2 The next-order energy

This section is a study of the energy of a system of charges in a background. It starts from the splitting formula, which isolates the *next-order energy* of the GβE. Everything rests on one identity, the *master formula*, which writes this energy as a Gagliardo seminorm of the discrepancy function. It produces a quantity $\widetilde{\mathsf{G}}_W$, the *energy of a window W at unit resolution*, which controls the number of points, the discrepancies and the linear statistics in W, and which is itself controlled by the energy (Onsager's lemma).

2.1 The next-order energy

Let $\zeta(x) := \int -\log|x-y|\mathtt{d}\mu_{\mathrm{sc}}(y) + \frac{x^2}{4} - \frac{1}{2}$, which is non-negative and vanishes exactly on $[-2,2]$ (the Euler–Lagrange equation of $\mathcal{I}$).

Lemma 2.1 (Splitting formula)**.** *Let $x_i' := Nx_i$ and $\nu_N' := \sum_i \delta_{x_i'} - \rho_{\mathrm{sc}}(\frac{\cdot}{N})\mathtt{d}x'$, a signed measure of total mass zero. Then*

$$\mathsf{H}_N(\mathsf{X}_N) = N^2\mathcal{I}(\mu_{\mathrm{sc}}) - \frac{N}{2}\log N + N\sum_{i=1}^{N}\zeta(x_i) + \frac{1}{2}\iint_{x\neq y} -\log|x-y|\,\mathtt{d}\nu_N'(x)\mathtt{d}\nu_N'(y).$$

The proof is an expansion of the square, see Exercise 2.9 and Figure 1. In blown-up coordinates, the

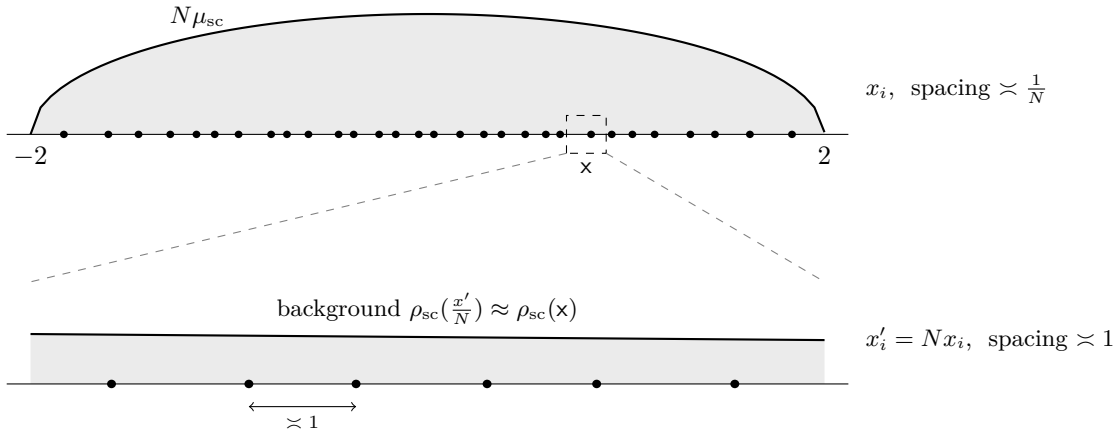


Figure 1: The blow-up $x' = Nx$. Top: the N charges at the macroscopic scale, distributed according to $N\mu_{\mathrm{sc}}$. Bottom: a window of length $O(1)$ around $N\mathsf{x}$, in the coordinates x'. The charges have spacing of order one, and the background $\rho_{\mathrm{sc}}(\frac{\cdot}{N})$ is almost constant: it varies at scale N. The configuration $\mathbf{X}_{N,\mathsf{x}}$ also recenters at x and rescales by $\rho_{\mathrm{sc}}(\mathsf{x})$, so as to have intensity one.

GβE is thus a system of N unit charges in a *neutralizing background* of density $\rho_{\mathrm{sc}}(\frac{\cdot}{N})$ – slowly varying, of order one – with an extra confinement which vanishes in $[-2N, 2N]$. The last term, the *next-order energy*, is of order N, like the entropy:

> *At the microscopic scale, energy and entropy compete, and the behavior depends on β.*

Energy as a distance between a configuration and a measure. The next-order energy is the energy $F(\mathbf{X},\mu) = \frac{1}{2}\iint_{x\neq y} -\log|x-y|\,\mathtt{d}(\mathbf{X}-\mu)^{\otimes 2}$ of a configuration $\mathbf{X}$ in a background μ of the same total mass. It measures how far the points are from being distributed according to μ: it is small when the points are well spread according to μ, and large when they accumulate or leave holes. The rest of this section is a study of this quantity for its own sake: we write it as a Gagliardo seminorm of the discrepancy function (the master formula), we see what it controls (numbers of points, discrepancies, linear statistics), and we bound it from below (Onsager's lemma).

2.2 Configurations, and notation

A *configuration* is a locally finite, integer-valued measure $\mathbf{X} = \sum_i \delta_{x_i}$ on $\mathbb{R}$: a collection of points, possibly infinite, with finitely many of them in each bounded set. For a set A, we write

$$\mathbf{X}(A) := \text{the number of points in } A, \qquad \mathbf{X}_A := \text{the restriction of } \mathbf{X} \text{ to } A.$$

Conf denotes the set of configurations, and $\mathsf{Conf}(\Lambda, n)$ the set of configurations of n points in Λ. Given a *background* μ (a measure with a bounded density, by default the Lebesgue measure), the *discrepancy* of an interval I is

$$\mathsf{Dis}(\mathbf{X}, I) := \mathbf{X}(I) - \mu(I), \qquad \text{by default } \mathsf{Dis}(\mathbf{X}, I) = \mathbf{X}(I) - |I| :$$

the number of points in I, minus the number expected. We write $\Lambda_L := [-\frac{L}{2}, \frac{L}{2}]$, and $\mathtt{C}, \mathtt{c}$ denote positive constants that may change from line to line.

The unit grid. We fix once and for all the *unit cells* $U_k := [k, k+1)$, $k \in \mathbb{Z}$, which partition $\mathbb{R}$, and write $\mathcal{U}$ for the set of unit cells. For an interval W, the *cells of* W are the non-empty pieces $U \cap W$, $U \in \mathcal{U}$; their set is denoted $\mathcal{U}(W)$. They partition W: the pieces inside W are whole unit cells, and at most two of them, at the ends of W, are cut, possibly very short ones. For a cell U (of the grid, or of W), $N_U := \mathbf{X}(U)$ is its number of points, and $\overline{U}$ its closure. Thus

$$\sum_U \quad \text{runs over all the unit cells } U \in \mathcal{U}, \qquad \sum_{U \in \mathcal{U}(W)} \quad \text{runs over the cells of } W,$$

and $\mathbf{X}(W) = \sum_{U \in \mathcal{U}(W)} N_U$, with $\#\mathcal{U}(W) \le |W| + 2$.

2.3 The energy of a finite system

A *finite system* is a finite configuration $\mathbf{X} = \sum_{i=1}^n \delta_{x_i}$ together with a *background* $\mu = m(t)\mathtt{d}t$, where $m \ge 0$ is bounded and compactly supported. Its energy is

$$F(\mathbf{X}, \mu) := \frac{1}{2} \iint_{x \ne y} -\log|x-y| \, \mathtt{d}\nu(x) \, \mathtt{d}\nu(y), \qquad \nu := \mathbf{X} - \mu,$$

where the diagonal is removed for the point charges only. When $\mu = \mathrm{Leb}_\Lambda$, we write $F_\Lambda(\mathbf{X}) := F(\mathbf{X}, \mathrm{Leb}_\Lambda)$: the energy of the charges in Λ together with a uniform neutralizing background on Λ. The system is *neutral* if $\nu(\mathbb{R}) = 0$, i.e. $n = \mu(\mathbb{R})$.

The *discrepancy function* of the system is

$$D(x) := \nu((-\infty, x]), \qquad \text{so that} \qquad D(y) - D(x) = \mathbf{X}((x,y]) - \mu((x,y]).$$

It jumps by $+1$ at each point, decreases with slope $-m$ in between, and, if the system is neutral, vanishes outside a compact set. All the results of these notes are statements about D.

2.4 The master formula

Proposition 2.2 (Master formula)**.** *For a neutral finite system with n distinct points,*

$$F(\mathbf{X}, \mu) = \lim_{\kappa \to 0} \left(\frac{1}{4} \iint_{|x-y|>\kappa} \frac{(D(x) - D(y))^2}{(x-y)^2} \mathtt{d}x \, \mathtt{d}y - \frac{n}{2} \left(\log \frac{1}{\kappa} + 1 \right) \right).$$

The formal computation, and what it misses. Formally, $\nu = D'$, and two integrations by parts give

$$\iint -\log|x-y| \, D'(x) D'(y) \mathtt{d}x \, \mathtt{d}y = -\iint \frac{D(x)D(y)}{(x-y)^2} \mathtt{d}x \, \mathtt{d}y.$$

Expanding $-D(x)D(y) = \frac{1}{2}(D(x) - D(y))^2 - \frac{1}{2}D(x)^2 - \frac{1}{2}D(y)^2$ produces the Gagliardo integral, and two terms $\frac{1}{2}\iint \frac{D(x)^2 + D(y)^2}{(x-y)^2}$ which are *infinite*. For a regular kernel g they would vanish, because $\int g'' =$

$[g']_{-\infty}^{+\infty} = 0$. For $g = -\log$, $g'' = u^{-2}$ is positive and not integrable: "$\int g'' = 0$" only holds for the Hadamard finite part, which carries an infinite negative mass on the diagonal. The proof makes this negative mass visible by truncating the logarithm.

Proof of Proposition 2.2. Let $\delta := \min_{i\neq j} |x_i - x_j|$, fix $\kappa < \delta$, and let $g_\kappa(u) := -\log \max(|u|, \kappa)$.

Step 1: truncating the kernel. Expand $\frac{1}{2}\iint g_\kappa \mathtt{d}\nu\mathtt{d}\nu$, the diagonal included.

- The pairs of distinct points give $\frac{1}{2}\sum_{i\neq j} -\log|x_i - x_j|$, since $|x_i - x_j| > \kappa$.
- The diagonal gives $\frac{n}{2} g_\kappa(0) = \frac{n}{2}\log\frac{1}{\kappa}$.
- In the terms involving μ, replacing g_κ by $-\log$ costs little: $|g_\kappa(u) + \log|u|| = \log\frac{\kappa}{|u|}\mathtt{1}_{|u|<\kappa}$, whose integral is 2κ. So these terms change by at most $2\kappa\|m\|_\infty(n + \mu(\mathbb{R}))$.

Hence

$$\frac{1}{2}\iint g_\kappa \mathtt{d}\nu\mathtt{d}\nu = F(\mathbf{X}, \mu) + \frac{n}{2}\log\frac{1}{\kappa} + O(\kappa). \tag{2.1}$$

Step 2: integrating by parts. The function g_κ is even and Lipschitz: $g'_\kappa = 0$ on $(-\kappa, \kappa)$, $g'_\kappa = -\frac{1}{u}$ outside, and $g'_\kappa(\pm\infty) = 0$. Its second derivative is the finite measure

$$g''_\kappa = \frac{\mathtt{1}_{|u|>\kappa}}{u^2}\mathtt{d}u - \frac{1}{\kappa}\left(\delta_\kappa + \delta_{-\kappa}\right), \qquad \int \mathtt{d}g''_\kappa = \frac{2}{\kappa} - \frac{2}{\kappa} = 0.$$

The function D has bounded variation and compact support, with $\mathtt{d}D = \nu$. Integrating by parts in x and then in y (the boundary terms vanish, since D has compact support and g'_κ is bounded),

$$\iint g_\kappa(x-y)\mathtt{d}\nu(x)\mathtt{d}\nu(y) = -\int D(y)\,(g''_\kappa * D)(y)\,\mathtt{d}y = -\int \mathtt{d}g''_\kappa(u)\int D(y)D(y+u)\mathtt{d}y.$$

Now $\int D(y)D(y+u)\mathtt{d}y = \|D\|_2^2 - \frac{1}{2}\int (D(y+u)-D(y))^2\mathtt{d}y$. The constant $\|D\|_2^2$ integrates to zero against g''_κ, whose total mass vanishes: these are the two "infinite terms" of the formal computation, which disappear exactly. Hence, the two Dirac masses giving the same contribution,

$$\iint g_\kappa \mathtt{d}\nu\mathtt{d}\nu = \frac{1}{2}\iint_{|x-y|>\kappa} \frac{(D(x)-D(y))^2}{(x-y)^2}\mathtt{d}x\,\mathtt{d}y - \frac{1}{\kappa}\int (D(x+\kappa)-D(x))^2\mathtt{d}x. \tag{2.2}$$

Step 3: the jumps. We have $D(x+\kappa) - D(x) = \mathbf{X}((x, x+\kappa]) - \mu((x, x+\kappa])$. Since $\kappa < \delta$, the first term is the indicator of the disjoint union A of the intervals $[x_i - \kappa, x_i)$, of measure $n\kappa$. The second, $a(x)$, satisfies $0 \le a \le \|m\|_\infty\kappa$ and $\int a = \kappa\mu(\mathbb{R})$. Hence

$$\int (D(x+\kappa)-D(x))^2\mathtt{d}x = \int (\mathtt{1}_A - a)^2 = n\kappa - 2\int_A a + \int a^2 = n\kappa + O(\kappa^2).$$

Conclusion. By (2.1), (2.2) and Step 3, for every $\kappa < \delta$,

$$F(\mathbf{X}, \mu) = \frac{1}{4}\iint_{|x-y|>\kappa} \frac{(D(x)-D(y))^2}{(x-y)^2} - \frac{n}{2}\left(\log\frac{1}{\kappa} + 1\right) + O(\kappa).$$

This proves the proposition, and shows that the bracket converges with an error $O(\kappa)$. □

Meaning of the constants. The $\log\frac{1}{\kappa}$ is the truncated self-energy of each point. The $+1$ is the negative mass $-\frac{1}{\kappa}(\delta_\kappa + \delta_{-\kappa})$ of g''_κ, seen by the jumps of D. For a smooth D this negative mass contributes $O(\kappa)$, and no renormalization is needed.

2.5 What the energy controls

The master formula writes the energy with the Gagliardo integral of D. Its part at distance ≥ 1,

$$\mathsf{Q}_1(W) := \iint_{W^2,\,|x-y|\geq 1} \frac{(D(x)-D(y))^2}{(x-y)^2}\mathtt{d}x\,\mathtt{d}y,$$

together with the counts N_U of the cells (Section 2.2), which see the scale one, form the quantity that we will localize. For an interval W, the *local energy at unit resolution* is

$$\widetilde{\mathsf{G}}_W := \frac{1}{4}\mathsf{Q}_1(W) + \sum_{U\in\mathcal{U}(W)} N_U^2.$$

It only depends on the configuration and the background in W, and it is monotone: if $W \subset W'$, each cell of W is contained in a cell of W', so $\widetilde{\mathsf{G}}_W \leq \widetilde{\mathsf{G}}_{W'}$.

Proposition 2.3 (What the energy controls)**.** *Let the background density be bounded by $\bar{m}$.*

1. *(Number of points, and overcrowding.) For every interval W,*

$$\mathbf{X}(W)^2 \leq (|W|+2)\,\widetilde{\mathsf{G}}_W.$$

Moreover, if an interval I contains $k \geq 8(1+\bar{m})(|I|+1)$ points, then the window W of length k centered at I satisfies

$$\widetilde{\mathsf{G}}_W \geq \mathtt{c}\,k^2 \log\frac{k}{|I|+1}.$$

2. *(Discrepancies.) If I is an interval of length ≥ 1, then*

$$\mathsf{Dis}(\mathbf{X},I)^2 \leq \mathtt{C}\log(2+|I|)\left(\widetilde{\mathsf{G}}_I + 1\right). \tag{2.3}$$

3. *(Linear statistics.) If φ is Lipschitz and supported in W, then*

$$|\mathrm{Fluct}(\varphi)| \leq \|\varphi'\|_{L^2}\left(\mathtt{C}|W|(\widetilde{\mathsf{G}}_W+1)\right)^{1/2}, \qquad \mathrm{Fluct}(\varphi) := \int \varphi\,\mathtt{d}\nu.$$

For $\varphi = \theta(\cdot/\ell)$ and $|W| \asymp \ell$, this is $\leq \mathtt{C}\|\theta'\|_{L^2}(\widetilde{\mathsf{G}}_W+1)^{1/2}$.

What the energy is a norm of, and what we need

Up to the counts at scale one, the energy in W is $\frac{1}{4}\mathsf{Q}_1(W)$, the square of a *truncated $\dot{H}^{1/2}$ seminorm* of D (Gagliardo form, pairs at distance ≥ 1). The three items of the proposition ask for three different controls of D.

- *Counts*: $\mathbf{X}(W) = \sum_{U\in\mathcal{U}(W)} N_U$, which is directly in $\widetilde{\mathsf{G}}_W$.
- *Smooth statistics*: by integration by parts and Cauchy–Schwarz, for every constant c (free because $\int\varphi' = 0$),

$$|\mathrm{Fluct}(\varphi)| = \Big|\int_W \varphi'(D-c)\Big| \leq \|\varphi'\|_{L^2}\,\|D-c\|_{L^2(W)}.$$

  We need the *L^2 norm of D modulo constants*.
- *Box discrepancies*: $\mathsf{Dis}(\mathbf{X},[a,b]) = D(b)-D(a)$. We need the *values of D at points*, modulo constants.

These are two classical questions of analysis: does a (truncated) $\dot{H}^{1/2}$ seminorm control the L^2 norm modulo constants (Poincaré), and the L^∞ norm? We answer them for general functions, and then apply the answers to D.

Two lemmas of analysis

Setting. Let J be an interval and $f : J \to \mathbb{R}$ a function, bounded on bounded sets. Two quantities enter.

- *The truncated Gagliardo seminorm*: for $A, B \subset J$,

$$[f]^2_{A\times B} := \iint_{\substack{A\times B\\ |x-y|\geq 1}} \frac{(f(x)-f(y))^2}{(x-y)^2}\mathtt{d}x\,\mathtt{d}y, \qquad [f]^2_J := [f]^2_{J\times J}.$$

  It is a seminorm (the L^2 norm of $(x,y) \mapsto \frac{f(x)-f(y)}{x-y}$ on a set of pairs), it vanishes on the constants, and only on them as soon as $|J| \geq 2$. It does not see the values of f at individual points: changing f on a null set does not change it.

- *The oscillations at scale one*: for each cell $U \in \mathcal{U}(J)$,

$$\omega_U(f) := \sup_{\overline{U}\cap J} f - \inf_{\overline{U}\cap J} f.$$

  We take the closure so that two consecutive cells share an endpoint. If f has bounded variation, $\omega_U(f) \leq |f'|(\overline{U} \cap J)$: this is a control of $W^{1,1}$ (or BV) type, discretized at scale one.

The seminorm controls all the scales ≥ 1, the oscillations control the scale 1. The point of the two lemmas is *how much of each is needed.*

Lemma 2.4 (Poincaré inequality)**.** *For intervals $V \subset V' \subset J$ with $|V| \geq 1$,*

$$\frac{1}{|V||V'|}\iint_{V\times V'} (f(x)-f(y))^2\mathtt{d}x\,\mathtt{d}y \leq \frac{|V'|}{|V|}[f]^2_{V'} + \frac{\mathtt{C}}{|V||V'|}\sum_{U\in\mathcal{U}(V')}\omega_U(f)^2,$$

with an absolute constant $\mathtt{C}$. In particular, for $V = V' = W$, with $\langle f\rangle_W := \frac{1}{|W|}\int_W f$,

$$\frac{1}{|W|}\int_W (f-\langle f\rangle_W)^2 \leq \frac{1}{2}[f]^2_W + \frac{\mathtt{C}}{|W|^2}\sum_{U\in\mathcal{U}(W)}\omega_U(f)^2. \tag{2.4}$$

Proof. Split the pairs $(x,y) \in V \times V'$ according to $|x-y|$.

Far pairs, $|x-y| \geq 1$. Both points lie in V', so $|x-y| \leq |V'|$, that is, $1 \leq \frac{|V'|^2}{(x-y)^2}$. Multiplying by $(f(x)-f(y))^2$ and integrating, these pairs give at most $|V'|^2[f]^2_{V'}$. Divide by $|V||V'|$.

Near pairs, $|x-y| < 1$. Then x and y lie in the same cell of V' or in two consecutive cells U, U', whose closures share an endpoint, so $|f(x)-f(y)| \leq \omega_U(f)+\omega_{U'}(f)$. For x in a given cell, the y with $|x-y|<1$ have measure ≤ 2 and lie in at most three cells. Summing, the near pairs give at most $\mathtt{C}\sum_U \omega_U(f)^2$.

For (2.4), expand the square: $\iint_{W^2}(f(x)-f(y))^2 = 2|W|\int_W (f-\langle f\rangle_W)^2$. □

What it says. *The seminorm controls the L^2 distance to the constants, with a constant that does not depend on the size of W.* The oscillations at scale one are needed *everywhere in W, but only on average*: they enter through $\frac{1}{|W|^2}\sum_{U\in\mathcal{U}(W)}\omega_U^2$, which is $\frac{1}{|W|}$ times their mean square, a lower-order term for large windows. The proof is one observation: in a window of length ℓ, the weight $(x-y)^{-2}$ is at least ℓ^{-2}. This is efficient when the increments of f live at the scale of the window, and wasteful when they live at much smaller scales, where the true weight is much larger. This is why, below, windows are only compared with windows of comparable size, $|V'| \leq 8|V|$.

Lemma 2.5 (Point values: the embedding into L^∞ at a cost $\sqrt{\log}$)**.** *Let $I = [a,b] \subset J$ with $|I| \geq 1$. For $p \in I$, let*

$$\Omega_p(f) := \sum_{U\in\mathcal{U}(I)} \frac{\omega_U(f)^2}{(1+\operatorname{dist}(p,U))^2}$$

be the oscillations at scale one, weighted by the inverse square distance to p. *Then*

$$(f(b)-f(a))^2 \leq \mathtt{C}\log(2+|I|)\left([f]^2_I + \Omega_a(f) + \Omega_b(f)\right), \tag{2.5}$$

with an absolute constant $\mathtt{C}$.

Proof. If $|I| < 8$, I has at most 9 cells, consecutive ones sharing an endpoint, so $|f(b) - f(a)| \le \sum_{U\in\mathcal{U}(I)} \omega_U(f)$, and Cauchy–Schwarz gives $(f(b) - f(a))^2 \le 9\sum_{U\in\mathcal{U}(I)} \omega_U(f)^2 \le \mathtt{C}\,\Omega_a(f)$. Let $|I| \ge 8$, and let $K \ge 1$ be maximal with $2^K \le \frac{|I|}{4}$. We use repeatedly the following elementary fact. For two intervals A, B, the difference of the averages is the average of the increments over the pairs, $\langle f\rangle_A - \langle f\rangle_B = \frac{1}{|A||B|}\iint_{A\times B}(f(x)-f(y))\mathtt{d}x\,\mathtt{d}y$, and the square of an average is at most the average of the square (Cauchy–Schwarz):

$$\left(\langle f\rangle_A - \langle f\rangle_B\right)^2 \le \frac{1}{|A||B|}\iint_{A\times B}(f(x)-f(y))^2\mathtt{d}x\,\mathtt{d}y. \tag{2.6}$$

We *zoom out from* a, by factors 2:

$$V_k := [a, a+2^k] \quad (0 \le k \le K), \qquad A_k := [a+2^k, a+2^{k+1}], \quad \text{so that } V_{k+1} = V_k \cup A_k,$$

and symmetrically from b, with $V'_k := [b-2^k, b]$ and A'_k (Figure 2). Telescoping,

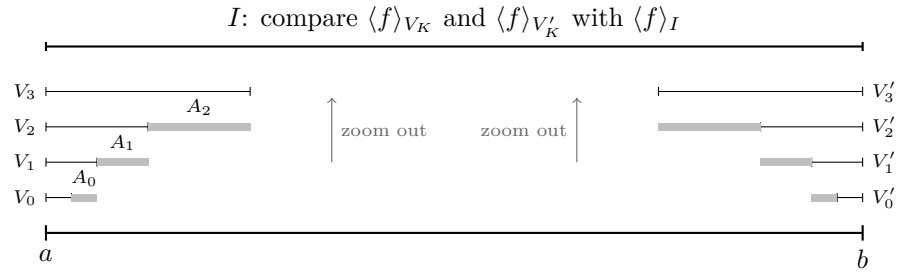


Figure 2: Zooming out from the endpoints of $I = [a, b]$, here with $K = 3$. Each window $V_{k+1} = V_k \cup A_k$ doubles the previous one. The added pieces A_k (grey) are disjoint, so the steps use disjoint sets of pairs $V_k \times A_k$: they share the seminorm. The last windows V_K, V'_K, of length $\asymp |I|$, are compared with I.

$$f(b) - f(a) = \big(f(b) - \langle f\rangle_{V'_0}\big) + \sum_{k<K}\big(\langle f\rangle_{V'_k} - \langle f\rangle_{V'_{k+1}}\big) + \big(\langle f\rangle_{V'_K} - \langle f\rangle_I\big) - \big(\text{the same from } a\big),$$

a sum of $2K + 4$ terms. We bound their squares.

- *Endpoints.* $|f(a) - \langle f\rangle_{V_0}| \le \sup_{V_0} f - \inf_{V_0} f$, and V_0 meets at most two cells of I, at distance ≤ 1 from a: the square is $\le \mathtt{C}\,\Omega_a(f)$.
- *One zoom step.* V_{k+1} is the union of V_k and A_k, which have the same length 2^k, so $\langle f\rangle_{V_{k+1}} = \frac{1}{2}(\langle f\rangle_{V_k} + \langle f\rangle_{A_k})$ is the midpoint of the two averages. Hence $\langle f\rangle_{V_{k+1}} - \langle f\rangle_{V_k} = \frac{1}{2}(\langle f\rangle_{A_k} - \langle f\rangle_{V_k})$, and (2.6) with $A = A_k$, $B = V_k$ gives

$$\delta_k^2 := \big(\langle f\rangle_{V_{k+1}} - \langle f\rangle_{V_k}\big)^2 = \tfrac{1}{4}\left(\langle f\rangle_{A_k} - \langle f\rangle_{V_k}\right)^2 \le \frac{1}{4\cdot 4^k}\iint_{V_k\times A_k}(f(x)-f(y))^2.$$

  The far pairs have $|x - y| \le 2^{k+1}$, so they give at most $[f]^2_{V_k\times A_k}$. The near pairs lie within distance 1 of the junction point $a + 2^k$, hence in the at most three cells of I that meet $[a + 2^k - 1, a + 2^k + 1]$, which are at distance $\le 2^k + 1$ from a. Hence they give at most $\mathtt{C}$ times the sum of $\frac{\omega_U(f)^2}{(1+\mathrm{dist}(a,U))^2}$ over these three cells.
- *The steps share the seminorm.* The sets of pairs $V_k \times A_k$ are disjoint, because the A_k are. They lie in $[a, a + \frac{|I|}{4}]^2$, and those of the chain from b lie in $[b - \frac{|I|}{4}, b]^2$. A cell meets $[j - 1, j + 1]$ for at most three junction points j. Hence

$$\sum_{k<K}\left(\delta_k^2 + \delta_k'^2\right) \le [f]_I^2 + \mathtt{C}\left(\Omega_a(f) + \Omega_b(f)\right).$$

- *Top step.* $V_K \subset I$ and $|I| < 8|V_K|$. By (2.6) and Lemma 2.4, $(\langle f\rangle_{V_K} - \langle f\rangle_I)^2 \le 8[f]_I^2 + \frac{\mathtt{C}}{|I|^2}\sum_{U\in\mathcal{U}(I)}\omega_U(f)^2 \le 8[f]_I^2 + \mathtt{C}\,\Omega_a(f)$, since every cell meeting I is at distance $\le |I|$ from a.

By Cauchy–Schwarz over the $2K+4 \le \mathtt{C}\log(2+|I|)$ terms, $(f(b)-f(a))^2 \le (2K+4)\sum(\text{terms})^2$, which gives (2.5). □

What it says. Three remarks.

- *What is assumed at scale one.* The oscillations ω_U enter with the weight $(1+\operatorname{dist}(p,U))^{-2}$: what matters is the scale one *near the two endpoints*, and far cells count less and less. In particular $\Omega_p(f) \le \mathtt{C}\sup_{U\in\mathcal{U}(I)}\omega_U(f)^2$, and $\Omega_p(f) \le \sum_{U\in\mathcal{U}(I)}\omega_U(f)^2$. Some control at scale one near the endpoints is unavoidable: the seminorm does not see the value $f(a)$ itself.
- *Where the logarithm comes from.* The only loss is the Cauchy–Schwarz inequality $(\sum_{k<K}\delta_k)^2 \le K\sum_k\delta_k^2$ over the $K \approx \log_2|I|$ scales. It is an equality when all the δ_k are equal: the worst f has the same jump of averages at each scale between 1 and $|I|$, that is, $f(x) \approx \frac{\log(|I|/|x-a|)}{\log|I|}$. This f has $f(a)-f(b)\approx 1$, $[f]_I^2 \asymp \frac{1}{\log|I|}$ and small oscillations (Exercise 2.10, item 4): *the factor* $\log(2+|I|)$ *cannot be removed.* If the steps did not share the seminorm, we would get $\log^2$.
- *Why it is not scale invariant.* Both ingredients refer to the scale 1: the truncation $|x-y|\ge 1$ and the grid of cells. The constant depends on the ratio of $|I|$ to this scale, and grows like its logarithm. Compare with Lemma 2.4, whose constant does not depend on $|W|$. Without any control at small scales, there is no bound at all: the untruncated $\dot H^{1/2}$ seminorm does not control L^∞ (think of $\log\log\frac{1}{|x|}$ near 0). This is the critical case of the Sobolev embedding, as $H^1 \not\subset L^\infty$ in dimension two.

Application to the energy

Dictionary. For the discrepancy function, $[D]_W^2 = \mathtt{Q}_1(W)$. Let $U\in\mathcal{U}(W)$, and U^+ the next cell of W (if any). For $x<y$ in $\overline{U}\cap W$, $D(y)-D(x) = \mathbf{X}((x,y]) - \mu((x,y])$ lies between $-\bar m$ and $N_U + N_{U^+}$ (the right endpoint of U belongs to U^+), so

$$\omega_U(D) \le N_U + N_{U^+} + \bar m.$$

The oscillations of D at scale one are controlled by the counts. Hence, for $p\in I$ and $|W|\ge 1$,

$$\Omega_p(D) \le \mathtt{C}\sum_{U\in\mathcal{U}(I)} N_U^2 + \mathtt{C}\bar m^2 \sum_{U\in\mathcal{U}(I)}\frac{1}{(1+\operatorname{dist}(p,U))^2} \le \mathtt{C}\left(\widetilde{\mathsf{G}}_I + 1\right), \qquad \frac{1}{|W|^2}\sum_{U\in\mathcal{U}(W)}\omega_U(D)^2 \le \mathtt{C}\left(\widetilde{\mathsf{G}}_W + 1\right).$$

Consequences. For every window W, Lemma 2.4 and the dictionary give

$$\frac{1}{|W|}\int_W (D - \langle D\rangle_W)^2 \le \mathtt{C}\left(\widetilde{\mathsf{G}}_W + 1\right). \tag{2.7}$$

For nested windows, since $\langle D\rangle_V - \langle D\rangle_{V'} = \frac{1}{|V||V'|}\iint_{V\times V'}(D(x)-D(y))$, the inequality (2.6) and Lemma 2.4 give the form used in Sections 4–5:

Corollary 2.6 (Comparing averages at comparable scales)**.** *For intervals $V\subset V'$ with $|V|\ge 1$ and $|V'|\le 8|V|$,*

$$(\langle D\rangle_V - \langle D\rangle_{V'})^2 \le \mathtt{C}\left(\widetilde{\mathsf{G}}_{V'} + 1\right).$$

Proof of Proposition 2.3. We prove the three items in turn.

Item 1. The first claim is Cauchy–Schwarz over the at most $|W|+2$ cells of W: $\mathbf{X}(W) = \sum_{U\in\mathcal{U}(W)} N_U \le (|W|+2)^{1/2}(\sum_{U\in\mathcal{U}(W)} N_U^2)^{1/2}$. For the second, the other points only increase D, so $D(y) - D(x) \ge k - \bar{m}(y-x)$ for $x < I < y$. The pairs with

$$x \in \left[\inf I - \tfrac{k}{8(1+\bar{m})},\, \inf I\right], \qquad y \in \left[\sup I,\, \sup I + \tfrac{k}{8(1+\bar{m})}\right]$$

lie in W and have $y-x \le |I| + \frac{k}{4(1+\bar{m})} \le \frac{k}{2(1+\bar{m})}$, hence $D(y)-D(x) \ge \frac{k}{2}$. Keeping those with $y-x \ge |I|+1$, they give

$$\widetilde{\mathsf{G}}_W \ge \frac{1}{4}\cdot\frac{k^2}{4}\iint \frac{\mathtt{d}x\,\mathtt{d}y}{(y-x)^2} \ge \mathtt{c}k^2 \log\frac{k}{|I|+1},$$

with $\mathtt{c}$ depending on $\bar{m}$ (the double integral is $\ge \mathtt{c}\log(1+\frac{s}{|I|+1})$, with $s = \frac{k}{8(1+\bar{m})}$ the length of the ranges). These are *long-range* pairs: the counts alone would only give $\sum_{U\in\mathcal{U}(W)} N_U^2 \ge \frac{k^2}{|I|+2}$, with no logarithm.

Item 2. Lemma 2.5 with $f = D$, $[D]_I^2 = \mathsf{Q}_1(I) \le 4\widetilde{\mathsf{G}}_I$, and the dictionary.

Item 3. Take $c = \langle D\rangle_W$ in the Cauchy–Schwarz bound above, and apply (2.7). □

In words: *smooth statistics only need the L^2 norm modulo constants, which the energy controls at no cost (Poincaré); box discrepancies need point values, which the energy controls up to a factor $\sqrt{\log}$ (the critical embedding into L^∞).* With the local laws, $\widetilde{\mathsf{G}}_W \lesssim |W|$, this gives $|\mathrm{Fluct}| \lesssim \sqrt{\ell}$ and $|\mathsf{Dis}| \lesssim \sqrt{\ell\log\ell}$ at scale ℓ.

2.6 Lower bounds: Onsager's lemma

The master formula splits the energy into the pairs at distance ≥ 1, which form the non-negative Gagliardo integral $\frac{1}{4}\mathsf{Q}_1$, and the renormalized pairs at distance < 1, which see the points one by one. We now show that the energy controls $\widetilde{\mathsf{G}}$: this is what makes the previous subsection useful.

Proposition 2.7 (Onsager's lower bound)**.** *For a neutral system with n points and background density $\le \bar{m}$,*

$$F(\mathbf{X},\mu) \ge \frac{1}{4}\mathsf{Q}_1(\mathbb{R}) + \mathtt{c}\sum_U N_U^2 - \mathtt{C}(1+\bar{m})n \ge -\mathtt{C}(1+\bar{m})n. \tag{2.8}$$

Proof. Split the integral of the master formula at $|x-y| = 1$. The pairs at distance ≥ 1 give $\frac{1}{4}\mathsf{Q}_1(\mathbb{R})$, and the renormalized pairs at distance < 1 are bounded below by Lemma 2.8 below. □

Note that (2.8) is *deterministic*, and that it controls the counts at scale one: a neutral system of energy F has $\sum_U N_U^2 \le \mathtt{C}(F + (1+\bar{m})n)$.

The pairs at distance < 1: what to expect. In the master formula, the pairs at distance < 1 carry the divergent part $\frac{n}{2}\log\frac{1}{\kappa}$, which is removed by the renormalization. The next lemma computes what is left, and the answer is simple:

- *a point alone* produces exactly its self-energy $\frac{1}{2}\log\frac{1}{\kappa}$, which the renormalization removes;
- *two points* at distance $\delta < 1$ produce $\log\frac{1}{\delta} - 1 + \delta \ge 0$: their interaction $-\log\delta$, truncated at distance 1;
- *the background* costs at most $\bar{m}$ per point.

So the short-distance part of the energy is non-negative up to $O(n)$, and it penalizes the accumulation of points: N points in a unit cell cost at least $\mathtt{c}N^2$.

Lemma 2.8 (Near-diagonal pairs)**.** *For a neutral system with background density $\le \bar{m}$ and n points,*

$$\lim_{\kappa\to 0}\left(\frac{1}{4}\iint_{\kappa<|x-y|<1}\frac{(D(x)-D(y))^2}{(x-y)^2}-\frac{n}{2}\left(\log\frac{1}{\kappa}+1\right)\right)\ge\frac{1}{16}\sum_U N_U^2-(1+\bar{m})n.$$

Proof. By symmetry, the integral over $|x-y|\in(\kappa,1)$, with the factor $\frac{1}{4}$, is $\frac{1}{2}$ times the integral over $x<y$. For $x<y$, let $h:=y-x$ and $J:=(x,y]$. Then $D(y)-D(x)=\mathbf{X}(J)-\mu(J)$, with $0\le\mu(J)\le\bar{m}h$, so

$$(D(y)-D(x))^2\ge\mathbf{X}(J)^2-2\bar{m}h\,\mathbf{X}(J),\qquad\mathbf{X}(J)^2=\sum_p \mathtt{1}_{p\in J}+\sum_{p\neq p'}\mathtt{1}_{p,p'\in J},$$

where p,p' run over the points, and the second sum is over ordered pairs. Each term is a sum over points or pairs of points of an explicit integral. We only need one observation: *for a fixed length h, the set of x such that $J=(x,x+h]$ contains a given point p is $[p-h,p)$, of length h; and J contains two given points at distance δ for a set of x of length $(h-\delta)_+$.*

One point: the self-energy. Each point contributes

$$\frac{1}{2}\int_\kappa^1\frac{h}{h^2}\mathtt{d}h=\frac{1}{2}\log\frac{1}{\kappa}.$$

The n points give exactly $\frac{n}{2}\log\frac{1}{\kappa}$: the renormalization, up to the $-\frac{n}{2}$ left in the bracket.

Two points: the truncated interaction. Two points at distance δ (counted twice, as two ordered pairs) contribute

$$g(\delta):=\int_\delta^1\frac{h-\delta}{h^2}\mathtt{d}h=\log\frac{1}{\delta}-1+\delta\quad\text{if }\delta<1,\qquad g(\delta)=0\text{ otherwise.}$$

Note that $g\ge 0$.

The background. The term $-2\bar{m}h\,\mathbf{X}(J)$ contributes $-\bar{m}\sum_p\int_\kappa^1\frac{h\cdot h}{h^2}\mathtt{d}h\ge-\bar{m}n$.

Collecting. The bracket is at least $\sum_{\{p,p'\}}g(|p-p'|)-(\frac{1}{2}+\bar{m})n$, the sum running over unordered pairs of points. Since $g\ge 0$, we may keep only some pairs. Cut each unit cell U into 8 intervals of length $\frac{1}{8}$, containing $N_1,\dots,N_8$ points. Two points in the same small interval are at distance $\delta<\frac{1}{8}$, so $g(\delta)\ge\log 8-1\ge 1$. By Cauchy–Schwarz, $\sum_j N_j^2\ge\frac{N_U^2}{8}$, so there are at least $\frac{1}{2}\sum_j N_j(N_j-1)\ge\frac{N_U^2}{16}-\frac{N_U}{2}$ such pairs. Summing over the cells gives at least $\frac{1}{16}\sum_U N_U^2-\frac{n}{2}$, and the lemma. □

Non-neutral windows. Neutrality is essential in Proposition 2.7. For a window with no point at all, $F_\Lambda=\frac{1}{2}\iint_{\Lambda^2}-\log|x-y|\approx-\frac{|\Lambda|^2}{2}\log|\Lambda|$, so no bound of the form $F_\Lambda\ge-\mathtt{C}n$ can hold. We will only need a crude bound. Expand F_Λ into the point–point, point–background and background–background terms, and use $-\log|x_i-x_j|\ge-\log_+|\Lambda|$, $\int_\Lambda-\log|x-t|\mathtt{d}t\le\int_{-1}^1-\log|u|\mathtt{d}u=2$, and $\frac{1}{2}\iint_{\Lambda^2}-\log|s-t|\ge-\frac{1}{2}|\Lambda|^2\log_+|\Lambda|$:

$$F_\Lambda(\mathbf{X})\ge-\mathtt{C}(n+|\Lambda|)^2(1+\log_+|\Lambda|)\qquad\text{for every configuration of }n\text{ points in }\Lambda.$$

The global energy is of order N. For a neutral system, Proposition 2.7 gives the deterministic bound

$$\widetilde{\mathsf{G}}_{\mathbb{R}}\le\mathtt{C}\left(F(\mathbf{X},\mu)+(1+\bar{m})n\right).$$

Apply it to the GβE in microscopic coordinates: N points, with the background $\rho_{\mathrm{sc}}(\cdot/N)$, of density $\le\frac{1}{\pi}$. We will prove in Section 3.4 (Lemma 3.5), by an elementary computation of free energies and without any exact formula, that the energy is $O(N)$ with exponential tails. Hence

$$\mathbb{P}_{N,\beta}\left(\widetilde{\mathsf{G}}_{\mathbb{R}}\ge tN\right)\le e^{-(\mathtt{c}t-\mathtt{C})N}\qquad(t>0):$$

globally, the energy at unit resolution is of the order of the volume.

This is an *average* statement over the system. It allows, for instance, $\sqrt{N}$ points in a single unit cell, or a net charge $\sqrt{N}$ in a window of length $\sqrt{N}$, at a cost $O(N)$. The local laws (Theorem 4.11) *localize* it:

$\widetilde{\mathsf{G}}_W \lesssim |W|$ *for every given bulk window* W*, at every scale from* $O(1)$ *to* N*, with exponential tails.*

Proposition 2.3 explains why this is the right quantity to localize: a bound $\widetilde{\mathsf{G}}_W \lesssim |W|$ has concrete consequences at the scale of W. The global bound only gives them for windows of size comparable to N: for a box of length $\ell \ll N$, it only gives $\mathsf{Dis}^2 \lesssim N \log \ell$, which is useless.

2.7 Comments

The comments are grouped in three themes, in the order of the section: the master formula seen from Fourier analysis, the sharpness of the controls of Proposition 2.3, and the dictionary with the electric formalism. None of them is used in the proofs.

The master formula from a higher viewpoint

Fourier. For a smooth neutral $\nu = D'$ with D smooth and compactly supported, use the convention $\widehat{f}(\xi) = \int f e^{-\mathrm{i}\xi x}$, so that $\int |f|^2 = \frac{1}{2\pi}\int |\widehat{f}|^2$. On neutral measures, the Fourier transform of $-\log|\cdot|$ is $\frac{\pi}{|\xi|}$ (up to a multiple of δ_0, which does not see neutral measures), and $\widehat{\nu} = \mathrm{i}\xi\widehat{D}$. Hence

$$\frac{1}{2}\iint -\log|x-y|\mathtt{d}\nu\mathtt{d}\nu = \frac{1}{2}\cdot\frac{1}{2\pi}\int \frac{\pi}{|\xi|}|\widehat{\nu}(\xi)|^2\mathtt{d}\xi = \frac{1}{4}\int |\xi||\widehat{D}(\xi)|^2\mathtt{d}\xi.$$

On the other hand,

$$\iint \frac{(D(x)-D(y))^2}{(x-y)^2} = \int \frac{\mathtt{d}h}{h^2}\int |D(x+h)-D(x)|^2\mathtt{d}x = \frac{1}{2\pi}\int |\widehat{D}|^2\int \frac{2(1-\cos\xi h)}{h^2}\mathtt{d}h\,\mathtt{d}\xi = \int |\xi||\widehat{D}|^2\mathtt{d}\xi,$$

since $\int \frac{2(1-\cos\xi h)}{h^2}\mathtt{d}h = 2\pi|\xi|$. The two sides agree.

Operator language. $\frac{1}{2}\iint -\log|x-y|\mathtt{d}\nu\mathtt{d}\nu = \frac{1}{2}\langle D, \pi(-\Delta)^{1/2}D\rangle$, since $-\partial_x\partial_y\log|x-y|$ is the kernel of $\pi|\partial| = \pi(-\Delta)^{1/2}$. The singular integral representation $\pi(-\Delta)^{1/2}D(x) = \text{p.v.}\int \frac{D(x)-D(y)}{(x-y)^2}\mathtt{d}y$ makes the difference $D(x)-D(y)$ appear at once, and symmetrizing gives $\frac{1}{4}\iint \frac{(D(x)-D(y))^2}{(x-y)^2}$. For point charges, the high-level proof must be completed by a renormalization at the charges, which is exactly Steps 1 and 3 of the proof of Proposition 2.2.

Sharpness of the controls: worst case versus average

With an energy $\widetilde{\mathsf{G}}_W$ of order $|W| = \ell$, (2.3) gives $|\mathsf{Dis}| \lesssim \sqrt{\ell\log\ell}$ for a sharp box. This is *worse than Poisson* by a logarithm, while we know that finite energy forces hyperuniformity (Section 6). There is no contradiction: the two statements are about different quantities.

The worst box: $\sqrt{\ell\log\ell}$, and this is sharp. Inequality (2.3) holds deterministically, for *every* box $I \subset W$, wherever it is placed. At this level it cannot be improved. A net charge q in a box of length ℓ can be produced at energy cost $\asymp \frac{q^2}{\log\ell}$, by spreading it evenly over the dyadic scales near an endpoint (Exercise 2.10, item 4). With an energy of order ℓ, this gives $q \asymp \sqrt{\ell\log\ell}$. But this profile is bad *only for the boxes whose endpoint sits on the logarithmic layers*: shift the box by a fraction of ℓ, and its discrepancy drops to $O(q/\log\ell)$.

The average box: Poisson level, with no logarithm. The energy is an integral over *all* pairs, that is, over all intervals $[x,y]$, with the weight $(y-x)^{-2}$ and with $D(y)-D(x) = \mathsf{Dis}(\mathbf{X},[x,y])$. Restricting $\mathsf{Q}_1(W)$ to the pairs with $y-x\in[u,2u]$, $1\le u\le\frac{|W|}{4}$, and using $(y-x)^{-2}\ge(2u)^{-2}$,

$$a(u) := \frac{1}{u}\cdot \operatorname*{avg}_{\substack{x\in W,\, v\in[u,2u]\\ [x,x+v]\subset W}} \mathsf{Dis}(\mathbf{X},[x,x+v])^2 \le \mathtt{C}\,\frac{\mathsf{Q}_1(W)}{|W|}.$$

On average over the position of the box, the discrepancy is at the Poisson level $\sqrt{u}$, without logarithm. Better: the dyadic scales $u = 2^k$ use disjoint sets of pairs, so $\sum_k a(2^k) \le \mathtt{C}\,\frac{\mathsf{Q}_1(W)}{|W|}$. The energy density is *shared* between the $\log_2 \ell$ scales, and for most of them $a(2^k)$ is much smaller than the energy density.

Stationary processes: hyperuniformity. For a stationary process, $\mathbb{E}[\mathsf{Dis}(\mathbf{X}, [x, x+u])^2] = \sigma^2(u)$ does not depend on x, so *the average over positions is the expectation for each box.* The previous bound becomes the formula of Theorem 6.6: up to a bounded term, the energy per unit length is $\frac{1}{2}\int_1^\infty \frac{\sigma^2(u)}{u^2}\mathtt{d}u$, a convergent sum over the dyadic scales of contributions of order $\frac{\sigma^2(2^k)}{2^k}$. These contributions tend to zero: $\sigma^2(u) = o(u)$ (Theorem 6.6, item 1). *Hyperuniformity is the tail of a convergent series over all scales up to infinity*: it is an averaged and infinite-volume statement, which no deterministic bound in a finite window can see.

What the course does with it. The deterministic bound (2.3) is only an *a priori* input, uniform over all boxes. The rigidity of the log-gas is a property of the Gibbs measure, not of the energy: Section 5 improves the discrepancies to $O(\log \ell)$ for sharp boxes and to $O(1)$ for smooth statistics, far below $\sqrt{\ell}$.

Dictionary with the electric formalism

What the electric field computes. The electric potential $u(z) = \int -\log|z-y|\mathtt{d}\nu(y)$ of ν is harmonic off the line, and on the line its derivative is, up to a factor $-\pi$, the Hilbert transform of ν. Its Dirichlet energy is $\int_{\mathbb{R}^2} |\nabla u|^2 = 2\pi \iint -\log|x-y|\mathtt{d}\nu\mathtt{d}\nu = \pi \iint \frac{(D(x)-D(y))^2}{(x-y)^2}$: the plane only computes, by harmonic extension, a quantity that already lives on the line. The truncation $|x-y| > \kappa$ plays the role of the smearing of the charges at scale κ. In the electric formalism, Onsager's lemma is proven by smearing the charges on small disks and using the superharmonicity of $-\log$, and the controls of Proposition 2.3 use the electric energy of a box. None of the arguments above uses the plane.

Discrepancy bounds from the electric energy. In the electric formalism, the classical bound on the discrepancy of a single box is one-sided and uses a single scale. If I carries an excess of n points, so do (up to $\frac{n}{2}$) the slightly larger boxes, of lengths in $[\ell, \ell + \mathtt{c}n]$, since points only increase D. By Gauss's law, each corresponding half-circle carries a flux $\ge \frac{n}{2}$, hence a field energy $\gtrsim \frac{n^2}{\ell}$, on a range of radii of length $\asymp n$. The total cost is $\gtrsim \frac{n^3}{\ell}$, and an energy of order ℓ gives $|\mathsf{Dis}| \lesssim \ell^{2/3}$. **[Exponent reconstructed from the argument: to be checked against Sandier–Serfaty.]** The chaining over all scales of Proposition 2.3 improves this to $\sqrt{\ell \log \ell}$, which is optimal for the worst box.

Averaging windows versus good contours. In the electric formalism, a discrepancy is a *flux*: by Gauss's law, $\mathsf{Dis}(\mathbf{X}, I)$ is the flux of the field through the boundary of a domain of the plane containing I. To bound it, one averages over the position of the contour (Fubini in the radius, or in the height of a strip), finds a good contour on which the field is not too large, and applies Cauchy–Schwarz with $\int |\nabla u|^2$. On the line there is no contour to move, and the comparison of averages (Lemma 2.5) plays both roles: *the average of D over a window replaces the flux through a contour, and the averaging over the window replaces the choice of a good contour.* Differences of averages over nested windows replace fluxes through nested contours, and the chaining over scales, which the geometry of the plane performed implicitly, is done by hand.

References

The master formula is the real-line form of the identity between the logarithmic energy and the $\dot{H}^{1/2}$ norm, see [CS07] for the extension point of view and [DNPV12] for fractional Sobolev spaces and their embeddings. The splitting formula, Onsager's lemma and the next-order energy of log-gases are in [SS15, Ser24].

Exercises

Exercise 2.9. Prove Lemma 2.1. *Hint: write* $\sum_i \delta_{x_i} = N\mu_{\mathrm{sc}} + \nu_N$*, expand, use the Euler–Lagrange equation, and* $-\log|x-y| = -\log|x'-y'| + \log N$.

Exercise 2.10. Let $q \in \mathbb{Z}$ and $L \geq 2$.

1. Let $D = q\mathbb{1}_{[0,L]}$, smoothed at unit scale at both ends (a charge q near 0, compensated near L). Show that $\frac{1}{4}\iint_{|x-y|>1} \frac{(D(x)-D(y))^2}{(x-y)^2} = q^2 \log L + O(q^2)$.
2. Show that the same holds, with $\frac{1}{2}q^2 \log L$, if the charge q is localized near 0 and compensated uniformly on $[0, L]$ ($D(x) = q(1 - x/L)$ on $[0, L]$).
3. Show that if $D = q\,\theta(x/L)$ with θ a smooth bump, the cost is $O(q^2)$, without logarithm.
4. Let $D(x) = q$ for $|x| \leq 1$, $D(x) = q\frac{\log(L/|x|)}{\log L}$ for $1 \leq |x| \leq L$, and $D = 0$ for $|x| \geq L$. The charge $+q$ sits in $[-L, -1]$ with density $\frac{q}{|x| \log L}$, the same mass in each dyadic shell, and is compensated symmetrically in $[1, L]$. Show that $\mathsf{Dis}([-L, 0]) = q$ and that the cost is $O(\frac{q^2}{\log L})$ (use the dilation invariance of the kernel). Deduce that, for a box of length L and an energy of order L, a discrepancy of order $\sqrt{L \log L}$ is possible: the factor $\log|I|$ in (2.3) cannot be removed.

The cost of a net charge depends on how it is distributed: the logarithm comes from localizing the charge at a scale much smaller than the distance to its compensation. This will be the heart of Section 4.

References

[CS07] Luis Caffarelli and Luis Silvestre. An extension problem related to the fractional Laplacian. *Communications in Partial Differential Equations*, 32(8):1245–1260, 2007.

[DNPV12] Eleonora Di Nezza, Giampiero Palatucci, and Enrico Valdinoci. Hitchhiker's guide to the fractional Sobolev spaces. *Bulletin des Sciences Mathématiques*, 136(5):521–573, 2012.

[Ser24] Sylvia Serfaty. Lectures on Coulomb and Riesz gases, 2024. arXiv:2407.21194.

[SS15] Etienne Sandier and Sylvia Serfaty. 1D log gases and the renormalized energy: crystallization at vanishing temperature. *Probability Theory and Related Fields*, 162(3–4):795–846, 2015.