ if $|x-y| \le r_0$ for a pair $\{x,y\} \subset \zeta_\Lambda \cup \omega_{\Lambda^c}$ with $\{x,y\} \cap \Lambda \ne \emptyset$, and zero otherwise. Equivalently, one can replace the configuration space Ω by the space

$$\Omega^{\mathrm{hc}} = \big\{\omega \in \Omega : |x-y| > r_0 \text{ for any two distinct } x,y \in \omega\big\}$$

of all hard-core configurations. The free energy functional on $\mathscr{P}_\Theta$ then takes the form $F^{\mathrm{hc}}_z := I_z + \Phi + \Phi^{\mathrm{hc}}$, where

$$\Phi^{\mathrm{hc}}(P) = \infty\ P^0\big(\omega : 0 < |x| \le r_0 \text{ for some } x \in \omega\big) = \infty\ P(\Omega \setminus \Omega^{\mathrm{hc}})$$

for $P \in \mathscr{P}_\Theta$ with Palm measure P^0; here we use the convention $\infty 0 = 0$. We claim that our results can also be adapted to this setting. In particular, the minimisers of F^{hc}_z are Gibbsian for z and the combined triangle and hard-core pair interaction, and the tempered Gibbs measures for this interaction minimise F^{hc}_z. We will comment on the necessary modifications in Remarks 4.4 and 4.12.

Combining the extensions in the last two remarks we can include the following example of phase transition.

Example 3.11. *The Delaunay-Potts hard-core model for particles with $q \ge 2$ colours.* In the setup of Remark 3.9 we have $E = \{1,\dots,q\}$, and the triangle potential is

$$\varphi(\bar{\tau}) = \begin{cases} \beta & \text{if } \varrho(\tau) \le r_1 \text{ and } \sigma_\tau \text{ is not constant,} \\ 0 & \text{otherwise,} \end{cases}$$

where $\beta > 0$ is the inverse temperature and $r_1 > 0$ is an arbitrary interaction radius. If one adds a hard-core pair interaction with range $r_0 < r_1/\sqrt{2}$ as in Remark 3.10, this model is similar to the model considered in [4]. (Instead of a triangle potential, these authors consider an edge potential along the Delaunay edges that do not belong to a tile τ of radius $\varrho(\tau) > r_1$.) Using a random cluster representation of the triangle interaction as in [15] and replacing edge percolation by tile percolation one finds that the methods of [4] can be adapted to the present model. Consequently, if z and β are sufficiently large, then the simplex $\mathscr{M}_\Theta(z,\varphi) = \mathscr{G}_\Theta(z,\varphi)$ has at least q distinct extreme points.

4 Proofs

4.1 Energy and free energy

We begin with the proof of Proposition 2.2, which states that the centred tile distribution μ_P depends continuously on P. The continuity of the energy density Φ and the lower semicontinuity of the free energy density $I_z + \Phi$ then follow immediately.

Proof of Proposition 2.2: Let (P_α) be a net in $\mathscr{P}_\Theta$ that converges to some $P \in \mathscr{P}_\Theta$. We need to show that $\int g\,d\mu_{P_\alpha} \to \int g\,d\mu_P$ for all $g \in \mathscr{L}_0$. We can assume without loss of generality that $0 \le g \le 1$.

We first consider the case that g has bounded support, in that $g \le \mathbb{1}\{\varrho \le r\}$ for some $r > 0$. Let $\Delta \subset \mathbb{R}^2$ be any bounded set of Lebesgue measure $|\Delta| = 1$. Also, let

$$f(\omega) = \sum_{\tau \in \mathrm{D}(\omega)} g(\tau - c(\tau))\,\mathbb{1}_\Delta(c(\tau)).$$

In view of (2.3), we then have $\int g\,d\mu_Q = \int f\,dQ$ for all $Q \in \mathscr{P}_\Theta$, and in particular for $Q = P_\alpha$ and $Q = P$. By the bounded support property of g, f depends only on the configuration in the r-neigbourhood $\Delta^r := \{y \in \mathbb{R}^2 : |y - x| \le r \text{ for some } x \in \Delta\}$ of Δ. That is, f is measurable with respect to $\mathscr{F}_{\Delta^r}$. Moreover, $f \le 2N_{\Delta^r}$ by Lemma 2.1, so that $f \in \mathscr{L}$. In the present case, the result thus follows from the definition of the topology $\mathscr{T}_{\mathscr{L}}$.

If g fails to be of bounded support, we can proceed as follows. Let $\varepsilon > 0$ be given and $r > 0$ be so large that $\mu_P(\varrho > r) < \varepsilon$. Since $\|\mu_{P_\alpha}\| = 2z(P_\alpha) \to 2z(P) = \|\mu_P\|$ and $\mu_{P_\alpha}(\varrho \le r) \to \mu_P(\varrho \le r)$ by the argument above, we have $\mu_{P_\alpha}(\varrho > r) \to \mu_P(\varrho > r)$. We can therefore assume without loss of generality that $\mu_{P_\alpha}(\varrho > r) < \varepsilon$ for all α. Using again the first part of this proof, we can thus write

$$\begin{aligned}
&\int g\,d\mu_P - \varepsilon \le \int g\,\mathbb{1}_{\{\varrho \le r\}}\,d\mu_P = \lim_\alpha \int g\,\mathbb{1}_{\{\varrho \le r\}}\,d\mu_{P_\alpha} \\
&\le \liminf_\alpha \int g\,d\mu_{P_\alpha} \le \limsup_\alpha \int g\,d\mu_{P_\alpha} \\
&\le \lim_\alpha \int g\,\mathbb{1}_{\{\varrho \le r\}}\,d\mu_{P_\alpha} + \varepsilon = \int g\,\mathbb{1}_{\{\varrho \le r\}}\,d\mu_P + \varepsilon \le \int g\,d\mu_P + \varepsilon\,.
\end{aligned}$$

Since ε was chosen arbitrarily, the result follows. $\Diamond$

We now turn to the properties of the free energy density.

Proof of Proposition 3.1: As φ belongs to $\mathscr{L}_0$, the continuity of Φ follows immediately from Proposition 2.2. By Lemma 2.3, we can also conclude that $I_z + \Phi$ is lower semicontiuous. Moreover, hypothesis (2.10) implies that the level set $\{I_z + \Phi \le c\}$ is contained in $\{I_z \le c + 2c_\varphi z(\cdot)\}$, which by (2.9) coincides with the compact set $\{I_{z'} \le c + z' - 1\}$ for $z' = z\exp(2c_\varphi)$.

Let $P = \delta_\emptyset \in \mathscr{P}_\Theta$ be the Dirac measure at the empty configuration. Then $\mu_P \equiv 0$ and thus $\Phi(P) = 0$. On the other hand, $I_z(P) = z$. This means that $I_z + \Phi$ is not identically equal to $+\infty$ on $\mathscr{P}_\Theta$ and thus, by the compactness of its level sets, attains its infimum. To see that the minimisers form a face of $\mathscr{P}_\Theta$, it is sufficient to note that $I_z + \Phi$ is measure affine; cf. Theorem (15.20) of [10]. $\Diamond$

Next we show that the minimisers of the free energy are nondegenerate.

Proof of Proposition 3.2: The second statement follows from the first because $\mathscr{M}_\Theta(z,\varphi)$ is a face of $\mathscr{P}_\Theta$. For, suppose there exists some $P \in \mathscr{M}_\Theta(z,\varphi)$ with $P(\{\emptyset\}) > 0$. Then $\delta_\emptyset$ appears in the ergodic decomposition of P, which would only be possible if $\delta_\emptyset \in \mathscr{M}_\Theta(z,\varphi)$.

To prove the first statement we note that $\Phi(\Pi^u) \le c_\varphi \|\mu_{\Pi^u}\| = 2c_\varphi u$ for all $u > 0$. Therefore, if $z > 0$ is given and u is small enough then

$$I_z(\Pi^u) + \Phi(\Pi^u) \le z - u + u\log(u/z) + 2c_\varphi u < z = I_z(\delta_\emptyset) + \Phi(\delta_\emptyset), \tag{4.1}$$

so that $\delta_\emptyset$ is no minimiser of the free energy. ◊

Finally we show that the energy density Φ is the infinite volume limit of the mean energy per volume.

Proof of Proposition 3.6: We begin with the case of periodic boundary conditions. For every $P \in \mathscr{P}_\Theta$, we have $v_n^{-1} \int H_{n,\mathrm{per}}\, dP = \int \Phi(R_n)\, dP = \Phi(PR_n)$. It is easy to see that $PR_n \to P$, cf. Remark 2.4 of [14]. Since Φ is continuous, it follows that $\Phi(PR_n) \to \Phi(P)$.

Next we consider the case of configurational boundary conditions and suppose that P is tempered. Applying (2.3) we obtain for each n

$$\begin{aligned}
\int P(d\omega) H_{n,\omega}(\omega) &= \int P(d\omega) \sum_{\tau \in \mathrm{D}(\omega)} \varphi(\tau - c(\tau))\, \mathbb{1}_{\{B(\tau - c(\tau)) \cap (\Lambda_n - c(\tau)) \neq \emptyset\}} \\
&= \int \mu_P(d\tau)\, \varphi(\tau) \int dx\, \mathbb{1}_{\{B(\tau) \cap (\Lambda_n - x) \neq \emptyset\}} \\
&= \int \mu_P(d\tau)\, \varphi(\tau) \left|\Lambda_n^{\varrho(\tau)}\right|,
\end{aligned} \tag{4.2}$$

where $\Lambda_n^{\varrho(\tau)}$ is the $\varrho(\tau)$-neigbourhood of Λ_n. Now, for each τ we have

$$\left|\Lambda_n^{\varrho(\tau)}\right| / v_n = 1 + 4\varrho(\tau)/\sqrt{v_n} + \pi\varrho(\tau)^2/v_n \to 1 \text{ as } n \to \infty.$$

In view of (2.4) and (2.10), we can apply the dominated convergence theorem to conclude that

$$\Phi(P) = \lim_{n\to\infty} v_n^{-1} \int H_{n,\omega}(\omega)\, P(d\omega),$$

as desired. ◊

4.2 The variational principle: first part

In this section we shall prove that each minimiser of the free energy is a Delaunay-Gibbs measure. We start with an auxiliary result on the 'range of influence' of the boundary condition on the events within a bounded set. Let $\Delta \subset \mathbb{R}^2$ be a bounded Borel set and $\omega \in \Omega$. Writing $B(x, r)$ for the open disc in $\mathbb{R}^2$ with center x and radius r, we let

$$\mathscr{R}_\Delta(\omega) = \left\{r > 0 :\ \omega_{B(x,r)\setminus\Delta} = \emptyset \text{ for some } x \in \mathbb{R}^2 \text{ s.t. } B(x,r) \cap \Delta \neq \emptyset\right\}$$

be the set of possible radii of circumcircles hitting Δ in the Delaunay triangulation of a configuration of the form $\zeta_\Delta \cup \omega_{\Delta^c}$ with arbitrary ζ, and $\mathbf{r}_\Delta(\omega) = \sup \mathscr{R}_\Delta(\omega)$. Let $\Delta^{2r} = \bigcup_{x\in\Delta} B(x, 2r)$ be the open $2r$-neigbourhood of Δ. We then observe the following.

Lemma 4.1. *(a) For all $r > 0$, $\{\mathbf{r}_\Delta < r\} \in \mathscr{F}_{\Delta^{2r}\setminus\Delta}$. In particular, $\mathbf{r}_\Delta$ is $\mathscr{F}_{\Delta^c}$-measurable.*
(b) For all $P \in \mathscr{P}_\Theta$ we have $P(\mathbf{r}_\Delta = \infty) = P(\{\emptyset\})$.

Proof: (a) Let $\tilde{\mathscr{R}}_\Delta(\omega)$ be defined as $\mathscr{R}_\Delta(\omega)$, except that x is required to belong to $\mathbb{Q}^2$. Then $\tilde{\mathscr{R}}_\Delta(\omega) \subset \mathscr{R}_\Delta(\omega)$. Moreover, if $r \in \mathscr{R}_\Delta(\omega)$ then $]0, r[\subset \tilde{\mathscr{R}}_\Delta(\omega)$, so that $\mathbf{r}_\Delta(\omega) = \sup \tilde{\mathscr{R}}_\Delta(\omega)$. I follows that

$$\{\mathbf{r}_\Delta < r\} = \bigcup_{s \in \mathbb{Q}:\, 0<s<r} \; \bigcap_{x \in \mathbb{Q}^2:\, B(x,s)\cap\Delta \neq \emptyset} \{N_{B(x,s)\setminus\Delta} \geq 1\}\,,$$

and the last set certainly belongs to $\mathscr{F}_{\Delta^{2r}\setminus\Delta}$.
(b) Since $\{\emptyset\} = \bigcap_{r\in\mathbb{N}}\{N_{B(0,r)} = 0\} \subset \{\mathbf{r}_\Delta = \infty\}$, it is sufficient to prove that $\{\mathbf{r}_\Delta = \infty, N_{B(0,r)} \geq 1\}$ has measure zero for all $r > 0$ and $P \in \mathscr{P}_\Theta$. Now, if $\{\mathbf{r}_\Delta = \infty\}$ occurs then Δ is hit by a sequence of discs B_n with $N_{B_n} = 0$ and $\mathrm{diam}(B_n) \to \infty$. Select points $z_n \in B_n \cap \Delta$. Passing to a subsequence one can assume that the sequence (z_n) converges to some z. Let ψ_n be the direction from z to the centre of B_n and ψ an accumulation point of the sequence (ψ_n). Finally, let C be a cone with apex z, direction ψ and opening angle less than π. It is then clear $N_C = 0$. On the other hand, Poincaré's recurrence theorem implies that $N_C = \infty$ almost surely on $\{N_{B(0,r)} \geq 1\}$. This contradiction gives the desired result. $\Diamond$

As an immediate consequence we obtain that each nondegenerate stationary point process is concentrated on the set Ω^* of admissible configurations.

Corollary 4.2. *The set Ω^* of admissible configurations is measurable (in fact, shift invariant and tail measurable), and $P(\Omega^*) = 1$ for all $P \in \mathscr{P}_\Theta$ with $P(\{\emptyset\}) = 0$.*

Proof: This is immediate from Lemma 4.1 because $\Omega^* = \bigcap_{n\geq 1}\{\mathbf{r}_{\Lambda_n} < \infty\}$. $\Diamond$

Next we state a consequence of Proposition 3.7.

Corollary 4.3. *For every $P \in \mathscr{M}_\Theta(z, \varphi)$, we have*

$$\lim_{n\to\infty} v_n^{-1} I_n(P; G_{n,z,\mathrm{per}}) = 0\,.$$

Proof: By the definition of relative entropy,

$$I_n(P; G_{n,z,\mathrm{per}}) = I_n(P; \Pi^z) + \int H_{n,\mathrm{per}}\, dP + \log Z_{n,z,\mathrm{per}}\,.$$

Together with Propositions 3.6 and 3.7, this gives the result. $\Diamond$

We are now ready to show that the minimisers of $I_z + \Phi$ are Gibbsian.

Proof of Theorem 3.3, first part: We follow the well-known scheme of Preston [19] (in the variant used in [13], Section 7). Let $P \in \mathscr{M}_\Theta(z, \varphi)$, f be a bounded local function, Δ a bounded Borel set, and

$$f_\Delta(\omega) = \int f(\zeta)\, G_{\Delta,z,\omega}(d\zeta)\,, \quad \omega \in \Omega.$$

We need to show that $\int f\,dP = \int f_\Delta\,dP$. Let $\mathbf{r}_\Delta$ be the range function defined above, and for each $r > 0$ let $\mathbb{1}_{\Delta,r} = \mathbb{1}\{\mathbf{r}_\Delta < r\}$ and Δ^{2r} be the $2r$-neigbourhood of Δ. By Lemma 4.1 (a), $\mathbb{1}_{\Delta,r}$ is measurable with respect to $\mathscr{F}_{\Delta^{2r}\setminus\Delta}$. Moreover, if $\mathbf{r}_\Delta(\omega) < r$ then $H_{\Delta,\omega}(\zeta) = H_{\Delta,\omega_{\Delta^{2r}\setminus\Delta}}(\zeta)$ for all $\zeta \in \Omega_\Delta$. So, if r is so large that f is $\mathscr{F}_{\Delta^{2r}}$-measurable, we can conclude that $\mathbb{1}_{\Delta,r} f_\Delta$ is $\mathscr{F}_{\Delta^{2r}\setminus\Delta}$-measurable.

Now we apply Corollary 4.3, which states that $\lim_{n\to\infty} v_n^{-1} I_{\Lambda_n}(P; G_{n,z,\mathrm{per}}) = 0$. By shift invariance, this implies that $P_\Lambda \ll G_{\Lambda,z,\mathrm{per}}$ with a density g_Λ for each sufficiently large square Λ. In particular, for any $\Delta' \subset \Lambda$ we have $P_{\Delta'} \ll (G_{\Lambda,z,\mathrm{per}})_{\Delta'}$ wih density $g_{\Lambda,\Delta'}(\omega) = \int G_{\Lambda\setminus\Delta',z,\omega\cap\Delta'}(d\zeta)\, g_\Lambda(\zeta)$. Corollary 4.3 implies further that for each $\delta > 0$ there exists a square Λ and a Borel set Δ' with $\Delta^{2r} \subset \Delta' \subset \Lambda$ such that

$$\int |g_{\Lambda,\Delta'} - g_{\Lambda,\Delta'\setminus\Delta}|\, dG_{\Lambda,z,\mathrm{per}} < \delta;$$

cf. Lemma 7.5 of [13]. Now we consider the difference

$$\int \mathbb{1}_{\Delta,r}\,(f - f_\Delta)\,dP = \int \mathbb{1}_{\Delta,r}\big(g_{\Lambda,\Delta'} f - g_{\Lambda,\Delta'\setminus\Delta} f_\Delta\big)\, dG_{\Lambda,z,\mathrm{per}}\,.$$

Since $G_{\Lambda,z,\mathrm{per}} = \int G_{\Lambda,z,\mathrm{per}}(d\omega)\, G_{\Delta,z,\omega}$ and $\mathbb{1}_{\Delta,r}\, g_{\Lambda,\Delta'\setminus\Delta}$ is $\mathscr{F}_{\Lambda\setminus\Delta}$-measurable, we can conclude that

$$\int \mathbb{1}_{\Delta,r}\, g_{\Lambda,\Delta'\setminus\Delta} f_\Delta\, dG_{\Lambda,z,\mathrm{per}} = \int \mathbb{1}_{\Delta,r}\, g_{\Lambda,\Delta'\setminus\Delta} f\, dG_{\Lambda,z,\mathrm{per}}\,.$$

By the choice of Λ and Δ', we can replace the density $g_{\Lambda,\Delta'\setminus\Delta}$ in the last expression by $g_{\Lambda,\Delta'}$ making an error of at most δ. We thus find that $|\int \mathbb{1}_{\Delta,r}\,(f - f_\Delta)\,dP\,| < \delta$. Letting $\delta \to 0$ and $r \to \infty$, we finally obtain by the dominated convergence theorem that

$$\int_{\{\mathbf{r}_\Delta < \infty\}} (f - f_\Delta)\,dP = 0\,.$$

This completes the proof because $P(\mathbf{r}_\Delta = \infty) = P(\{\emptyset\}) = 0$ by Lemma 4.1(b) and Proposition 3.2. $\Diamond$

Remark 4.4. Here are some comments on the necessary modifications in the hard-core setup of Remark 3.10. In analogy to Proposition 3.7, one needs that

$$\lim_{n\to\infty} v_n^{-1} \log \int e^{-H_{n,\mathrm{per}} - H^{\mathrm{hc}}_{n,\mathrm{per}}}\, d\Pi^z_n = -\min_{P\in\mathscr{P}_\Theta} \big[I_z(P) + \Phi(P) + \Phi^{\mathrm{hc}}(P)\big]\,.$$

This follows directly from Propositions 4.1 and 5.4 of [12] because Φ is continuous. Corollary 4.3 therefore still holds for the periodic Gibbs distributions with additional hard-core pair interaction. One also needs to modify the proof of Proposition 3.2, in that the Poisson processes Π^u should be replaced by the Gibbs measure P^u with activity u and pure hard-core interaction. P^u is defined as the limit of the Gibbs distributions $G^{\mathrm{hc}}_{n,u,\mathrm{per}}$ for the periodic hard-core Hamiltonians $H^{\mathrm{hc}}_{n,\mathrm{per}}$. By Proposition 7.4 of [12], P^u exists and satisfies

$$I_u(P^u) = -\lim_{n\to\infty} v_n^{-1} \log \int e^{-H^{\mathrm{hc}}_{n,\mathrm{per}}}\, d\Pi^u_n \le -\lim_{n\to\infty} v_n^{-1} \log \Pi^u_n(\{\emptyset\}) = u\,.$$

Together with (2.9) we find that

$$I_z(P^u) + \Phi(P^u) \le z + z(P^u)\left[\log(u/z) + 2c_\varphi\right],$$

which is strictly less than z when u is small enough. Since $\Phi^{\mathrm{hc}}(P^u) = \Phi^{\mathrm{hc}}(\delta_\emptyset) = 0$, it follows that the minimisers of $I_z + \Phi + \Phi^{\mathrm{hc}}$ are non-degenerate. No further changes are required for the proof of the first part of the variational principle.

4.3 Boundary estimates

We now work towards a proof of the reverse part of the variational principle. In this section, we control the boundary effects that determine the difference of $H_{n,\mathrm{per}}$ and $H_{n,\omega}$. The resulting estimates will be crucial for the proof of Proposition 4.11. For every $\omega \in \Omega$ and every Borel set Δ let

$$\mathsf{S}_\Delta(\omega) = \left\{\tau \in \mathsf{D}(\omega) : B(\tau) \cap \Delta \ne \emptyset \text{ and } B(\tau) \setminus \Delta \ne \emptyset\right\}$$

be the set of all triangles $\tau \in \mathsf{D}(\omega)$ for which $B(\tau)$ crosses the boundary of Δ. We start with a lemma that controls the influence on S_Δ when two configurations are pasted together.

Lemma 4.5. *Let Δ be a (not necessarily bounded) Borel set in $\mathbb{R}^2$, $\zeta \in \Omega^* \cup \{\emptyset\}$ a configuraton with $\zeta_{\partial\Delta} = \emptyset$, and $\omega \in \Omega$. Then for each $\tau \in \mathsf{S}_\Delta(\zeta_\Delta \cup \omega_{\Delta^c})$ and each $x \in \tau \cap \Delta$ there exists some $\tau' \in \mathsf{S}_\Delta(\zeta)$ with $x \in \tau'$.*

Proof: Let Δ, ζ and ω be given. If ζ is empty, there exists no $x \in \tau_\Delta \subset \zeta_\Delta$, so that the statement is trivially true. So let $\zeta \in \Omega^*$ and suppose there exists some $\tau \in \mathsf{S}_\Delta(\zeta_\Delta \cup \omega_{\Delta^c})$ with $\tau_\Delta \ne \emptyset$. Let $x \in \tau_\Delta$. Since $\zeta_{\partial\Delta} = \emptyset$, x does in fact belong to the interior of Δ. This implies that $B(\tau') \cap \Delta \ne \emptyset$ for each $\tau' \in \mathsf{D}(\zeta)$ containing x. Therefore we only need to show that $B(\tau') \setminus \Delta \ne \emptyset$ for at least one such τ'. Suppose the contrary. Then $B(\tau') \subset \Delta$ whenever $x \in \tau' \in \mathsf{D}(\zeta)$. This means that the Delaunay triangles containing x are completely determined by ζ_Δ. This gives the contradiction

$$\emptyset \ne \{\tau \in \mathsf{S}_\Delta(\zeta_\Delta \cup \omega_{\Delta^c}) : \tau \ni x\} = \{\tau' \in \mathsf{S}_\Delta(\zeta) : \tau' \ni x\} = \emptyset,$$

and the proof is complete. $\Diamond$

The following proposition is the fundamental boundary estimate. It bounds the difference of Hamiltonians with periodic and configurational boundary conditions in terms of $S_n := \operatorname{card} \mathsf{S}_{\Lambda_n}$.

Proposition 4.6. *There exists a universal constant $\gamma < \infty$ such that*

$$|H_{n,\mathrm{per}}(\zeta) - H_{n,\omega}(\zeta)| \le \gamma c_\varphi \left(S_n(\omega) + S_n(\zeta)\right)$$

for all $n \ge 1$ and all $\zeta, \omega \in \Omega^* \cup \{\emptyset\}$ with $\zeta_{\partial\Lambda_n} = \omega_{\partial\Lambda_n} = \emptyset$.

Proof: Let n, ζ, ω be fixed and

$$\mathsf{A} = \left\{\tau \in \mathsf{D}(\zeta_{\Lambda_n} \cup \omega_{\Lambda_n^c}) : B(\tau) \cap \Lambda_n \ne \emptyset\right\}, \quad \mathsf{B} = \left\{\tau \in \mathsf{D}(\zeta_{n,\mathrm{per}}) : c(\tau) \in \Lambda_n\right\}.$$

In view of (3.1) and (3.5) we have $H_{n,\omega}(\zeta) = \sum_{\tau \in \mathsf{A}} \varphi(\tau)$ and $H_{n,\mathrm{per}}(\zeta) = \sum_{\tau \in \mathsf{B}} \varphi(\tau)$. Since φ is bounded by c_φ, we only need to estimate the cardinalities of $\mathsf{A} \setminus \mathsf{B}$ and $\mathsf{B} \setminus \mathsf{A}$. We note that

$\mathsf{A}\setminus\mathsf{B}\subset\mathsf{S}_{\Lambda_n}(\zeta_{\Lambda_n}\cup\omega_{\Lambda_n^c})$ and $\mathsf{B}\setminus\mathsf{A}\subset\mathsf{S}_{\Lambda_n}(\zeta_{n,\mathrm{per}})$. So we can apply Lemma 4.5 to both $\Delta=\Lambda_n$ and $\Delta=\Lambda_n^c$ to obtain that the set of points belonging to a triangle in $\mathsf{A}\setminus\mathsf{B}$ is contained in the set of points belonging to a triangle of $\mathsf{S}_{\Lambda_n}(\zeta)\cup\mathsf{S}_{\Lambda_n}(\omega)$. Hence, $\mathrm{card}\left(\bigcup_{\tau\in\mathsf{A}\setminus\mathsf{B}}\tau\right)\leq 3(S_n(\zeta)+S_n(\omega))$. By Lemma 2.1, it follows that $\mathrm{card}(A\setminus B)\leq 6(S_n(\omega)+S_n(\zeta))$.

To estimate the cardinality of $\mathsf{B}\setminus\mathsf{A}$ we may assume that $\zeta_n\neq\emptyset$. The periodic continuation $\zeta_{n,\mathrm{per}}$ then contains a lattice, and this implies that every triangle of $\mathsf{D}(\zeta_{n,\mathrm{per}})$ has a circumscribed disc of diameter at most $\sqrt{2v_n}$. Hence, each $\tau\in\mathsf{B}\setminus\mathsf{A}$ is contained in Λ_{5n+2}, the union of 5^2 translates of Λ_n (up to their boundaries). Applying Lemma 4.5 to each of these translates we conclude that the number of points that belong to a triangle of $\mathsf{B}\setminus\mathsf{A}$ is bounded by $3\cdot 5^2 S_n(\zeta)$. Using Lemma 2.1 again we find that $\mathrm{card}(B\setminus A)\leq 150(S_n(\omega)+S_n(\zeta))$, and the result follows with $\gamma=156$. $\Diamond$

The following immediate corollary will be needed in the proof of Theorem 3.3.

Corollary 4.7. *There exists a constant $C<\infty$ such that $|H_{n,\omega}(\emptyset)|\leq C\,S_n(\omega)$ for all $n\geq 1$ and $\omega\in\Omega^*$ with $\omega_{\partial\Lambda_n}=\emptyset$.*

The next proposition exhibits the fundamental role of the temperedness condition (2.4) combined with stationarity for controlling the boundary effects.

Proposition 4.8. *For every $P\in\mathscr{P}_\Theta^{\mathrm{tp}}$, $v_n^{-1}S_n\to 0$ in $L^1(P)$ and P-almost surely.*

Proof: For each $i\in\mathbb{R}^2$ we consider the shifted unit square $C(i)=\Lambda_0+i$ and define the random variable

$$Z_i=\mathrm{card}\left\{\tau\in\mathsf{D}(\cdot):B(\tau)\cap C(i)\neq\emptyset\right\}.$$

Then

$$S_n\leq\sum_{i\in I_n\setminus I_{n-1}}Z_i\,,\tag{4.3}$$

where $I_n=\Lambda_n\cap(\mathbb{Z}^2+(\frac{1}{2},\frac{1}{2}))$. Note that $\mathrm{card}\,I_n=v_n$. As in (4.2) we have

$$\int Z_0\,dP=\int\mu_P(d\tau)\left|\Lambda_0^{\varrho(\tau)}\right|=\int\mu_P(d\tau)\,\left(1+4\varrho(\tau)+\pi\varrho(\tau)^2\right).$$

The last term is finite by the temperedness of P. So, each Z_i is P-integrable. Since $Z_i=Z_0\circ\vartheta_i$, the two-dimensional ergodic theorem implies that $v_n^{-1}\sum_{i\in I_n}Z_i$ converges to a finite limit $\bar{Z}$, both P-almost surely and in $L^1(P)$. This implies that $v_n^{-1}\sum_{i\in I_n\setminus I_{n-1}}Z_i$ tends to zero P-almost surely and in $L^1(P)$. The result thus follows from (4.3). $\Diamond$

4.4 Temperedness and block average approximation

Our first result in this subsection is a sufficient condition for temperedness in terms of vacuum probabilities. For $P\in\mathscr{P}_\Theta$ let

$$V_k(P)=\operatorname{ess\,sup}P(N_{\Lambda_k}=0|\mathscr{F}_{\Lambda_k^c})\tag{4.4}$$

be the essential supremum of the conditional probability that Λ_k contains no particle given the configuration outside.

Proposition 4.9. *Every $P \in \mathscr{P}_\Theta \setminus \{\delta_\emptyset\}$ satisfying*

$$\sum_{k \geq 0} v_k \, V_k(P) < \infty \tag{4.5}$$

is tempered.

Proof: For each $k \geq 1$ we consider the shifted squares $\Lambda_k(i) = \Lambda_k + (2k+1)i$, $i \in \mathbb{Z}^2$, as well as the event

$$A_k = \{N_{\Lambda_k(i)} \neq 0 \text{ for all } i \in \mathbb{Z}^2 \text{ with } \|i\|_\infty \leq 1\}.$$

Since $P \neq \delta_\emptyset$, it is clear that $P(A_k) \to 1$ as $k \to \infty$. Thus we can write

$$\begin{aligned}&\int \operatorname{card}\big(\tau \in \mathsf{D}(\omega) : 0 \in B(\tau)\big)\, P(d\omega)\\ &\qquad \leq \sum_{k \geq 1} \int_{A_k \setminus A_{k-1}} \operatorname{card}\big(\tau \in \mathsf{D}(\omega) : 0 \in B(\tau)\big) P(d\omega)\end{aligned}$$

with the convention $A_0 = \emptyset$. Now, if $\omega \in A_k$ then each circumscribed disc containing 0 of a triangle $\tau \in \mathsf{D}(\omega)$ has a diameter not larger than $2\sqrt{2v_k}$, so that each such τ in fact belongs to $\mathsf{D}(\omega_{\Lambda_{7k+3}})$. Lemma 2.1 thus shows that the number of such τ is at most $2N_{7k+3}(\omega)$. The last sum is therefore not larger than

$$\sum_{k \geq 1} \int_{A_{k-1}^c} 2N_{7k+3} \, dP \leq 2 \sum_{k \geq 1} \sum_{i \in \mathbb{Z}^2 : \|i\|_\infty \leq 1} \int \mathbb{1}_{\{N_{\Lambda_{k-1}(i)} = 0\}} N_{7k+3} \, dP \,.$$

In view of the stationarity of P, the last integral is bounded by $V_{k-1}(P)\, v_{7k+3}\, z(P) = 7^2\, v_k \, V_{k-1}(P)\, z(P)$. So we arrive at the estimate

$$\int \operatorname{card}\big(\tau \in \mathsf{D}(\omega) : 0 \in B(\tau)\big)\, P(d\omega) \leq 2 \cdot 7^2\, 3^2\, z(P) \sum_{k \geq 1} v_k \, V_{k-1}(P)\,.$$

Together with (2.5) and assumption (4.5), this implies the temperedness of P because $v_k \sim v_{k-1}$ as $k \to \infty$. $\Diamond$

The second result concerns the approximation of stationary measures in terms of tempered ergodic measures. This approximation uses the block average construction first introduced by Parthasarathy for proving that the ergodic measures are dense in $\mathscr{P}_\Theta$; cf. [10, Theorem (14.12)], for example.

Proposition 4.10. *Let $z > 0$ and $Q \in \mathscr{P}_\Theta$ be such that $I_z(Q) + \Phi(Q) < \infty$. Then for each $\varepsilon > 0$ there exists some tempered Θ-ergodic $\hat{Q} \in \mathscr{P}_\Theta$ such that $I_z(\hat{Q}) + \Phi(\hat{Q}) < I_z(Q) + \Phi(Q) + \varepsilon$*

Proof: Let $Q \in \mathscr{P}_\Theta$ be given. We can assume that $Q \neq \delta_\emptyset$ because otherwise we can choose $\hat{Q} = Q$. For $n \geq 1$ let Q_n^{iid} denote the probability measure on Ω relative to which the particle configurations in the blocks $\Lambda_n + (2n+1)i$, $i \in \mathbb{Z}^2$, are independent with identical distribution $Q_n = Q \circ \mathrm{pr}_{\Lambda_n}^{-1}$. (In particular, this means that the boundaries of these blocks contain no particles.) Consider the spatial average

$$Q_n^{\text{iid-av}} = v_n^{-1} \int_{\Lambda_n} Q_n^{\mathrm{iid}} \circ \vartheta_x^{-1} \, dx.$$

It is clear that $Q_n^{\text{iid-av}} \in \mathscr{P}_\Theta$. It is also well-known that $Q_n^{\text{iid-av}}$ is Θ-ergodic; cf. [10, Theorem (14.12)], for example. By an analogue of [10, Proposition (16.34)] or [14, Lemma 5.5], its entropy density satisfies

$$I_z(Q_n^{\text{iid-av}}) \le v_n^{-1} I_{\Lambda_n}(Q;\Pi^z) \le I_z(Q). \tag{4.6}$$

Next, the same argument as in [14, Lemma 5.7] shows that $Q_n^{\text{iid-av}} \to Q$ in the topology $\mathscr{T}_{\mathscr{L}}$. By Proposition 2.2, Φ is continuous. We thus conclude that $\Phi(Q_n^{\text{iid-av}}) \to \Phi(Q)$, whence $\Phi(Q_n^{\text{iid-av}}) < \Phi(P)+\varepsilon$ for large n. It remains to prove that each $Q_n^{\text{iid-av}}$ is tempered. Let n be fixed and $k \ge \ell(2n+1)$ for some $\ell \ge 1$. We claim that $V_k(Q_n^{\text{iid-av}}) \le q^{v_{\ell-1}}$ with $q = Q(N_n = 0)$. Indeed, for each $x \in \Lambda_n$ we have $\Lambda_k + x \supset \Lambda_{n+(\ell-1)(2n+1)}$, and the latter set consists of $v_{\ell-1} = (2\ell-1)^2$ distinct blocks as above. Letting g be any nonnegative $\mathscr{F}_{\Lambda_k^c}$-measurable function and using the independence of block configurations, we thus conclude that

$$\begin{aligned}\int \mathbb{1}_{\{N_k=0\}} g\, dQ_n^{\text{iid-av}} &= v_n^{-1}\int_{\Lambda_n} dx \int dQ_n^{\text{iid}}\, \mathbb{1}_{\{N_{\Lambda_k+x}=0\}}\, g\circ\vartheta_x \\ &\le v_n^{-1}\int_{\Lambda_n} dx \int dQ_n^{\text{iid}}\, \mathbb{1}_{\{N_{n+(\ell-1)(2n+1)}=0\}}\, g\circ\vartheta_x = q^{v_{\ell-1}} \int g\, dQ_n^{\text{iid-av}},\end{aligned}$$

which proves the claim. Now, we have $q<1$ because $Q \ne \delta_\emptyset$. It follows that

$$\sum_{k>2n} v_k\, V_k(Q_n^{\text{iid-av}}) \le \sum_{\ell\ge1} q^{v_{\ell-1}} \sum_{\ell(2n+1)\le k<(\ell+1)(2n+1)} v_k \le C_n \sum_{\ell\ge0} v_\ell\, q^{v_\ell} < \infty$$

for some constant $C_n<\infty$. Together with Proposition 4.9, this gives the temperedness of $Q_n^{\text{iid-av}}$. ◊

4.5 The variational principle: second part

In this section we will complete the proof of the variational principle. The essential ingredient is the following counterpart of Proposition 3.7 involving configurational instead of periodic boundary conditions. We only state the lower bound we need.

Proposition 4.11. *For every $P \in \mathscr{P}_\Theta^{\text{tp}}$ with $P(\{\emptyset\}) = 0$ and P-almost every ω we have*

$$\liminf_{n\to\infty} v_n^{-1} \log Z_{n,z,\omega} \ge p(z,\varphi).$$

Proof: By (3.6) and Lemma 4.10, it is sufficient to show that

$$\liminf_{n\to\infty} v_n^{-1} \log Z_{n,z,\omega} \ge -I_z(Q) - \Phi(Q)$$

for every ergodic $Q \in \mathscr{P}_\Theta^{\text{tp}}$. We can assume without loss that the right-hand side is finite. Now, since $I_z(Q)$ is finite, Q is locally absolutely continuous with repect to Π^z. We fix some $\varepsilon > 0$, let $f_n = dQ_n/d\Pi_n^z$, and consider for every $\omega \in \Omega^*$ the set

$$A_{n,\omega} = \left\{ |H_{n,\omega} - H_{n,\text{per}}|/v_n \le \varepsilon,\ \Phi(R_n) < \Phi(Q)+\varepsilon,\ v_n^{-1}\log f_n < I_z(Q)+\varepsilon \right\}.$$

Then for sufficiently large n we have

$$\begin{aligned} Z_{n,z,\omega} &\geq \int_{A_{n,\omega}} e^{-H_{n,\omega}} f_n^{-1}\, dQ \\ &\geq \int_{A_{n,\omega}} e^{-H_{n,\mathrm{per}}} e^{-v_n \varepsilon} f_n^{-1}\, dQ \\ &\geq e^{-v_n [I_z(Q)+\Phi(Q)+3\varepsilon]}\, Q(A_{n,\omega}). \end{aligned}$$

It is therefore sufficient to show that for P-almost every ω, $Q(A_{n,\omega}) \to 1$ as $n \to \infty$. By the ergodic theorem, $\Phi(R_n)$ converges to $\Phi(Q)$ in Q-probability; cf. Remark 2.4 of [14]. By McMillan's theorem [9; 18], $Q(v_n^{-1} \log f_n < I_z(Q) + \varepsilon) \to 1$ when n tends to infinity. Moreover, Propositions 4.6 and 4.8 imply that, for P-almost all ω, $|H_{n,\omega} - H_{n,\mathrm{per}}|/v_n$ converges to 0 in $L^1(Q)$. This gives the result. $\Diamond$

We can now show that every tempered stationary Gibbs measure minimises the free energy density.

Proof of Theorem 3.3, second part: We follow the argument of [13], Proposition 7.7. Let $P \in \mathscr{G}_\Theta^{\mathrm{tp}}(\varphi, z)$. On each $\mathscr{F}_{\Lambda_n}$, P is absolutely continuous w.r.t. to Π^z with density

$$g_n(\zeta) = \int P(d\omega)\, \frac{dG_{n,z,\omega}}{d\Pi_n^z}(\zeta) = \int P(d\omega)\, e^{-H_{n,\omega}(\zeta)}/Z_{n,z,\omega}\,.$$

Using Jensen's inequality and the Gibbs property of P we thus find that

$$\begin{aligned} I_n(P;\Pi^z) &= \int g_n \log g_n \, d\Pi^z \\ &\leq \int \Pi^z(d\zeta) \int P(d\omega) \frac{dG_{n,z,\omega}}{d\Pi_n^z}(\zeta) \big[-H_{n,\omega}(\zeta) - \log Z_{n,z,\omega} \big] \\ &= -\int P(d\omega) H_{n,\omega}(\omega) - \int P(d\omega)\, \log Z_{n,z,\omega}\,. \end{aligned}$$

Next we divide by v_n and let $n \to \infty$. We know from Proposition 3.6 that $v_n^{-1} \int P(d\omega) H_{n,\omega}(\omega) \to \Phi(P)$. On the other hand, Corollary 4.7 implies that $v_n^{-1} \log Z_{n,z,\omega} \geq -z - C v_n^{-1} S_n(\omega)$. Using Propositions 4.8 and 4.11 together with Fatou's Lemma, we thus find that

$$\liminf_{n\to\infty} v_n^{-1} \int P(d\omega)\, \log Z_{n,z,\omega} \geq p(z,\varphi).$$

Therefore $I_z(P) \leq -\Phi(P) + \min[I_z + \Phi]$, as required. $\Diamond$

Remark 4.12. In the hard-core setting of Remark 3.10, a slight refinement of Proposition 4.10 is needed. Namely, under the additional assumption that $\Phi^{\mathrm{hc}}(Q) = 0$ one needs to achieve that also $\Phi^{\mathrm{hc}}(\hat{Q}) = 0$. To this end we fix an integer $k > r_0/2$ and define Q_n^{iid} in such a way that the particle configurations in the blocks $\Lambda_n + (2n+1)i$, $i \in \mathbb{Z}^2$, are independent with identical distribution Q_{n-k}, rather than Q_n. This means that the blocks are separated by corridors of width $2k > r_0$ that contain no particles. It follows that $\Phi^{\mathrm{hc}}(Q_n^{\text{iid-av}}) = 0$, and it is still true that $\limsup_{n\to\infty} I_z(Q_n^{\text{iid-av}}) \leq I_z(Q)$; cf. [12, Lemma 5.1]. We thus obtain the refined Proposition 4.10 as before.

A similar refinement is required in the proof of Proposition 4.11. One can assume that $\Phi^{\mathrm{hc}}(P)=0$ and $\Phi^{\mathrm{hc}}(Q)=0$, and in the definition of $A_{n,\omega}$ one should introduce an empty corridor at the inner boundary of Λ_n to ensure that $H^{\mathrm{hc}}_{n,\omega}=H^{\mathrm{hc}}_{n,\mathrm{per}}=0$ on $A_{n,\omega}$, see [12, Proposition 5.4] for details. In the proof of the second part of Theorem 3.3, one then only needs to note that $\Phi^{\mathrm{hc}}(P)=0$ when P is a Gibbs measure P for the combined triangle and hard-core pair interaction. The proof of Theorem 3.4 carries over to the hard-core case without any changes.

4.6 Temperedness of Gibbs measures

Here we establish Theorem 3.4. By Proposition 4.9 it is sufficient to show the following.

Proposition 4.13. *Let φ be bounded and eventually increasing, $z>0$, and P be any stationary Gibbs measure for φ and z. Then there exists a constant $C>0$ such that*

$$P(N_k=0|\mathscr{F}_{\Lambda_k^c})\le C\,v_k^{-2} \tag{4.7}$$

for all $k\ge 1$.

To prove this we need an auxiliary result which states that the radii of all circumcircles in the Delaunay tessellation must decrease when a point is added to the configuration. Specifically, let $\omega\in\Omega^*$ and $x\in\mathbb{R}^2\setminus\omega$ be such that $\omega\cup\{x\}$ is in general circular position and x is not collinear with two points of ω. We consider the sets

$$\mathsf{C}_x(\omega):=\mathsf{D}(\omega)\setminus\mathsf{D}(\omega\cup\{x\})=\big\{\tau\in\mathsf{D}(\omega):B(\tau)\ni x\big\}$$

and

$$\mathsf{C}_x^+(\omega):=\mathsf{D}(\omega\cup\{x\})\setminus\mathsf{D}(\omega)=\big\{\tau\in\mathsf{D}(\omega\cup\{x\}):\tau\ni x\big\}.$$

If $\langle\tau\rangle$ denotes the convex hull of a triangle τ,

$$\Delta_x(\omega):=\bigcup_{\tau\in\mathsf{C}_x(\omega)}\langle\tau\rangle=\bigcup_{\tau\in\mathsf{C}_x^+(\omega)}\langle\tau\rangle \tag{4.8}$$

is the region on which the triangulations $\mathsf{D}(\omega)$ and $\mathsf{D}(\omega\cup\{x\})$ differ; see Fig. 1. Up to the point x, the interior $\Delta_x^o(\omega)$ of $\Delta_x(\omega)$ is covered by the discs $B(\tau)$ with $\tau\in\mathsf{C}_x^+(\omega)$, which by definition are free of particles. Consequently, $\Delta_x^o(\omega)$ contains no particle of ω, so that the vertices of each $\tau\in\mathsf{C}_x(\omega)$ belong to the boundary $\partial\Delta_x(\omega)$. Next, Lemma 2.1 shows that

$$\operatorname{card}\mathsf{C}_x^+(\omega)=\operatorname{card}\mathsf{C}_x(\omega)+2\,, \tag{4.9}$$

and for every $\Lambda\ni x$ we have

$$H_{\Lambda,\omega}(\omega\cup\{x\})-H_{\Lambda,\omega}(\omega)=\sum_{\tau\in\mathsf{C}_x^+(\omega)}\varphi(\tau)-\sum_{\tau\in\mathsf{C}_x(\omega)}\varphi(\tau). \tag{4.10}$$

Here is the monotonicity result announced above.

Lemma 4.14. *Under the conditions above, there exist a subset $\mathsf{C}_x'(\omega)\subset\mathsf{C}_x(\omega)$ with $\operatorname{card}\big(\mathsf{C}_x(\omega)\setminus\mathsf{C}_x'(\omega)\big)\le 4$ and an injection I from $\mathsf{C}_x'(\omega)$ to $\mathsf{C}_x^+(\omega)$ such that*

$$\varrho(I(\tau))\le\varrho(\tau)\quad\textit{for all }\tau\in\mathsf{C}_x'(\omega). \tag{4.11}$$

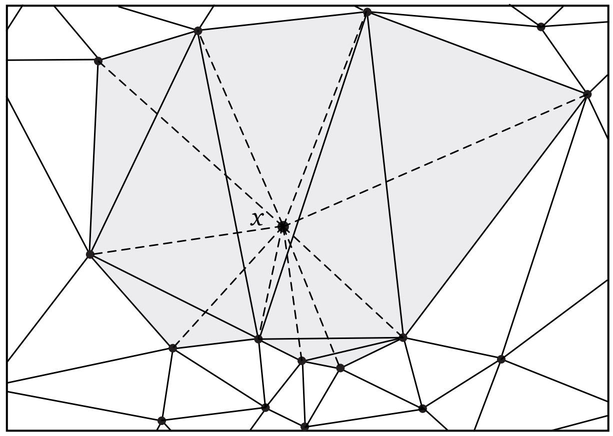

Figure 1: $\mathsf{D}(\omega)$ (solid lines) and $\mathsf{D}(\omega\cup\{x\})\setminus\mathsf{D}(\omega)$ (dashed lines). The difference region $\Delta_x(\omega)$ is shaded in grey.

We postpone the proof of this lemma until the end, coming first to its use.

Proof of Proposition 4.13: By assumption, φ is eventually increasing. So there exists some $r_\varphi < \infty$ and a nondecreasing function ψ such that $\varphi(\tau) = \psi(\varrho(\tau))$ when $\varrho(\tau) \geq r_\varphi$. Combining Lemma 4.14 and Equations (4.9) and (4.10) we thus find that

$$H_{k,\omega}(\omega\cup\{x\}) \leq H_{k,\omega}(\omega) + 10\,c_\varphi \tag{4.12}$$

for all $\omega \in \Omega^*$, $k \geq 1$, and Lebesgue-almost all $x \in \Lambda_k \setminus \omega$ that have at least the distance $2r_\varphi$ from all points of ω. Next, let $P \in \mathscr{G}_\Theta(z,\varphi)$. By definition,

$$P(N_k = 0|\mathscr{F}_{\Lambda_k^c})(\omega) = Z_{k,z,\omega}^{-1}\, e^{-zv_k}\, e^{-H_{k,\omega}(\emptyset)}$$

for all $\omega \in \Omega^*$. Let $\Lambda_k^{(2)} = \big\{(x,y) \in \Lambda_{k-2r_\varphi}^2 : |x-y| \geq 2r_\varphi\big\}$. Applying (4.12) twice (viz. to $\omega_{\Lambda_k^c}$ and x as well as $\omega_{\Lambda_k^c}\cup\{x\}$ and y) and recalling (3.2) we find that

$$\begin{aligned} Z_{k,z,\omega} &\geq e^{-zv_k}\,\frac{z^2}{2}\int_{\Lambda_k^{(2)}} e^{-H_{k,\omega}(\{x\}\cup\{y\})}\,dx\,dy \\ &\geq z^2\,|\Lambda_k^{(2)}|\, e^{-zv_k}\, e^{-H_{k,\omega}(\emptyset)-20\,c_\varphi}/\,2\,. \end{aligned}$$

Since $|\Lambda_k^{(2)}| \sim v_k^2$ as $k \to \infty$, the result follows. $\Diamond$

Finally we turn to the proof of Lemma 4.14.

Proof of Lemma 4.14: Let τ_x be the unique triangle of $\mathsf{C}_x(\omega)$ containing x in its interior, and $\mathsf{C}_x^{+\wedge}(\omega)$ the set of all $\tau \in \mathsf{C}_x^+(\omega)$ that have an acute or right angle at x. Note that $\mathrm{card}(\mathsf{C}_x^+(\omega)\setminus\mathsf{C}_x^{+\wedge}(\omega)) \leq 3$ because the angles at x of all $\tau \in \mathsf{C}_x^+(\omega)$ add up to 360 degrees. We will associate to each triangle $\tau \in \mathsf{C}_x(\omega)$ a triangle $I(\tau) \in \mathsf{C}_x^+(\omega)$, except possibly when $\tau = \tau_x$ or the candidate for $I(\tau)$ does

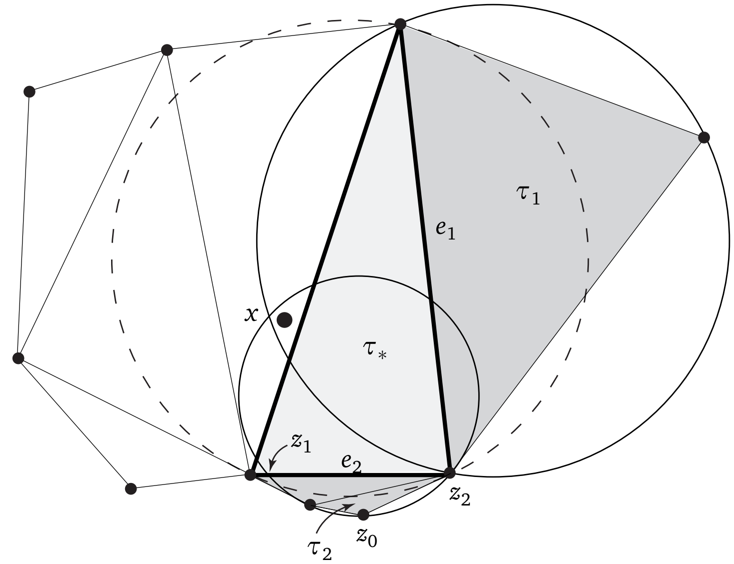

Figure 2: The set $\mathsf{C}_x(\omega)$ for the configuration ω of Fig. 1, with a tile $\tau_* \in \mathsf{C}_x^{(0)}(\omega)$ (light grey), its circumcircle $B(\tau_*)$ (dashed), the associated edges e_i and regions W_i (dark grey), and two triangles $\tau_i \in \mathsf{C}_x(\omega)$ with $\tau_i \subset W_i$ with their circumcircles (solid). The construction in the proof gives $\tilde{I}(\tau_*) = \tau_2$.

not belong to $\mathsf{C}_x^{+\wedge}(\omega)$. Our definition of $I(\tau)$ depends on the number $k = k(\tau)$ of edges $e \subset \tau$ with $\langle e \rangle \subset \partial \Delta_x(\omega)$. Let $\mathsf{C}_x^{(k)}(\omega)$ be the set of all $\tau \in \mathsf{C}_x(\omega)$ that have k such edges. Since $\mathsf{C}_x^{(3)}(\omega) = \emptyset$ except when $\mathsf{C}_x(\omega) = \{\tau_x\}$, we only need to consider the three cases $k = 0, 1, 2$.

The cases $k = 1$ and 2 are easy: For every $\tau \in \mathsf{C}_x^{(1)}(\omega)$ there exists a unique edge $e(\tau)$ such that $e(\tau) \cup \{x\} \in \mathsf{C}_x^{+}(\omega)$. If in fact $e(\tau) \cup \{x\} \in \mathsf{C}_x^{+\wedge}(\omega)$ we set $I(\tau) = e(\tau) \cup \{x\}$; otherwise we leave $I(\tau)$ undefined. Likewise, every $\tau \in \mathsf{C}_x^{(2)}(\omega)$ has two edges $e_1(\tau)$ and $e_2(\tau)$ in $\partial \Delta_x(\omega)$ (in clockwise order, say) and can be mapped to $I(\tau) = e_1(\tau) \cup \{x\}$, provided this triangle belongs to $\mathsf{C}_x^{+\wedge}(\omega)$. The resulting mapping I is clearly injective. Moreover, τ and $I(\tau)$ have the edge $e(\tau)$ (resp. $e_1(\tau)$) in common, and $x \in B(\tau)$ because $\tau \in \mathsf{C}_x(\omega)$. Since $I(\tau) \in \mathsf{C}_x^{+\wedge}(\omega)$ whenever it is defined, we can conclude that $\varrho(I(\tau)) \le \varrho(\tau)$.

The case $k = 0$ is more complicated because the tiles $\tau \in \mathsf{C}_x^{(0)}(\omega)$ are not naturally associated to a tile of $\mathsf{C}_x^{+}(\omega)$. To circumvent this difficulty we define an injection $\tilde{I}$ from $\mathsf{C}_x^{(0)}(\omega) \setminus \{\tau_x\}$ to $\mathsf{C}_x^{(2)}(\omega)$ such that $\varrho(\tilde{I}(\tau)) \le \varrho(\tau)$. Each triangle $\tau \in \mathsf{C}_x^{(0)}(\omega)$ different from τ_x can then be mapped to the triangle $I(\tau) = e_2(\tilde{I}(\tau)) \cup \{x\}$, provided the latter belongs to $\mathsf{C}_x^{+\wedge}(\omega)$; otherwise $I(\tau)$ remains undefined. This completes the construction of I. (Note that τ_x does not necessarily belong to $\mathsf{C}_x^{(0)}(\omega)$. However, if it does we have no useful definition of $\tilde{I}(\tau_x)$.)

To construct $\tilde{I}$ we turn $\mathsf{C}_x(\omega)$ into the vertex set of a graph $G_x(\omega)$ by saying that two tiles are adjacent if they share an edge. The set $\mathsf{C}_x^{(2)}(\omega)$ then coincides with the set of all leaves of $G_x(\omega)$, and $\mathsf{C}_x^{(0)}(\omega)$ is the set of all triple points (= points of degree 3) of $G_x(\omega)$. Consider a fixed $\tau_* \in \mathsf{C}_x^{(0)}(\omega) \setminus \{\tau_x\}$. Since $\tau_* \subset \partial \Delta_x(\omega)$, the set $\Delta_x(\omega) \setminus \langle \tau_* \rangle$ splits into three connected components.

Let $W_i = W_i(\tau_*, x, \omega)$ be the closure of the ith component, $i = 1,2,3$. Any two of these sets intersect at a point of τ_*, and one of them contains x because $\tau_* \neq \tau_x$. Suppose $x \in W_3$. For $i = 1,2$ let $e_i = \tau_* \cap W_i$ be the edge of τ_* that separates W_i from the rest of $\Delta_x(\omega)$; see Fig. 2. We claim that there exists some $i = i(\tau_*) \in \{1,2\}$ such that

$$\varrho(\tau) \leq \varrho(\tau_*) \quad \text{for all } \tau \in \mathsf{C}_x(\omega) \text{ with } \tau \subset W_i\,. \tag{4.13}$$

The image $\tilde{I}(\tau_*)$ of τ_* can then be defined inductively. First we pick a triple point τ_* for which $W_{i(\tau_*)}$ contains no further triple point of $G_x(\omega)$ and let $\tilde{I}(\tau_*)$ be the leaf of $G_x(\omega)$ in $W_{i(\tau_*)}$. Then we remove the path connecting τ_* with $\tilde{I}(\tau_*)$ from the graph $G_x(\omega)$ and proceed in the same way for the remaining graph.

It remains to prove (4.13). Since $\tau_* \neq \tau_x$, there exists at least one i such that the triangle $\{x\} \cup e_i$ has an acute angle at x. We fix such an i and consider any $\tau \in \mathsf{C}_x(\omega)$ with $\tau \subset W_i$. There exists at least one point $z_0 \in \tau$ that is not contained in the closed disc $\overline{B}(\tau_*)$. Since $x \in B(\tau)$ and $\langle \tau \rangle$ is covered by the tiles $\langle \tau' \rangle$ for $\tau' \in \mathsf{C}_x^+(\omega)$ with $\langle \tau' \rangle \cap \langle e_i \rangle \neq \emptyset$, we conclude that the line segment s from z_0 to x is contained in $\overline{B}(\tau)$ and hits both the circle $\partial B(\tau_*)$ and the edge $\langle e_i \rangle$. In particular, $B(\tau) \cap \langle e_i \rangle \neq \emptyset$. Since $B(\tau)$ contains no points of ω, we deduce further that the circle $\partial B(\tau)$ hits the edge $\langle e_i \rangle$ in precisely two points z_1 and z_2. By the choice of i, the angle of the triangle $\{z_1, x, z_2\}$ at x is acute. Since $x \in B(\tau)$, it follows that the angle of the triangle $\{z_1, z_0, z_2\}$ at z_0 is obtuse. Consequently, if we consider running points y_k such that y_0 runs from z_0 to the point $s \cap \partial B(\tau_*)$ and the edge $\{y_1, y_2\}$ from $\{z_1, z_2\}$ to e_i, the associated circumcircles $B(\{y_1, y_0, y_2\})$ run from $B(\tau)$ to $B(\tau_*)$, and their radii $\varrho(\{y_1, y_0, y_2\})$ must grow. This proves that $\varrho(\tau) \leq \varrho(\tau_*)$, and the proof of (4.13) and the lemma is complete. $\Diamond$

Acknowledgment. We are grateful to Remy Drouilhet who brought us together and drew the interest of H.-O. G. to the subject. We also thank the referee for his useful comments.

References

[1] M. Aigner and G. Ziegler, *Proofs from the book*, Springer, Berlin etc., 3rd ed., 2004. MR2014872

[2] E. Bertin, J.-M. Billiot and R. Drouilhet, Existence of “nearest-neighbour” spatial Gibbs models, *Adv. in Appl. Probab.* **31**:895–909 (1999) MR1747447

[3] E. Bertin, J.-M. Billiot and R. Drouilhet, Existence of Delaunay pairwise Gibbs point process with superstable component, *J. Statist. Phys.* **95**:719–744 (1999) MR1700922

[4] E. Bertin, J.-M. Billiot and R. Drouilhet, Phase transition in the nearest-neighbor continuum Potts model, *J. Statist. Phys.* **114**:79–100 (2004) MR2032125

[5] E. Bertin, J.-M. Billiot and R. Drouilhet, *R*-local Delaunay inhibition model, *J. Stat. Phys.* **132**:649–667 (2008) MR2429701

[6] A. Dembo and O. Zeitouni, *Large Deviations: Techniques and Applications*. Springer, New York, 1998. MR1619036

[7] D. Dereudre, Gibbs Delaunay tessellations with geometric hard core conditions, *J. Statist. Phys.* **131**:127–151 (2008). MR2394701

[8] D. Dereudre, R. Drouilhet and H.-O. Georgii, Existence of Gibbsian graphs for stable interactions, *in preparation*.

[9] J. Fritz, Generalization of McMillan's theorem to random set functions. *Stud. Sci. Math. Hung.* **5**:369–394 (1970). MR0293956

[10] H.O. Georgii, *Gibbs Measures and Phase Transitions*, de Gruyter, Berlin, 1988. MR0956646

[11] H.O. Georgii, Large deviations and maximum entropy principle for interacting random fields on $\mathbb{Z}^d$, *Ann. Probab.* **21**:1845–1875 (1993). MR1245292

[12] H.O. Georgii, Large deviations and the equivalence of ensembles for Gibbsian particle systems with superstable interaction, *Probab. Theory Relat. Fields* **99**:171–195 (1994). MR1278881

[13] H.O. Georgii, The equivalence of ensembles for classical systems of particles, *J. Statist. Phys.* **80**:1341-1378 (1995). MR1349785

[14] H.O. Georgii and H. Zessin, Large deviations and the maximum entropy principle for marked point random fields, *Probab. Theory Relat. Fields* **96**:177–204 (1993). MR1227031

[15] G. Grimmett, Potts models and random-cluster processes with many-body interactions, *J. Statist. Phys.* **75**:67–121 (1994). MR1273054

[16] R. Holley and D. W. Stroock, Nearest neighbor birth and death processes on the real line, *Acta Math.*Ê**140**:103–154 (1978).

[17] J. Mfiller, *Lectures on Random Voronoi Tessellations*, Lecture Notes in Statistics Vol. 87, Springer, Berlin etc., 1994. MR1295245

[18] X.X. Nguyen and H. Zessin, Ergodic theorems for spatial processes, *Z. Wahrscheilichkeitstheor. Verw. Geb.* **48**:133–158 (1979). MR0534841

[19] C. Preston, *Random fields*, Lect. Notes Math. 534, Springer, Berlin, 1976. MR0448630

[20] R. Schneider and W. Weil, *Stochastic and integral geometry*, Springer, Berlin, 2008. MR2455326

[21] H. Zessin, Specific index and curvature for random simplicial complexes, *Inst. Math. Natl. Acad. Sci. Armenia* **37**:64–81 (2002). MR1964589