1 Introduction

These lecture notes are devoted to the mathematical study of two physical phenomena that have close mathematical connections: vortices in the Ginzburg–Landau model of superconductivity on the one hand, and classical Coulomb gases on the other hand. A large part of the results we shall present originates in joint work with Etienne Sandier (for Ginzburg–Landau and two dimensional Coulomb gases) and in joint work with Nicolas Rougerie (for higher dimensional Coulomb gases), recently revisited in work with Mircea Petrache. In order to simplify the presentation, we have chosen to present the material in reverse chronological order, starting with the more recent results on Coulomb gases which are simpler to present, and finishing with the more complex study of vortices in the Ginzburg–Landau model. But first, in this introductory chapter, we start by briefly presenting the two topics and the connection between them.

1.1 From the Ginzburg–Landau model to the 2D Coulomb gas

1.1.1 Superconductivity and the Ginzburg–Landau model. The Ginzburg–Landau model is a very famous physics model for superconductivity. Superconductors are certain metallic alloys, which, when cooled down below a very low *critical temperature*, lose their resistivity and let current flow without loss of energy. This is explained at the microscopic level by the creation of superconducting electron pairs called Cooper pairs (Bardeen–Cooper–Schrieffer or BCS theory), and superconductivity is a macroscopic manifestation of this quantum phenomenon. The Ginzburg–Landau theory, introduced on phenomenological grounds by Ginzburg and Landau in the 1950's, some forty years after superconductivity had first been discovered by Kammerling Ohnes in 1911, has proven amazingly effective in describing the experimental results and predicting the behavior of superconductors. It is only very recently that the Ginzburg–Landau theory [GL] has been rigorously (mathematically) derived from the microscopic theory of Bardeen–Cooper–Schrieffer [BCS], also dating from the 50's, by Frank, Hainzl, Seiringer and Solovej [FHSS].

These superconducting alloys exhibit a particular behavior in the presence of a magnetic field: the superconductor levitates above the magnet. This is explained by the *Meissner effect*: the superconductor expells the magnetic field. This only happens when the external field h_{ex} is not too large. There are three critical fields $H_{c_1}, H_{c_2}, H_{c_3}$ for which *phase transitions* occur. Below the *first critical field H_{c_1}*, the material is everywhere superconducting. At H_{c_1}, one first observes local defects of superconductivity, called *vortices*, around which a superconducting loop of current circulates. As h_{ex} increases, so does the number of vortices, so that they
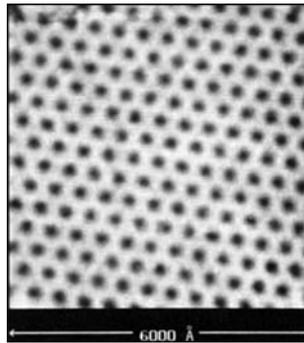
Figure 1.1. Abrikosov lattice, H. F. Hess et al. Bell Labs *Phys. Rev. Lett.* 62, 214 (1989).

become densely packed in the sample. The vortices repel each other, while the magnetic field confines them inside the sample, and the result of the competition between these two effects is that they arrange themselves in a particular *triangular lattice* pattern. It was predicted by Abrikosov [Abri], and later observed experimentally, that there should be periodic arrays of vortices appearing in superconductors, and this was later observed experimentally (Abrikosov and Ginzburg earned the 2003 Nobel Prize for their discoveries on superconductivity), cf. Fig. 1.1 (for more pictures, see www.fys.uio.no/super/vortex/).

These triangular lattices (originally Abrikosov predicted a square lattice but he had made a small mistake) then became known as *Abrikosov lattices*. A part of our study, detailed in this course, is aimed towards understanding why this particular lattice appears.

The second and third critical fields correspond respectively to the loss of superconductivity in the sample bulk and to the complete loss of superconductivity. These two transitions are not the focus of our study, and for more mathematical details on them, we refer to the monograph by Fournais and Helffer [FH1]. For a physics presentation of superconductivity and the Ginzburg–Landau model we refer to the standard texts [SST,DeG,Ti], for a mathematical presentation one can see [SS4,FH1] and references therein.

In non-dimensionalized form and in a simply connected domain Ω of the plane, the model proposed by Ginzburg–Landau can be written as the functional

$$G_\varepsilon(u, A) = \frac{1}{2} \int_\Omega |(\nabla - iA)u|^2 + |\text{curl } A - h_{\text{ex}}|^2 + \frac{(1 - |u|^2)^2}{2\varepsilon^2}. \qquad (1.1)$$

This may correspond to the idealized situation of an infinite vertical cylindrical sample of cross-section Ω and a vertical external field of intensity h_{ex}, or to a thin film.

Here

- $u \colon \Omega \to \mathbb{C}$, usually denoted by ψ in the physics literature, is called the *order parameter*. Its modulus (the density of Cooper pairs of superconducting electrons in the BCS theory) indicates the local state of the material: where

$|u| \approx 1$, the material is in the superconducting phase, where $|u| \approx 0$ is in the normal phase. The vortices correspond to isolated zeroes of u, and since u is complex-valued each zero carries an integer topological degree, like a "topological charge."

- $A \colon \Omega \to \mathbb{R}^2$ is the vector potential of the magnetic field $h = \operatorname{curl} A$ (defined by $\operatorname{curl} A := \partial_1 A_2 - \partial_2 A_1$), which is thus a real-valued function.

- The parameter $h_{\mathrm{ex}} > 0$ is the intensity of the applied (or external) magnetic field.

- The parameter $\varepsilon > 0$ is a material constant, corresponding to the ratio between characteristic lengthscales of the material (the *coherence length* over the *penetration depth*). We will be interested in the asymptotic regime $\varepsilon \to 0$. The functional is generally expressed in the physics literature in terms of the inverse of the constant ε, denoted by κ, and called the *Ginzburg–Landau parameter*. Materials with high-κ (the case we are interested in) are sometimes called "extreme type-II superconductors," and the limit $\kappa \to \infty$ is often called the *London limit*.

When considering the problem of minimizing the functional G_ε, a heuristic examination leads to observing that:

- The term $(1 - |u|^2)^2$ favors u close to 1, hence u should not vanish too often, especially as $\varepsilon \to 0$. A dimensional analysis in fact shows that the regions where $|u|$ is small have lengthscale ε.

- The quantity $|\operatorname{curl} A - h_{\mathrm{ex}}|^2$ is smaller when $\operatorname{curl} A = h \approx h_{\mathrm{ex}}$, that is, when the magnetic field penetrates the material so that the induced magnetic field equals the external magnetic field.

Minimizers and critical points of the Ginzburg–Landau functional without boundary constraints solve the associated set of Euler–Lagrange equations, called the Ginzburg–Landau equations:

$$
\text{(GL)} \quad \begin{cases} -\nabla_A^2 u = \frac{1}{\varepsilon^2} u(1 - |u|^2) & \text{in } \Omega \\ -\nabla^\perp h = \langle i u, \nabla_A u \rangle & \text{in } \Omega \end{cases}
$$

where $\nabla_A := \nabla - iA$, $\langle \cdot, \cdot \rangle$ denotes the scalar product in $\mathbb{C}$ as identified with $\mathbb{R}^2$, $\nabla^\perp = (-\partial_2, \partial_1)$ and again $h = \operatorname{curl} A$; with natural boundary conditions

$$
\begin{cases} \nabla_A u \cdot \nu = 0 & \text{on } \partial\Omega \\ h = h_{\mathrm{ex}} & \text{on } \partial\Omega. \end{cases}
$$

1.1.2 Reduction to a Coulomb interaction.
More details on the analysis of the Ginzburg–Landau model will be given in Chapter 7, which will be devoted to it, but for now let us try to explain the Coulombic flavor of the phenomenon.

In the regime with vortices (for $H_{c_1} \leq h_{\text{ex}} \ll H_{c_2}$), formal computations that will be better detailed in Chapter 7 show that in the asymptotic regime $\varepsilon \to 0$, the functional $G_\varepsilon(u, A)$ behaves as if it were:

$$G_\varepsilon(u, A) \approx \frac{1}{2} \int_\Omega |\nabla h|^2 + |h - h_{\text{ex}}|^2, \text{ where } h = \text{curl } A \qquad (1.2)$$

with what is known in the physics literature as the *London equation*:

$$\begin{cases} -\Delta h + h \approx 2\pi \sum d_i \delta_{a_i}^{(\varepsilon)} & \text{in } \Omega \\ h = h_{\text{ex}} & \text{on } \partial\Omega, \end{cases} \qquad (1.3)$$

where the a_i's are the centers of the vortices of u and the coefficients $d_i \in \mathbb{Z}$ their (topological) degrees. One should think of $\delta_{a_i}^{(\varepsilon)}$ as being formally a Dirac mass at a_i, smoothed out at the scale ε, or some approximation of it. A large part of our analysis in [SS4, SS7] is devoted to giving rigorous statements and proofs of these heuristics.

Inserting the London equation (1.3) into the approximation (1.2) leads to the following electrostatic analogy:

$$G_\varepsilon(u, A) \approx \frac{1}{2} \int_{\Omega \times \Omega} G_\Omega(x, y) \Big(2\pi \sum_i d_i \delta_{a_i}^{(\varepsilon)} - h_{\text{ex}}\Big)(x) dx \Big(2\pi \sum_i d_i \delta_{a_i}^{(\varepsilon)} - h_{\text{ex}}\Big)(y) dy \qquad (1.4)$$

where G_Ω is a Green kernel (or more accurately, Yukawa or screened Green kernel), solution to

$$\begin{cases} -\Delta G_\Omega + G_\Omega = \delta_y & \text{in } \Omega \\ G_\Omega = 0 & \text{on } \partial\Omega. \end{cases} \qquad (1.5)$$

This kernel is logarithmic to leading order: we may write

$$G_\Omega(x, y) = -\frac{1}{2\pi} \log|x - y| + R_\Omega(x, y) \qquad (1.6)$$

where R_Ω is a nonsingular function of (x, y). Approximating G_Ω by $-\frac{1}{2\pi} \log$ gives that the leading terms in (1.4) are

$$G_\varepsilon(u, A) \approx -\pi \sum_{i,j} d_i d_j \log|a_i - a_j| \qquad (1.7)$$

which is a sum of pairwise logarithmic or Coulombic interactions, weighted by the degrees d_i. Two such topological charges repel each other when they have the same sign, and attract each other if they have different signs. Rigorously, this is of course wrong, because we have replaced the smoothed out Diracs by true Dirac masses, leading to infinite contributions when $i = j$ in (1.7). One needs to analyze more carefully the effects of the smearing out, and to remove the infinite self-interaction of each "charge" at a_i in (1.7). One also needs to retain the interaction of these charges with the "background charge" $-h_{\text{ex}} dx$ appearing in (1.4). This is what leads to the analogy with the Coulomb gas that we will define and describe just below.

When looking for a model that retains these features: Coulombic interactions of points, combined with the confinement by a background charge, the simplest is to consider a discrete model with all charges equal to 1, and consider the Hamiltonian of a Coulomb gas with confining potential in dimension 2:

$$H_n(x_1, \ldots, x_n) = -\sum_{i \neq j} \log |x_i - x_j| + n \sum_{i=1}^{n} V(x_i), \tag{1.8}$$

where $x_i \in \mathbb{R}^2$, V is the confining potential (smooth, growing faster than $\log |x|$ at infinity), and the number of points n tends to infinity.

It turns out that this much simpler model (compared to G_ε) does retain many of the essential features of the vortex interaction, and is also of independent interest for physics and mathematics, as we will see. The study of (1.8) and its higher-dimensional analogues will occupy the largest part of these notes. We will then see how to use the perspective and knowledge gained on this to analyze the Ginzburg–Landau model (again, this is the reverse of the literature chronology, since we first studied the Ginzburg–Landau model and then adapted our analysis to the Coulomb gas situation!).

1.2 The classical Coulomb gas

1.2.1 The general setting. The Hamiltonian given by (1.8) corresponds to the energy of a gas of charged particles in $\mathbb{R}^2$ interacting via the Coulomb kernel in two dimensions. To be more precise, $-\log |x - y|$ is a multiple of the Coulomb kernel (or the fundamental solution of the Laplacian in the plane) in dimension 2. The counterpart in higher dimension corresponds to the d-dimensional Coulomb kernel, which is a multiple of $|x|^{2-d}$ for $d \geq 3$. The Hamiltonian of a classical Coulomb gas in any dimension $d \geq 2$ is thus given by

$$H_n(x_1, \ldots, x_n) = \sum_{i \neq j} g(x_i - x_j) + n \sum_{i=1}^{n} V(x_i) \tag{1.9}$$

where

$$g(x) = \begin{cases} -\log |x| & \text{for } d = 2 \\ \frac{1}{|x|^{d-2}} & \text{for } d \geq 3. \end{cases} \tag{1.10}$$

The statistical mechanics of a Coulomb gas, also called in physics a *two-dimensional one-component plasma*, is described by the corresponding Gibbs measure:

$$d\mathbb{P}_{n,\beta}(x_1, \ldots, x_n) := \frac{1}{Z_{n,\beta}} e^{-\frac{\beta}{2} H_n(x_1, \ldots, x_n)} dx_1 \ldots dx_n \tag{1.11}$$

where $\beta > 0$ is the *inverse temperature* and $Z_{n,\beta}$ is a normalization constant, the *partition function*, defined by

$$Z_{n,\beta} = \int_{(\mathbb{R}^d)^n} e^{-\frac{\beta}{2} H_n(x_1,\ldots,x_n)} dx_1 \ldots dx_n.$$

The probability measure $\mathbb{P}_{n,\beta}$ gives the probability of finding the particles at $(x_1, \ldots, x_n)$ at (inverse) temperature β. The object of statistical mechanics is then to analyze possible transitions in the types of states that can be effectively observed (i. e. those that have probability 1 or almost 1), according to the value of the inverse temperature β (e. g. transitions from ordered to disordered states at critical temperatures, such as liquid to solid phases etc). For general reference, we refer to standard statistical mechanics textbooks such as [Huan], and with increasing order of specificity to the books [HMD, Fo].

This model is one of the most basic statistical mechanics models not confined to a lattice, and it is considered difficult because of the long-range nature of the electrostatic interaction. Moreover, it can play the role of a toy model of the structure of matter, even if it is a purely classical – and not quantum – model. Studies in this direction include [SM, LieOx, AJ, JLM, PenSm].

The macroscopic distribution of the points as their number n goes to infinity is well understood and relatively simple to derive. This will be the object of Chapter 2. On the other hand, their microscopic distribution, more precisely the one seen at the scale $n^{-1/d}$, is less understood, and will be the main object of these lectures.

Let us now see some more specific motivations for studying the classical Coulomb gas, many of them being specific to dimension $d = 2$.

1.2.2 Two-dimensional Coulomb gas.

This is the setting that is the closest to the Ginzburg–Landau setting, as we discussed above. In this setting, the microscopic distribution of the points in the plane is expected to crystallize (most likely in the Abrikosov lattice triangular pattern) at low temperature. In fact there is some controversy in the physics literature as to whether there is a finite temperature phase transition for this crystallization which is numerically observed, cf. e. g. [BST, Sti, AJ].

Vortices in superfluids and superconductors   A first motivation for studying the two-dimensional Coulomb gas is the analysis of vortices in the Ginzburg–Landau model of superconductivity, but also more generally of vortex systems in classical fluids [CLMP], in quantum fluids such as in superfluids or Bose–Einstein condensates [CPRY], and in fractional quantum Hall physics [Gir, RSY1, RSY2]. All these systems share a lot of mathematics in common, and it is also of interest to understand their statistical physics (critical temperatures and phase transitions).

Fekete sets   This motivation no longer comes from physics but rather from a very different area of mathematics: interpolation theory. Fekete points are defined to be

points that maximize the quantity

$$\prod_{1 \leq i < j \leq n} |x_i - x_j|, \tag{1.12}$$

among all families of n points defined on a certain subset of $\mathbb{R}^d$, or a manifold (or any metric space, replacing moduli by distances). The Fekete points have the property of minimizing the error when interpolating a function by its value at points, see [SaTo] for reference, or [SK] for more details on the motivation, and [BrGr, Gra] for surveys of recent results on the sphere. A whole literature is also devoted to understanding Fekete points on complex manifolds, possibly in higher dimension, see e. g. [Ber, BBN, LevOrC] and references therein.

Of course, in the setting of Euclidean space, maximizing (1.12) is equivalent to minimizing the logarithmic interaction

$$- \sum_{1 \leq i \neq j \leq n} \log |x_i - x_j|$$

which takes us back to the setting of a two-dimensional Coulomb gas. Indeed, Fekete sets confined to a set $K \subset \mathbb{R}^d$ correspond to minimizers of H_n with V taken to be 0 in K and $+\infty$ in K^c. Minimizers of (1.8) for general V's are in fact called *weighted Fekete sets*, also defined as maximizers of

$$\prod_{1 \leq i < j \leq n} |x_i - x_j| \prod_{i=1}^{n} e^{-\frac{n}{2} V(x_i)}$$

where V is the weight. For definitions and the connection to logarithmic potential theory, see again [SaTo] and references therein. Weighted Fekete sets are also naturally related to the theory of weighted orthogonal polynomials (cf. the surveys [Sim, Ko] or again [SaTo]).

The correspondence can also be made via a mapping, e. g. the important question of finding the Fekete points on the (2-)sphere is equivalent, by stereographical projection, to studying the weighted Fekete sets on $\mathbb{R}^2$ with weight $V(x) = \frac{1}{2} \log(1 + |x|^2)$, for details see [Ha, Bet].

Random matrix theory Random matrix theory (RMT) is a relatively old theory, pioneered by statisticians and physicists such as Wishart, Wigner and Dyson, and originally motivated by an understanding of the spectrum of heavy atoms, see [Me]. For a more recent mathematical reference see [AGZ, D, Fo]. An important model of random matrices is the so-called *Ginibre ensemble* [Gin]: the law is that of an $n \times n$ complex matrix whose coefficients are i.i.d. complex normal random variables. The main question asked by RMT is: what is the law of the spectrum of a large random matrix? In the case of the Ginibre ensemble, the law is known exactly: upon rescaling the (complex) eigenvalues $x_1, \ldots, x_n$ by a factor $\frac{1}{\sqrt{n}}$, it is given by the following

density:

$$d\mathbb{P}_n(x_1, \ldots, x_n) = \frac{1}{Z_n} e^{-H_n(x_1, \ldots, x_n)} dx_1 \ldots dx_n \qquad (1.13)$$

with

$$H_n(x_1, \ldots, x_n) = -\sum_{i \neq j} \log |x_i - x_j| + n \sum_{i=1}^{n} |x_i|^2 \qquad (1.14)$$

and Z_n a normalization constant. We recognize in (1.14) the 2D Coulomb gas Hamiltonian with potential $V(x) = |x|^2$, and the law $d\mathbb{P}_n$ is the Gibbs measure (1.11) at inverse temperature $\beta = 2$. This analogy between random matrices and the statistical mechanics of Coulomb gases was first noticed by Wigner [Wi] and Dyson [Dy], see [Fo] for more on this link. Writing the law in the form (1.13) immediately displays the phenomenon of *repulsion of eigenvalues*: eigenvalues in the complex plane interact like Coulomb particles, i. e. they do not "like" to be too close and repel each other logarithmically.

At this specific temperature $\beta = 2$, the law of the spectrum acquires a special algebraic feature: it becomes a determinantal process, part of a wider class of processes (see [HKPV, Bor]) for which the correlation functions are explicitly given by certain determinants. This allows for many explicit algebraic computations. However, many relevant quantities that can be computed explicitly for $\beta = 2$ are not exactly known for the $\beta \neq 2$ case, even in the case of the potential $V(x) = |x|^2$. In this course, in contrast, we will work for any β, and with a wide class of potentials.

1.2.3 The one-dimensional Coulomb gas and the log gas.
We have not mentioned yet the one-dimensional Coulomb gas, which corresponds to (1.9) with the Coulomb kernel (up to a constant) $g(x) = |x|$. The reason we will not be interested in it is because it has already been well-understood [Le1, Le2, Ku, BraLie, AlMu]. It can be "solved" almost explicitly and crystallization at zero temperature is established.

We are interested however in another one-dimensional model (i. e. with points $x_i \in \mathbb{R}$), where the two-dimensional logarithmic interaction $g(x) = -\log |x|$ is used in (1.9). This is usually called a *log gas*, and its motivation also comes from Random Matrix Theory (see [Fo]): one-dimensional counterparts to the Ginibre ensemble are the Gaussian Unitary Ensemble (GUE) and the Gaussian Orthogonal Ensemble (GOE), which are symmetric analogues of it. The law of the GUE (resp. the GOE) is that of an $n \times n$ matrix whose coefficients are complex (resp. real) normal random variables, independent up to a Hermitian (resp. symmetry) condition. Because of the Hermitian or symmetric nature of the matrix, its eigenvalues lie on the real line (hence the one-dimensionality of the model), but they still repel each other logarithmically: again, the law of the spectrum (the distribution of eigenvalues) can be given explicitly by the following density on $\mathbb{R}$:

$$d\mathbb{P}_n(x_1, \ldots, x_n) = \frac{1}{Z_n} e^{-\frac{\beta}{2} H_n(x_1, \ldots, x_n)} dx_1 \ldots dx_n, \qquad (1.15)$$

where H_n is still defined as

$$H_n(x_1, \ldots, x_n) = -\sum_{i \neq j} \log|x_i - x_j| + n \sum_{i=1}^{n} |x_i|^2, \quad x_i \in \mathbb{R} \qquad (1.16)$$

with $\beta = 1$ for the GOE and $\beta = 2$ for GUE. This is thus a particular case of a log gas, at specific temperature $\beta = 1$ or 2 and with quadratic potential, and the phenomenon of *repulsion of eigenlevels* is visible in the same way as for the Ginibre ensemble. Again, in these cases of the GOE and GUE, a lot about (1.15) can be understood and computed explicitly thanks to the underlying random matrix structure and its determinantal nature. In fact the global and local statistics of eigenvalues are completely understood.

Considering the coincidence between a statistical mechanics model and the law of the spectrum of a random matrix model for several values of the inverse temperature, it is also natural to ask whether such a correspondence exists for any value of β. The answer is yes for $\beta = 4$. It corresponds to the Gaussian Symplectic Ensemble (GSE) of Hermitian matrices with quaternionic coefficients, and for any β a somehow complicated model of tridiagonal matrices can be associated to the Gibbs measure of the one-dimensional log gas at inverse temperature β, see [DE]. This and other methods allow us again to compute a lot explicitly, and to derive that the microscopic laws of the eigenvalues are those of a so-called *sine-β process* [VV].

Generally speaking, much is known for log gases in one dimension, for any value of β and a wide class of potentials V. In particular, a lot of attention has been devoted to proving that many of the features of the system at the microscopic scale are *universal*, i.e. independent of the particular choice of V. For recent results, see [BEY1, BEY2, Shch1, Shch2, Shch3, BoGui, BG2, BFG]. In contrast, the analogue is true in dimension 2 only for $\beta = 2$ [Gin, BSi, AHM]. Thus the topic of log/Coulomb gases does not seem reducible to just a subset of Random Matrix Theory. This is of course even more true in dimension 3 and higher, which we will also treat, and where we leave the realm of RMT.