ETH Zürich

DUALITY

Bachelor's Thesis

written by

Franco de Bardeci

under the supervision of

Prof. Richard Pink

Abstract

This thesis has the purpose of illustrating the Principle of Duality. We present various examples in different fields of mathematics, including Geometry, Algebra, Graph Theory and Probability Theory.

February 16, 2016

Contents

1 Introduction

Duality is a principle that appears throughout mathematics. It can manifest itself in great elegance and be very useful at the same time. It often starts out as a certain symmetry at the axiomatic level, extending across the theory. In many fields of mathematics duality plays a fundamental role that might also remain present in the applications of those fields.

An undergraduate student may encounter this principle many times, for example when considering dual vector spaces or dual norms. Also perhaps, in projective geometry. Sometimes without explicit mention of the fact that a duality is involved, as for the case of the Fourier transform.

There is no precise definition of what a duality is, but the following is the prototype for many of them. Say we have a mathematical structure that in particular involves the sets A and B, and is defined by a set of axioms. Each axiom is a statement that may refer to the sets A and B. For each axiom we consider the *dual* statement, which is obtained by interchanging the roles of A and B. This means that whenever A is mentioned, it is replaced by B, and vice versa. For example, if at some point one of the axioms reads: "For each element of A there is an element of B such that...", then the dual statement reads: "For each element of B there is an element of A such that...". One has a duality, if the dual statement of each of the axioms is derivable from the axioms. When this is the case, any statement implies its dual. Thus, by proving a theorem, unless the dual statement of the theorem is identical to the theorem itself, we get a second theorem for free! In the following sections we consider various examples of dualities. The sets A and B will for example consist of the points and lines in the projective plane, or the vectors and linear functionals of a finite vector space.

2 Projective Geometry

We begin by considering the Principle of Duality on the projective plane. A projective plane consists of a set of *points* and a set of *lines*, together with a relation between the two, such that a certain set of axioms is satisfied. This relation is called the *incidence relation*, and two elements satisfying this relation are called *incident*. Intuitively, a point and a line are incident if the point lies on the line. The following set of axioms for this geometry can be found in Coxeter [1], Chapter 3. We introduce them along with the relevant terms.

Axiom 2.1. Any two lines are incident with at least one point.

Axiom 2.2. Any two distinct points are incident with exactly one line.

From these first two axioms, it follows that also two distinct lines are incident with exactly one point. For any two distinct lines l and m, we can therefore define their intersection as being the unique point incident with both lines, which we denote by $l \cdot m$. Analogously, for any two distinct points P and Q, we define the line PQ, as the unique line incident with P and Q.

A set of points is called *collinear* if there is a line such that every point in the set is incident with this line. A set of lines is called *concurrent* if there is a point such that every line in the set is incident with this point.

Axiom 2.3. There exist four points of which no three are collinear.

Four points together with six distinct lines joining each pair of points are called a *complete quadrangle*. Pairs of these lines that do not meet at one of those four points are called *opposite*. We can form three different opposite pairs, the points where these pairs intersect are called the *diagonal points*.

Axiom 2.4. The three diagonal points of a complete quadrangle are never collinear.

The set of points incident with a given line is called the *range* of that line. The set of lines incident with a given point is called the *pencil* of that point. Given a line and a point which are not incident, there is a natural one-to-one correspondence between the elements of the range and the elements of the pencil, corresponding elements being incident. Consider a finite sequence whose elements alternate between lines and points, such that any two consecutive elements are never incident, and any two elements that differ by two positions are distinct. Considering the one-to-one correspondence between range and pencil of consecutive elements, we can define a transformation from the range or pencil of the first element to the range or pencil of the last element. Such transformation is called a *projectivity*.

Axiom 2.5. If a projectivity leaves invariant three distinct points on a line, it leaves invariant every point on the line.

Three non-collinear points A, B, C together with the lines AB, BC, and CA are called a *triangle* denoted by ABC.

Axiom 2.6. Let ABC and $A'B'C'$ be triangles such that all six points, as well as the six lines are different. If the lines AA', BB', and CC' are concurrent, then the points $AB \cdot A'B'$, $BC \cdot B'C'$, and $CA \cdot C'A'$ are collinear.

The last axiom is also known as Desargues's Theorem. Although we state it as an axiom here, it might be a theorem if another set of axioms was chosen.

The reader may have noticed that a certain symmetry emerges. As noted before, Axiom 2.2 still holds if we interchange the words points and lines. Namely, if two distinct lines are given, Axiom 2.1 asserts that there is at least one point of intersection. If there was more than one point, then by Axiom 2.2 the lines would be equal, and therefore the point of intersection is unique. Also, we observe that the definition of collinear and concurrent are related in the same way. This relation is precisely the *duality*, and pairs of statements, definitions or objects that are related in this way are said to be *dual*. If we were able to prove the dual statement of each axiom, we would know that the dual statement of any theorem must also hold.

This is indeed the case. The dual of Axiom 2.1 follows from Axiom 2.2 and the fact that more than one point exists, which follows from Axiom 2.3.

For the dual of Axiom 2.3, consider the four points given by this axiom. Choose an ordering of the points and join consecutive points with a line as well as the first point with the last point. We obtain four lines. First note that the lines are all distinct. For if two pairs were to be joined by the same line, there would be at least three collinear points. Suppose now that three of the lines are concurrent. For each line consider the two points that gave rise to them. The only way not to count more than four points is that the three lines share one of these points. But by the way these lines were constructed, such a point can be shared by two lines at most.

Instead of dealing with a complete quadrangle, the dual of Axiom 2.4 deals with a set of four lines, such that every pair of lines intersects at one of six distinct points. This is called a *complete quadrilateral*. The diagonal points correspond to the three other lines that join pairs of this six points, they are called the *diagonal lines*. We have to show that these are not concurrent. Consider a complete quadrilateral formed by the lines a, b, c, d. Assume the diagonal lines all meet at a point P. Say the quadrilateral has the vertices A, B, C, D, E, F. The vertices A, B, C, D define a complete quadrangle. Its edges are precisely the lines a, b, c, d, and two of the three diagonal lines of the quadrilateral. Its diagonal points are E, F and P. The diagonal line of the quadrilateral that is not an edge of the quadrangle is EF. The point P must also lie on this line, and therefore the diagonal points of the quadrangle are collinear, which is impossible by Axiom 2.4.

To prove the dual of Axiom 2.5, assume there is a projectivity on a pencil that leaves invariant every of three lines but not the entire pencil. Such a projectivity is given by a sequence $a_1, a_2, \ldots, a_n, a_1$. By shifting the sequence to $a_2, \ldots, a_n, a_1, a_2$, we get a counterexample to Axiom 2.5.

The dual statement of Axiom 2.6 is precisely its converse. We prove it. Say we have the triangles ABC and $A'B'C'$ as in the axiom. Set

$$P := AB \cdot A'B', \; Q := BC \cdot B'C', \; R := CA \cdot C'A'.$$

By assumption these points are collinear. Consider the triangles $AA'P$ and $CC'Q$. The lines AC, $A'C'$, and PQ meet at the point R. By Axiom 2.6 this means that B, B' and $S := AA' \cdot CC'$ are collinear. Therefore the lines AA', BB', and CC' meet at the point S.

This completes the proof that all dual statements of the axioms are true. Any theorem will therefore imply its dual, and hence every time we prove a theorem, unless it is self dual, we also get a second theorem. We illustrate this fact. Six distinct points and six distinct lines, where the points can be arranged in a sequence such that pairs of consecutive points, as well as the first and the last, are joined by one of the lines, and each line joins one of those pairs, are called a *hexagon*. The points of the hexagon are called its *vertices*, and the lines of the hexagon are called its *sides*. Two of the vertices that differ by three places in the sequence are called *opposite*. Likewise, two of the sides that join pairs that differer by three places are called *opposite*.

Theorem 2.7. *If the vertices of a hexagon lie alternately on two lines, the points of intersection of the three pairs of opposite sides of the hexagon are collinear.*

This is Pappus's Theorem. For a proof see Coxeter [1], Section 4.4. By duality one obtains:

Theorem 2.8. *If the sides of a hexagon alternately intersect two points, the lines that join the three pairs of opposite vertices of the hexagon are concurrent.*

3 Vector Spaces

Throughout let V be a vector space over a field K. A *linear functional* is a homomorphism from V to K. The set of all linear functionals, with pointwise addition and pointwise scalar multiplication, also becomes a vector space over K. It is called the *dual space* of V, and is denoted by V^*. How is here the term dual related to the previous discussion? Just as before points and

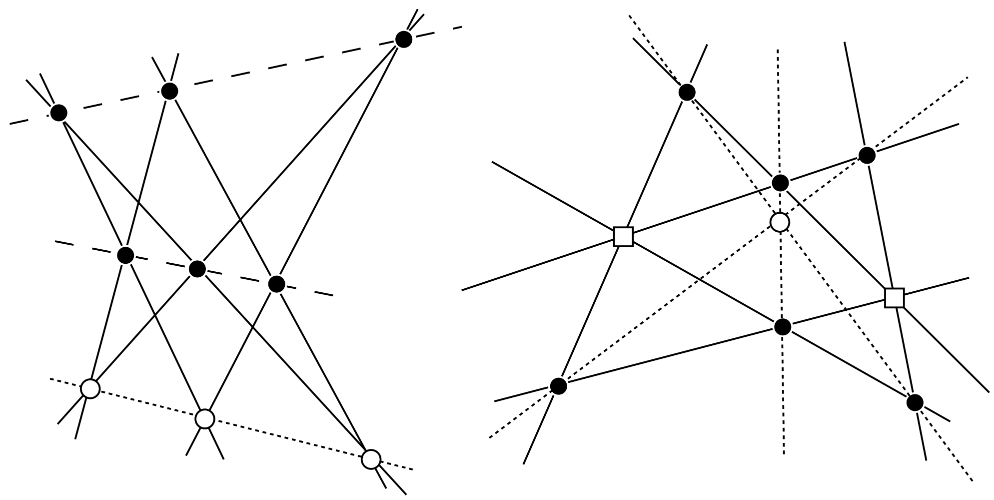

Figure 1: Two hexagons. The dark lines and dark points are their sides and vertices. The one the left satisfies the conditions in Pappus's Theorem. The one on the right, the conditions in the dual theorem.

lines were interchangeable, we will see that the same kind of relation holds between vectors and linear functionals. Consider the map

$$e : V \times V^* \to K, \ (v, f) \mapsto f(v).$$

By definition, for each $f \in V^*$ the map $V \to K$, $v \mapsto f(v)$ is a linear functional, and every linear functional can be obtained this way. Observe that for each $v \in V$ the map $\lambda_v : V^* \to K$, $f \mapsto f(v)$ is a linear functional as well. This already shows some symmetry. Define

$$\alpha : V \to V^{**}, \ v \mapsto \lambda_v.$$

This is an injective isomorphism. We will show that for finite dimensional vector spaces, it is an isomorphism. Therefore, in this case, the symmetry is perfect: V^* is the space of linear functionals on V by definition, and V is naturally isomorphic to the space of linear functionals of V^*.

Henceforth, let $v_1, ..., v_n$ form a basis for V. For a fixed index $i \in \{1, .., n\}$, consider the map $f_i : V \to K$ that takes a vector to its i-th coordinate. It is a linear functional and for any $i, j \in \{1, .., n\}$ we have: $f_i(v_j) = \delta_{ij}$.

Lemma 3.1. *The maps $f_1, ..., f_n$ form a basis for V^*.*

Proof. Let f be any element of V^*. By the homomorphism property for any $v \in V$ we have:

$$f(v) = f(f_1(v)v_1) + ... + f(f_n(v)v_n) = f(v_1)f_1(v) + ... + f(v_n)f_n(v).$$

Therefore f can be written as a linear combination in $f_1, ..., f_n$. Further, if $b_1 f_1 + ... + b_n f_n = 0$, then in particular we have:

$$b_1 f_1(v_i) + ... + b_n f_n(v_i) = 0, \ \forall i \in \{1, ..., n\}.$$

This implies $b_i = 0$ for every $i \in \{1, ..., n\}$, and hence $f_1, ..., f_n$ are lineary independent. $\square$

Proposition 3.2. *The above described map α is an isomorphism from V to V^{**}.*

Proof. It is directly verifiable that α is a homomorphism. Moreover, by Lemma 3.1, $f_1, ..., f_n$ form a basis for V^*. By the same process we implemented to obtain this basis from $v_1, ..., v_n$, a basis for V^{**} can be obtained from $f_1, ..., f_n$, say $\mu_1, ..., \mu_n$. We have $\mu_i(b_1 f_1 + ... + b_n f_n) = b_i$, for every $i \in \{1, ..., n\}$. Also

$$\lambda_{v_i}(b_1 f_1 + ... + b_n f_n) = b_1 f_1(v_i) + ... + b_n f_n(v_i) = b_i, \ \forall i \in \{1, ..., n\}.$$

Therefore μ_i and λ_{v_i} are identical and thus α sends basis elements to basis elements. $\square$

Note that even though when proving that α is an isomorphism we made use of bases, from its definition it is clear that α does not depend on any choice of basis. Note also that by Lemma 3.1 there is an isomorphism from V to V^*. However, this isomorphism may depend on the basis chosen.

4 Topological Groups

A *topological group* is a group with a topology such that the law of composition, as well as the inversion operation, are continuous functions. A *locally compact abelian group*, *LCA group* for short, is a topological group whose group structure is abelian and whose topology is locally compact and Hausdorff. Note that the latter is not represented in the name.

In a topological group, translation by an element of the group defines a homeomorphism from the topological group to itself. The open neighborhoods of the identity therefore describe the whole topology, other open sets being translations of them. To show that a homomorphism from a topological group to another is continuous, it suffices to show that it is continuous at the identity.

Throughout let G be an LCA group and $\mathbb{T}$ the topological group of complex numbers of absolute value 1 under multiplication. A *character* is a continuous homomorphism from G to $\mathbb{T}$. The set of all characters becomes a group by defining the composition of two characters to be the character obtained by pointwise multiplication. This group, together with a topology we describe next, is called the *dual group* of G and is denoted by $\hat{G}$.

Let X and Y be topological spaces, and F a set of functions from X to Y. For a compact set K of X, and an open set U of Y, we define $S(K, U)$ to be the set of all elements f of F such that $f(K) \subset U$. The collection of all sets of the form $S(K, U)$ is subbasis for a topology on F. It is called the *compact-open topology.* The dual group $\hat{G}$ is given the compact-open topology.

Lemma 4.1. *Let X be a topological space and Y a metric space. Let K be a compact set of X, let U be an open set of Y, and $f : X \to Y$ a continuous map that sends K to U. There is an $\epsilon > 0$ such that for any point of $f(K)$ the open ball of radius ϵ around that point is contained in U.*

Proof. For $y \in Y$ define $r(y)$ to be the supremum of all radii r', such that $B_{r'}(y) \subset U$. Suppose no such ϵ exists. Then there is a sequence $(y_n)_{n\in\mathbb{N}}$ in $f(K)$ such that $r(y_n) \to 0$ for $n \to \infty$. Since $f(K)$ is compact $(y_n)_{n\in\mathbb{N}}$ has a subsequence that converges to some point of $f(K)$, say y_0. But $r(y_0) > 0$, and hence we have a contradiction. □

Proposition 4.2. *The dual group is a topological group.*

Proof. Let $\phi : \hat{G} \times \hat{G} \to \hat{G}$ be the map that takes (χ_1, χ_2) to $\chi_1\chi_2^{-1}$. It suffices to show that ϕ is continuous. Consider an element of the subbasis, $S(K, U)$ say. Let (χ_1, χ_2) be any element of the preimage of $S(K, U)$ under ϕ. We will show that there is an open set in $\hat{G} \times \hat{G}$ that contains (χ_1, χ_2) and is mapped into $S(K, U)$ by ϕ.

As $\chi_1\chi_2^{-1}$ is continuous, the set $\chi_1\chi_2^{-1}(K)$ is compact. By Lemma 4.1 there is an $\epsilon > 0$ such that a ball of radius ϵ around any point of $\chi_1\chi_2^{-1}(K)$ is contained in U. Let $i \in \{1, 2\}$. As χ_i is continuous, for any point x in G there is a neighborhood such that the image under χ_i of its closure is contained in a ball of radius $\epsilon/4$. Let $(V_{i,j})_{j\in J_i}$ be a finite collection of such neighborhoods covering K. Let $K_{i,j}$ be the intersection of the closure of $V_{i,j}$ with K, and $U_{i,j}$ a ball of radius $\epsilon/4$ that contains the image of $K_{i,j}$ under χ_i. Consider the set:

$$\bigcap_{j\in J_1} S(K_{1,j}, U_{1,j}) \times \bigcap_{j\in J_2} S(K_{2,j}, U_{2,j}).$$

It is open in $\hat{G} \times \hat{G}$ and contains (χ_1, χ_2). Further, it is mapped into $S(K, U)$ by ϕ, since for every (ξ_1, ξ_2) in $\hat{G} \times \hat{G}$ and every $x \in G$ we have:

$$|\chi_1 \chi_2^{-1}(x) - \xi_1 \xi_2^{-1}(x)| \leq |\chi_1(x) - \xi_1(x)| + |\chi_2(x) - \xi_2(x)|.$$

□

Note that we have not used the fact that G is locally compact in the proof. Using this fact, we can prove that $\hat{G}$ is an LCA group. It follows directly from the definition that $\hat{G}$ is abelian. To see that it is Hausdorff, consider two different characters χ_1 and χ_2. They must differ at some element of G, say x. As $\mathbb{T}$ is Hausdorff, we can separate $\chi_1(x)$ from $\chi_2(x)$ with two open sets, U_1 and U_2. The open sets $S(\{x\}, U_1)$ and $S(\{x\}, U_2)$ then separate χ_1 from χ_2. To show local compactness, we use a generalization of Ascoli's Theorem which characterizes compact subspaces in certain function spaces. Before we state the theorem, we introduce a required definition.

For a topological space X and a metric space Y, a set of functions from X to Y is called *equicontinuous* at $x \in X$ if for every $\epsilon > 0$ there is a neighborhood V of x such that for every function f in the set, we have $f(V) \subset B_\epsilon(f(x))$. If the set of functions is equicontinuous at every $x \in X$, then its called *equicontinuous*. We define $C(X, Y)$ to be the set of all continuous functions from X to Y with the compact-open topology, and $\mathbb{T}^+$ to be the of elements of $\mathbb{T}$ that have positive real part.

Theorem 4.3. *Let X be a locally compact space and Y a compact metric space. A subset in $C(X, Y)$ has compact closure if and only if it is equicontinuous.*

A proof of the theorem can be found in Bourbaki [5]. It is obtained by combining Theorem 2 in Section 2 of Chapter 10, and Theorem 2 in Section 3 of the same chapter. Alternatively, it can be found in von Querenburg [6], Theorem 14.22.

Proposition 4.4. *The dual group is an LCA group.*

Proof. It sufficient to show that the identity element 1 of $\hat{G}$ has a compact neighborhood. Let C be a compact neighborhood of the identity element e of G. Consider the set $S(C, \mathbb{T}^+)$, it is an open neighborhood of 1. We show that it is equicontinuous, then the claim follows by Theorem 4.3.

Showing that $S(C, \mathbb{T}^+)$ is equicontinuous amounts to showing that it is equicontinuous at the identity e. Let φ be the branch of the argument function on $\mathbb{C}$ that takes values in $(-\pi, \pi]$. For $n \in \mathbb{N}$ let U_n be the set of all

$z \in \mathbb{T}$ such that $|\varphi(z)| < \pi/n$. Let $\phi : G \times G \to G$ be the composition map, the map that takes (χ_1, χ_2) to $\chi_1\chi_2$. Let V be an open neighborhood of e in G that is contained in C. As ϕ is continuous, $\phi^{-1}(V)$ is open, also it is a neighborhood of (e, e). Therefore, there is a standard basis element of the product topology of $G \times G$, say $W_1 \times W_2$, that is contained in $\phi^{-1}(V)$ and is a neighborhood of (e, e). The set $W_1 \cap W_2$ has the property, that the composition of any two of its elements is in C. By iterating this process, one can obtain for any $n \in \mathbb{N}$, a neighborhood V_n of e, such that the composition of any n elements of V_n is in C. Any open neighborhood of 1 in $\mathbb{T}$ contains a set of the form U_n. We show that V_n is mapped into U_n by any character in $S(C, \mathbb{T}^+)$, as any character maps e to 1, this shows that $S(C, \mathbb{T}^+)$ is equicontinuous. Assume the opposite, then there is a character χ in $S(C, \mathbb{T}^+)$ and $x \in V_n$ such that $|\varphi(\chi(x))| \geq \pi/n$. There is $k < n$ such that $\chi(x)^k \notin \mathbb{T}^+$. But $\chi(x)^k = \chi(x^k)$ and therefore, because x^k is an element of C, we have $\chi(x)^k \in \mathbb{T}^+$. □

In analogy to what we did for vector spaces we consider the map:

$$e : G \times \hat{G} \to \mathbb{T},\ (x, \chi) \mapsto \chi(x).$$

Lemma 4.5. *The map e is continuous.*

Proof. Consider the preimage of an arbitrary open set U under this map. Choose a point on the preimage, say (x_0, χ_0). Consider the open set $\chi_0^{-1}(U)$. Because G is locally compact we can choose a compact set K, and an open set V such that $\{x_0\} \subset V \subset K \subset \chi_0^{-1}(U)$. Then $V \times S(K, U)$ is an open set in $G \times \hat{G}$ and is mapped into U by e. □

By Proposition 4.4 the dual group is also an LCA group, we can therefore consider the dual of the dual: $\hat{\hat{G}}$. Consider the map e. For $x \in G$, the map $\lambda_x : \hat{G} \to \mathbb{T},\ \chi \mapsto \chi(x)$ is continuous. It is also a homomorphism and thus an element of $\hat{\hat{G}}$. We therefore have a well defined map:

$$\alpha : G \to \hat{\hat{G}},\ x \mapsto \lambda_x.$$

It is called the *Pontryagin map.*

Theorem 4.6. *The Pontryagin map is a homeomorphism and an isomorphism.*

This is the Pontryagin Duality Theorem. In contrast to its analog for finite vector spaces, it is hard to prove and we will not do so here. It was first proven by Soviet mathematician Lev Pontryagin, although in less generality. His results were reported at the International Mathematical Congress in Zurich in 1932, see Pontryagin [11]. There are many proofs available. For a proof involving abstract Fourier analysis see Rudin [8], Chapter 1, or Deitmar and Echterhoff [9], Chapter 3. However, the following partial result is not difficult.

Proposition 4.7. *The Pontryagin map is a continuous homomorphism.*

Proof. For $x_1, x_2 \in G$ we have $\lambda_{(x_1+x_2)} = \lambda_{x_1} + \lambda_{x_2}$ so that α is a homomorphism. We prove that it is also continuous. By definition, the sets of the form $S(K,U)$, where K is a compact set of $\hat{G}$ and U is an open set of $\mathbb{T}$, form a subbasis for the topology on $\hat{\hat{G}}$. For same K and U, let $\hat{S}(K,U)$ be the set of all $x \in G$ that are taken into U by any character in K. Then the preimage of $S(K,U)$ under α is precisely $\hat{S}(K,U)$. We show that every set of the form $\hat{S}(K,U)$, for K a compact set of $\hat{G}$ and U an open set of $\mathbb{T}$, is open. Take any point in $\hat{S}(K,U)$, say x_0. Consider the preimage of U under the map e. It is open. For any character χ in K we can therefore choose a standard basis set of the product topology around the point (x_0, χ) that is contained in $e^{-1}(U)$. Such a basis set consists of a product of two open sets, say $U_\chi \times V_\chi$. The set of all sets of the form V_χ, where $\chi \in K$, is an open cover of K. Choose a finite subcover indexed by $I \subset K$. Then $\cap_{\chi \in I} U_\chi$ is an open neighborhood of x_0 contained in $\hat{S}(K,U)$. $\square$

Consider the group of integers $\mathbb{Z}$ with the discrete topology. It is an LCA group. We calculate its dual. A homomorphism from $\mathbb{Z}$ to $\mathbb{T}$ is already defined by specifying the target of the element 1. As $\mathbb{Z}$ has the discrete topology, any such homomorphism is continuous. We therefore have a one-to-one correspondence between $\hat{\mathbb{Z}}$ and $\mathbb{T}$, defined by the map that takes a character to the target of the element 1. It is an isomorphism. We show that it is also a homeomorphism. For an arbitrary open set U of $\mathbb{T}$, the open set $S(\{1\}, U)$ corresponds to U, hence the map is continuous. As the topology of $\mathbb{Z}$ is discrete, its compact sets are exactly its finite sets. The sets of the form $S(\{a\}, U)$ therefore form a subbasis for the topology of $\hat{\mathbb{Z}}$. A character χ in $S(\{a\}, U)$ must take 1 to an $|a|$-th root of $\chi(a)$, or to the inverse of such a root. It follows that we can find an open set U' of $\mathbb{T}$ such that $S(\{1\}, U')$ contains χ and is contained in $S(\{a\}, U)$. As $S(\{1\}, U')$ corresponds to U', our correspondence also has continuous inverse, and thus $\mathbb{T}$ is the dual of $\mathbb{Z}$. By the Pontryagin Duality Theorem we just have also shown that $\mathbb{Z}$ is the dual of $\mathbb{T}$.

As another example, $\mathbb{R}$ turns out to be self dual. In abstract Fourier analysis, LCA groups are given a certain measure, and one considers complex valued Lebesgue integrable functions on them. The Fourier transform of such a function is then a function defined on the dual group. When taking the group to be $\mathbb{T}$, one obtains the special case of Fourier series. When taking $\mathbb{R}$, one obtains the classical Fourier transform. For more on abstract Fourier analysis see for example Rudin [8] or Deitmar and Echterhoff [9].

As the examples of $\mathbb{Z}$, $\mathbb{T}$ and $\mathbb{R}$ hint at, discrete and compact LCA groups are in duality. More precisely:

Proposition 4.8. *If G is compact then $\hat{G}$ is discrete.*

Proof. It suffices to show that the set containing only the trivial character is open in $\hat{G}$. Consider the open set $S(G, T^+)$. Let χ be a character in that set. Then $\chi(G)$ is a subgroup of $\mathbb{T}$ contained in $\mathbb{T}^+$. Let φ be the branch of the argument function on $\mathbb{C}$ that takes values in $(-\pi, \pi]$. Let a be any element of $\mathbb{T}^+$ different from 1. We have $0 < |\varphi(a)| < \pi/2$. There is an integer n such that $|\varphi(a^n)| = n|\varphi(a)| \in (\pi/2, \pi)$, and therefore $a^n \notin \mathbb{T}^+$. Hence, the only subgroup of $\mathbb{T}$ contained in $\mathbb{T}^+$ is the trivial subgroup. Therefore $S(G, T^+)$ contains the trivial character alone. $\square$

Proposition 4.9. *If G is discrete then $\hat{G}$ is compact.*

Proof. By Tychonoff's Theorem, the set of all functions from G to $\mathbb{T}$ endowed with the product topology is compact. The fact that G is discrete implies that the subbasis in the definition of the compact-open topology is the same as the standard subbasis of the product topology. Also because G is discrete, every function from G to $\mathbb{T}$ is continuous and thus $\hat{G}$ is the set of homomorphisms from G to $\mathbb{T}$. This is a closed subset of the set of all functions from G to $\mathbb{T}$ endowed with compact-open topology. For if a character χ is not a homomorphism, there are a and b in G such that $\chi(ab) \neq \chi(a)\chi(b)$. For a sufficiently small r we then have that

$$S\big(\{ab\}, B_r(\chi(ab))\big) \cap S\big(\{a\}, B_r(\chi(a))\big) \cap S\big(\{b\}, B_r(\chi(b))\big)$$

contains no homomorphism. $\square$

By the Pontryagin Duality Theorem the converse of each of the last two propositions holds as well.

5 Category Theory

A category is a collection of *objects*, together with a collection of *arrows*, and three operations that satisfy certain axioms. The first operation takes an arrow, say f, and returns an object called the *domain* of the arrow, denoted by $d(f)$. The second operation takes the arrow, and returns an object called the *codomain* of the arrow, denoted by $cod(f)$. The third operation, a composition of arrows, takes an ordered pair of arrows, say (f, g), such that $cod(f) = d(g)$, and returns an arrow, denoted by $g \circ f$. The axioms are the following.

Axiom 5.1. $d(g \circ f) = d(f)$.

Axiom 5.2. $cod(g \circ f) = cod(g)$.

Axiom 5.3. The composition of arrows is associative.

Axiom 5.4. For every object, there is an arrow that has the object as domain and codomain, such that the arrow is a left and right identity for the composition of arrows.

Here the duality emerges when interchanging the operations for the domain and codomain, as well as reversing the order of the composition operation. This means, we write $cod(f)$ for $d(f)$, $d(f)$ for $cod(f)$, and $f \circ g$ for $g \circ f$. We can proceed as we did for the projective plane, and prove the dual of every axiom, although this time there is almost nothing to prove. Axiom 5.1 is the dual of Axiom 5.2, and vice versa. Axiom 5.3 can be rewritten as: For any three objects: f, g and h, that satisfy $cod(f) = d(g)$ and $cod(f) = d(h)$, we have: $(h \circ g) \circ f = h \circ (g \circ f)$. Which is the same as: For any three objects: f, g and h, that satisfy $c(f) = cod(g)$ and $d(f) = cod(h)$, we have: $f \circ (g \circ h) = (f \circ g) \circ h$. Likewise for Axiom 5.4.

6 Plane Graphs

A *simple polygonal curve* is a subset of the real plane that is the union of finitely many straight line segments which can be arranged in a sequence such that any pair of these straight line segments, with exception of the pair consisting of the first and the last segment, intersects if and only if the two are consecutive, and intersections only occur at endpoints of the straight line segments.

A *plane multigraph* consists of a finite set of points in the real plane, called *vertices*, and a finite set of simple polygonal curves, called *edges*, such that the endpoints of any edge are vertices, and no two edges intersect at a point that is not an endpoint of both, nor an edge intersects a vertex at a point that is not an endpoint. Consider the complement of the union of all edges and all vertices of a plane multigraph. It divides the plane in finitely many connected components. Those are called the *faces* of a plane multigraph. The sets of vertices and edges form an underling multigraph. A plane multigraph is said to be *connected* if the underlying multigraph is connected. Likewise, we adopt other terms like *path* or *cycle*. We call the set of all points of an edge that are not its endpoints, the *interior* of the edge.

Throughout let G be a connected plane multigraph, and V, E, and F its sets of vertices, edges, and faces respectively. A *plane dual* for G, is a plane multigraph G^* with vertices V^*, edges E^*, and faces F^*, such that there is a one-to-one correspondence between F and V^*, as well as a one-to-one correspondence between E and E^*, such that the following conditions are satisfied:

- Any vertex in V^* is contained in the face in F to which it corresponds.
- Any edge in E intersects its corresponding edge in E^* at exactly one point, this point is not an endpoint of a straight line segment of either edge, and the edge does not intersect any other edge of E^*.

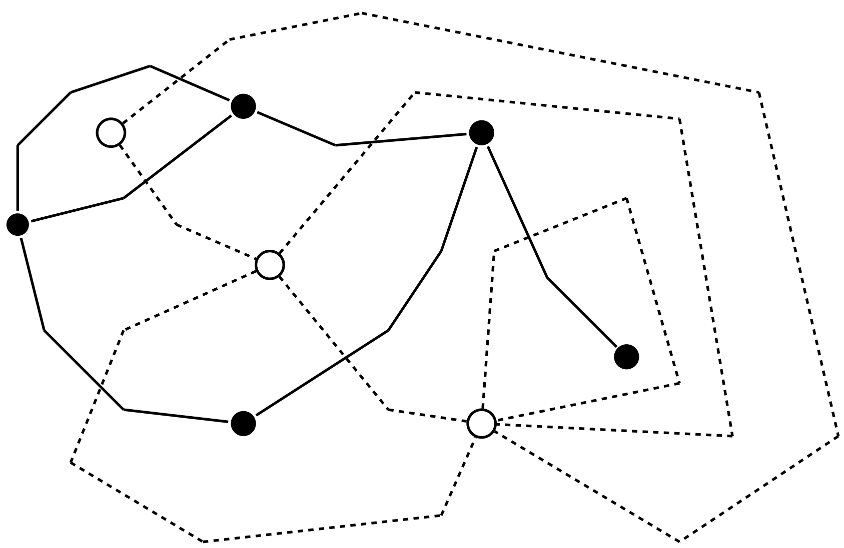

Figure 2: Example of a connected plane multigraph and a dual.

Any connected plane multigraph has a plane dual. To construct one, on each face f choose a point f^*. From the fact that each face is an open, connected subset of the plane, it follows that for each edge of G one can choose a corresponding edge, such that the second condition is satisfied. The endpoints are already determined.

A vertex and an edge are said to be *incident*, if the vertex is one of the endpoints of the edge. Two different edges are said to be *incident* if they share an endpoint. Choose a point on the interior of an edge. We can put an open ball of sufficiently small radius around that point, such that the intersection of the ball with the complement of the graph in the plane has two connected components, each of which is contained in one of the faces, possibly both in the same. These faces do not depend on the particular point chosen on the interior of the edge, and are said to be *incident* with the edge. The number of faces incident with a given edge is called its *degree*. An edge can have degree 1 or 2. The set of edges that are incident with a given face, is called the *frontier* of that face.

The plane with the vertices of G removed is connected, therefore we can join any two vertices of V^* with a curve that does not intersect any vertex of G. We can modify any such curve to intersect any edge of G at most once, and also, to enter a face at most once. Then, the edges that correspond to edges that the curve intersects is a path in G^* that joins the two vertices. Therefore G^* is connected.

Lemma 6.1. *Let v_0 be a vertex of G, let e_0 be an edge incident with v_0, and e_0^* the edge of G^* corresponding to e_0. If e_0 has two different endpoints, then there is a cycle in G^* that contains e_0^* such that every edge of the cycle corresponds to an edge incident with v_0.*

Proof. If e_0^* has only one endpoint, then it forms such a cycle. Assume therefore e_0^* has two different endpoints. Put a closed ball of sufficiently small radius around v_0 such that the ball only intersects the edges of G that are incident with v_0, these intersections are straight line segments, and the ball does not intersect any other vertex of G, nor any edge of G^*. Consider the straight line segments obtained by the intersection of the ball and the edges. Arrange them in a sequence $s_0, s_1, ..., s_n$ such that consecutive segments lie beside each other, in the sense that they can be joined by a path in the plane that stays inside the ball and does not intersect any other segment, and such that the sequence starts with a segment that belongs to e_0. Then s_0 and s_n also lie beside each other. As e_0 has two different endpoints by assumption, s_0 is the only segment in the sequence that belongs to e_0. The sequence has at least two elements: if s_0 was the only such segment, then the vertices of e_0^* could be joined by a path in the plane traveling alongside the edges e_0^* and e_0 that does not intersect G, which by the definition of a dual plane multigraph contradicts the assumption that the vertices are different. For $i \in \{0, ..., n\}$ let $\vartheta(s_i)$ be the edge of G to which the straight line segment s_i belongs, and $\vartheta(s_i)^*$ the corresponding edge in G^*.

For two different $i, j \in \{0, ..., n\}$, if s_i and s_j lie beside each other, then there is an endpoint v_1 of $\vartheta(s_i)^*$ and an endpoint v_2 of $\vartheta(s_j)^*$, such that both endpoints can be joined by a path in the plane that travels alongside the edges $\vartheta(s_i)^*$, $\vartheta(s_i)$, $\vartheta(s_j)$, and $\vartheta(s_j)^*$, which does not intersect any edge of G. Therefore v_1 and v_2 are vertices of G^* that lie on the same face of G, and thus we have $v_1 = v_2$. So $\vartheta(s_i)$ and $\vartheta(s_j)$ are incident.

In particular e_0^* is incident with $\vartheta(s_1)^*$ and $\vartheta(s_n)^*$. Further if e_0^* and $\vartheta(s_1)^*$ share a given vertex of e_0^*, then e_0^* and $\vartheta(s_n)^*$ share the other vertex of e_0^*. Consider the edges $\vartheta(s_1)^*$,...,$\vartheta(s_n)^*$. By the previous paragraph consecutive edges are incident, therefore the union of those edges forms a connected subgraph of G^*. It contains both endpoints of e_0^*, but it does not contain e_0^* itself. Take a path in this subgraph that connects both endpoints, adding e_0^* to the path yields a cycle with the required properties. □

Proposition 6.2. *The plane multigraph G is a dual for G^*.*

Proof. The second condition in the definition of a dual plane multigraph is self dual. It is therefore left to show that every face of G^* contains exactly one vertex of G.

Assume there are two vertices v_1 and v_2 of G that lie on same face of G^*. As G is connected, there is a path from v_1 to v_2 in G. Consider the first edge of the path starting from v_1. Consider its corresponding edge in G^*. By Lemma 6.1, there is a cycle in G^* that contains this edge, such that every edge in the cycle corresponds to an edge that is incident with v_1. This cycle divides the plane into two connected components. The vertices v_1 and v_2 must be in the same of those connected components. With the first edge the path exits the component in which v_1 is contained. In order to reach v_2 it must enter again. But any edge that crosses the cycle to enter this component, arrives at v_1. We therefore have a contradiction and thus a face of G^* contains at most one vertex of G.

Also, every face of G^* that is not the whole plane, must have an edge of G^* on its boundary. The corresponding edge of G to that edge has a vertex on the face. □

Let S be a plane multigraph that has the same vertices as G and its edges are a subset of the edges of G. Let S' be the plane multigraph that has the same vertices as G^* and its edges are those edges of G^* that correspond to edges of G that are not edges of S.

Lemma 6.3. *The plane multigraph S is connected if and only if S' has no cycles.*

Proof. By the Jordan Curve Theorem, a cycle in S' divides the plane into two connected components. An edge of G that corresponds to an edge of the cycle, has one endpoint in each connected component. Therefore they cannot be joined by a path that does not meet the given cycle. Thus if S' has a cycle, then S is not connected.

Conversely, assuming that S is not connected, we construct a cycle in S'. For the purpose of this proof we call an edge of G^* that corresponds to an edge of G which connects two vertices that are in different connected components of S, a *separating* edge. As S is not connected, there is at least one separating edge, say e_0^*. If the endpoints of e_0^* are equal, this edge is already a cycle in S'. Assume therefore the endpoints of e_0^* are distinct. We show that for each endpoint of e_0^* there is another separating edge that is incident with the endpoint. Given this, a cycle can be found by iteration. Let v_0 be any endpoint of e_0^*. Let e_0 be the edge of G that corresponds to e_0^*. By the dual statement of Lemma 6.1, there is a cycle in G that contains e_0 such that every edge in the cycle corresponds to an edge that is incident with v_0. Remove e_0 from the cycle. We obtain a path in G that connects the endpoints of e_0. Consider the sequence of vertices of this path. The first and the last vertices are the endpoints of e_0, and these belong to different connected components of S. Hence, there must be a pair of consecutive vertices in the sequence that are in different connected components of S. The edge in the path joining this pair of vertices corresponds to a separating edge that is incident with v_0. □

If S is connected and has no cycles it is called a *spanning tree* of G. By Lemma 6.3 and by duality, S is a spanning tree for G if and only if S' is a spanning tree of G^*.

A consequence of this, is that a connected plane multigraph satisfies Euler's Formula. Let T be a spanning tree for G. As the number of edges in a tree is always one less than its number of vertices we have:

$$[\text{ Number of edges of } T\ ] + 1 = |V|,$$

$$[\text{ Number of edges of } T'\ ] + 1 = |F|.$$

The number of edges in T and T' together, sum to $|E|$. Thus, putting both equations together we obtain:

$$|V| - |E| + |F| = 2.$$

Under the definition for plane multigraphs, two different plane multigraphs might still be such that we would consider them to be essentially the same. For example, when one is a translation of the other, or obtainable after certain deformations. One can introduce many different equivalence relations that capture this idea. But is there an equivalence relation that is compatible with duality? More precisely, we want an equivalence relation such that assigning a dual to an equivalence class by considering any representative, any dual of the representative, and finally the equivalence class of that dual, yields a well defined map. A first try could be to consider only the underlying multigraph structure, whithout the embedding into the plane, and therefore define two plane multigraphs to be equivalent if there is a one-to-one correspondence between their vertices and a one-to-one correspondence between their edges such that the incidence relations are conserved. However, Figure 3 shows that two plane multigraphs can be equivalent under this relation but have duals that are not. In a second attempt, we could ask for these correspondences to be induced by a homeomorphism of the plane to itself, so that the homeomorphism would map every vertex to its corresponding vertex, and every edge to its corresponding edge. This is not adequate either. Figure 4 shows that a graph can have two duals that are not equivalent under this relation. It turns out, that this would be actually a precise equivalence relation for our purpose if instead of drawing the graphs on the plane, we would do so on a sphere! We only need the topological properties of a sphere though. So we can keep our definitions for plane multigraphs on the plane, only adding a point at infinity.

We give $\mathbb{R}^2 \cup \{\infty\}$ the topology of the Riemann sphere, this means that a subset that does not contain infinity is open if it is open in $\mathbb{R}$, and a subset that contains infinity is open if its complement is compact in $\mathbb{R}$. We define two plane multigraphs to be *equivalent* if there is a homeomorphism from $\mathbb{R}^2 \cup \{\infty\}$ to itself that maps vertices to vertices and edges to edges.

Proposition 6.4. *Let G_1 and G_2 be connected plane multigraphs. Then G_1^* is equivalent to G_2^* if and only if G_1 is equivalent to G_2.*

We sketch a proof for the proposition. By duality it suffices to prove only one direction. Say ϕ is a homeomorphism from $\mathbb{R}^2 \cup \{\infty\}$ to itself that

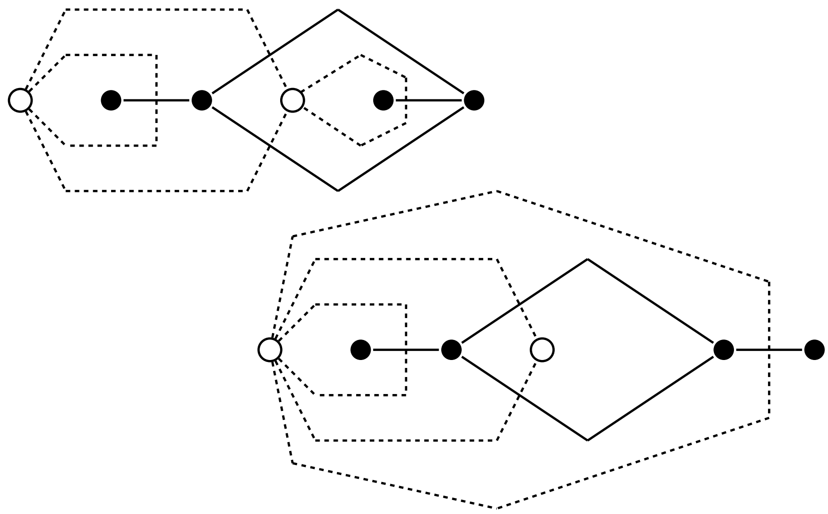

Figure 3: Counting the degree of the vertices of the dashed graphs, reveals that their elements cannot paired such that the incidence relations are conserved.

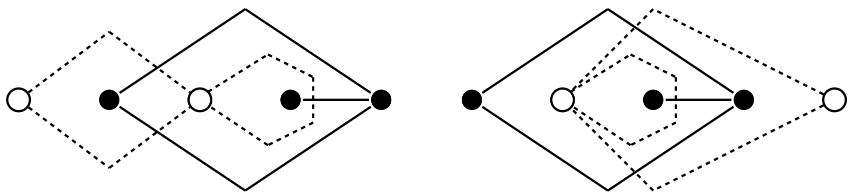

Figure 4: A homeomorphism from the plane to itself that maps the vertices of one of the dashed graph to the vertices to the other dashed graph, as well as their edges, would map a set that it compact to a set that is not.

induces an equivalence between G_1 and G_2. We construct a homeomorphism ϕ^* from $\mathbb{R}^2 \cup \{\infty\}$ to itself that induces an equivalence between G_1^* and G_2^*. We do so by defining it first on the vertices of G_1^*, then extending it on the edges, and finally to the faces. Because ϕ is a homeomorphism it maps faces of G_1 to faces of G_2. Consider the following diagram:

$$
\begin{tikzcd}
V_1 \arrow[d, dash] \arrow[rrrrr, dashed, dash, bend left] & E_1 \arrow[d, dash] \arrow[rrr, dashed, dash, bend left] & F_1 \arrow[d, dash] \arrow[r, dashed, dash, bend left] & V_1^* \arrow[d, dash] & E_1^* \arrow[d, dash] & F_1^* \arrow[d, dash] \\
V_2 \arrow[rrrrr, dashed, dash, bend right] & E_2 \arrow[rrr, dashed, dash, bend right] & F_2 \arrow[r, dashed, dash, bend right] & V_2^* & E_2^* & F_2^*
\end{tikzcd}
$$

The dashed lines represent one-to-one correspondences given by the dualities, while the non-dashed lines represent the one-to-one correspondences given by the equivalences. The correspondences between the elements of the duals are not defined yet, since they are determined by ϕ^*. First we define ϕ^* on V_1^* such that the diagram commutes. Again by the fact that ϕ is a homeomorphism, it follows that if we pair the edges of G_1^* with the edges of G_2^* such that the digram commutes, corresponding edges will have corresponding endpoints. We therefore can extend ϕ^* to the edges on E_1^* such that it maps edges to edges and induces the correspondence for which the diagram commutes. We proceed in the same fashion to extend ϕ^* to the faces. The key is that by pairing the faces such that the digram commutes, corresponding faces have corresponding frontiers. Also, one uses the fact that every face is homeomorphic to the open unit disk.

7 Percolation Theory

In Percolation Theory, a similar duality to that presented in the previous section plays a prominent role. Consider $\mathbb{Z}^2 \subset \mathbb{R}^2$. Consider the infinite graph, where the vertices are the points in $\mathbb{Z}^2$ and there is an edge between any two points that are separated by an Euclidean distance of 1. A function from the set of edges to $\{0, 1\}$ is called a *configuration*. We define a probability space on the set of all configurations. In a given configuration, an edge is called *open* if it is mapped to 1, and *closed* otherwise. For an edge e, define A_e to be subset of all configurations for which e is open. Our σ-algebra will be the one generated by all such subsets. Take an arbitrary $p \in [0, 1]$. We call it the *percolation parameter*. There is a unique probability measure $\mathbb{P}_p$ on our σ-algebra, such that given any edge e, the probability $\mathbb{P}_p(A_e)$ equals

p, and the set of all events of the form A_e is independent. This will be our probability measure.

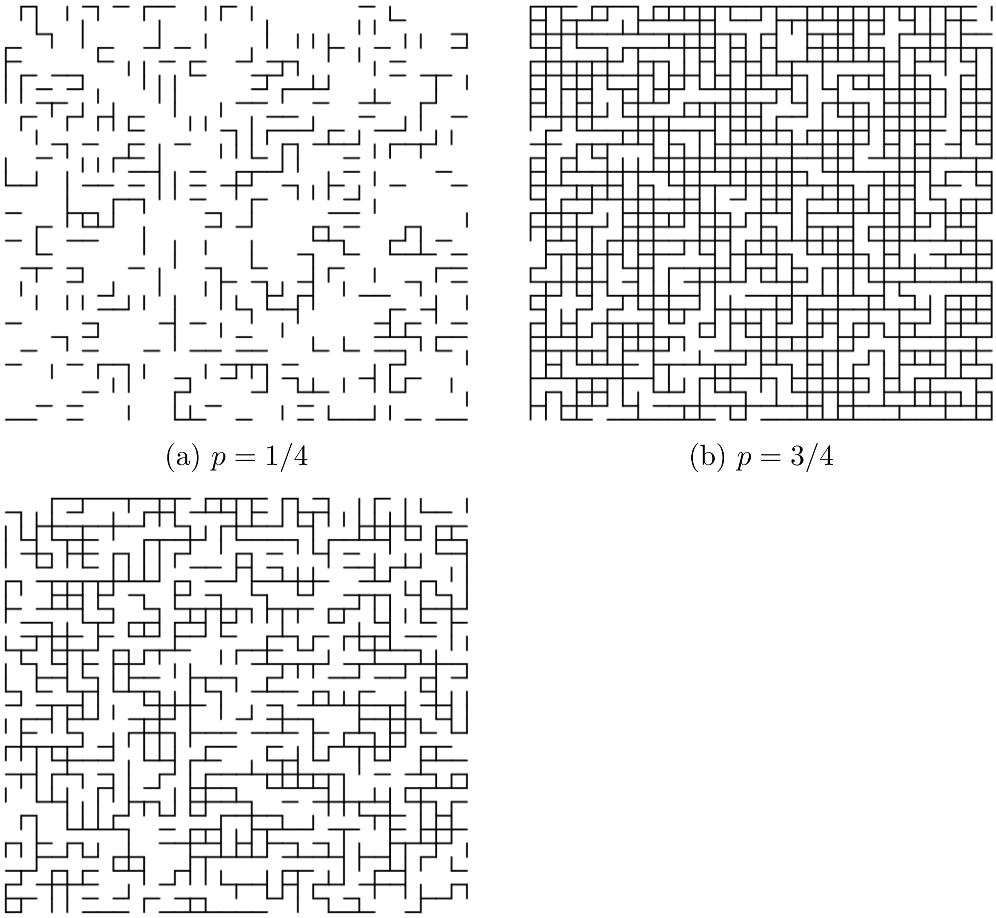

(a) $p = 1/4$

(b) $p = 3/4$

(c) $p = 1/2$

Figure 5: Simulations of for distinct values of p on an excerpt of the infinite graph. Only the open edges are shown.

A *path* is a finite sequence of distinct points in $\mathbb{Z}^2$, such that consecutive points have Euclidean distance 1. Consider a path of strictly more than two points, for which the first point has Euclidean distance of 1 to the last point. For such a path, consider the union of the set of edges that join consecutive points with the edge that joins the first point with the last. Such a set of edges is called a *cycle*. Given a configuration, a path for which consecutive points are joined by open edges is said to be *open*, or *closed* if those edges are closed. A cycle that consists of open edges is likewise said to be *open*, or *closed* if it consists of closed edges. An *open cluster* is a subset of $\mathbb{Z}^2$, such that any two points can be joined by an open path within the subset, and is maximal in the sense that there is no subset strictly containing it that also satisfies this condition.

The name Percolation Theory has the following explanation. Our model is a two-dimensional analog of what could be seen as a model for some porous material, the open edges being conduits where a liquid can pass through. An interesting question is how connected these conduits are, in dependance of p. Suppose this is a rock. If the conduits were only sparely connected, water would only penetrate superficially. On the contrary, if they were largely connected, forming large tunnels that go trough the rock, water would strain through.

A formal question we can ask in this regard is: Given a point in $\mathbb{Z}^2$, what is the probability that the open cluster containing this point is infinite? First, that probability is independent of the given point, so we can assume it to be the origin. Let C be the event: "The open cluster containing the origin is infinite". Our probability is then $\mathbb{P}_p(C)$. If $p = 0$, this probability is 0. If $p = 1$, this probability is 1. But what exactly happens in the middle? Consider yet another analogous model, for which we take $\mathbb{Z}$ instead of $\mathbb{Z}^2$. This case is easier. Here, if $p < 1$, the probability for having only open edges on at least one of both sides of the origin is 0, as is any probability for an event that fixes the state of an infinite number of edges. Therefore, unless p is 1, the probability for having the open cluster containing the origin to be infinite, is also 0. Is the same true for our model in $\mathbb{Z}^2$? The answer is no, as we will prove using duality.

Claim 7.1. There is $p < 1$ such that $\mathbb{P}_p(C) > 0$.

Recall how we defined the dual graph in the previous section. We proceed similarly here. Consider our infinite graph with vertices $\mathbb{Z}^2$ and the edges as previously defined. Let it be denoted by $\mathbb{L}$. Consider the infinite graph $\mathbb{L}^*$ that has vertices the points of the form $(1/2 + a, 1/2 + b)$, where $a, b \in \mathbb{Z}$, and edges the straight line segments that connect points of distance 1. Each edge of $\mathbb{L}$ intersects with exactly one edge of $\mathbb{L}^*$ and vice versa. We therefore have a correspondence between the edges of $\mathbb{L}$ and those of $\mathbb{L}^*$. For a configuration of $\mathbb{L}$, the *dual configuration* is the configuration of $\mathbb{L}^*$ where open edges correspond to open edges. A cycle in $\mathbb{L}^*$ divides the plane into two connected components, and only one is bounded. We say that a cycle in $\mathbb{L}^*$ *contains the origin*, if the origin is contained in the bounded connected component. Using a similar argument as in the proof of Lemma 6.3, one obtains the following lemma. Compare with Figure 6.

Lemma 7.2. *Given a configuration of $\mathbb{L}$, the open cluster containing the origin is finite if and only if there is a closed cycle in the dual configuration that contains the origin.*

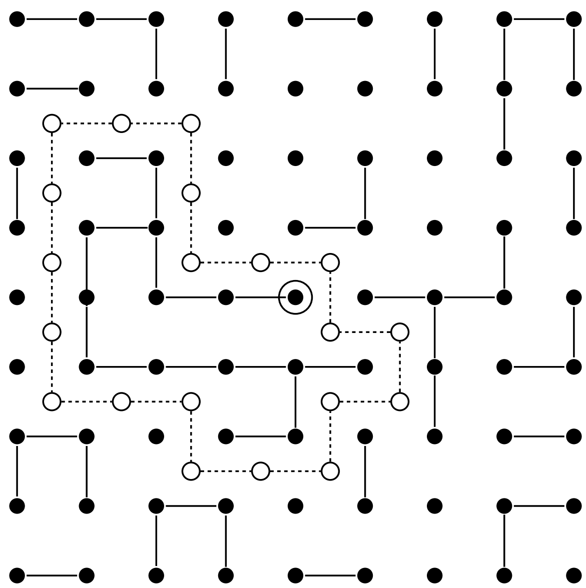

Figure 6: A depiction of the excerpt of a configuration around the origin. Only the open edges are shown. The cluster containing the origin is finite. The dotted edges form a closed cycle in the dual configuration, which contains the origin.

We are now ready to prove Claim 7.1. Let C' be the event: "There is a closed cycle in the dual configuration that contains the origin". By Lemma 7.2 we have:

$$\mathbb{P}_p(C) = 1 - \mathbb{P}_p(C').$$

We look for an estimate for $\mathbb{P}_p(C')$. Consider the positive part of the real line. Any cycle in $\mathbb{L}^*$ that contains the origin, must intersect this set at least once. This can only happen at the points of the form $(0, -1/2 + a)$, for $a \in \mathbb{N}$. A cycle of length n that contains the origin must intersect the positive part of the real line somewhere before n. Also, for a given point of the form $(0, -1/2 + a)$, with $a \in \mathbb{N}$, there are not more than 3^n cycles of length n that intersect that point. To see this, imagine constructing a cycle edge by edge, starting with the edge that intersects the given point on the positive real line, then placing the second edge incident on the upper endpoint of the first edge, and continuing to place each new edge such that it is incident with the previous one. Each time we add a new edge there are not more three possibilities for placing it. Therefore $n3^n$ is an upper bound for the numbers of cycles of length n that contain the origin. The probability for a given cycle of length n to be closed is $(1-p)^n$. We therefore have:

$$\mathbb{P}_p(C') \leq \sum_{n=1}^{\infty}(1-p)^n n3^n.$$

The right hand side converges for $p > 2/3$, and approaches 0 when p approaches 1. From this the claim follows. This is just a glimpse of how duality is used in Percolation Theory.

Much more can be proved about the behavior of $\mathbb{P}_p(C)$ in dependance of p. As a function in p, we have that $\mathbb{P}_p(C)$ is monotonically non-decreasing, and it turns out that there is $p_c \in [0,1]$, called the *critical value*, such that:

$$\mathbb{P}_p(C) \begin{cases} = 0, & \text{for } p < p_c. \\ > 0, & \text{for } p > p_c. \end{cases}$$

Claim 7.1 can be used as a key ingredient to prove this. By the use of duality one can show that in fact $p_c = 1/2$. This requires substantially more work than the proof of Claim 7.1. See Grimmett [14], Chapter 1 and Chapter 11, for a treatment of these results.

References

[1] H. S. M. Coxeter, *Projective Geometry*, Second Edition, Springer, 2003.

[2] A. N. Whitehead, *The Axioms of Projective Geometry*, Cambridge University Press, 1913.

[3] S. Lang, *Algebra*, Springer, 2002.

[4] G. Fischer, *Lineare Algebra*, Springer, 2008.

[5] N. Bourbaki, *General Topology*, Volume 2, Hermann, Paris, 1966.

[6] B. von Querenburg, *Mengentheoretische Topologie*, Springer, 2001.

[7] J. Munkres, *Topology*, Second Edition, Pearson, 2000.

[8] W. Rudin, *Fourier Analysis on Groups*, Intersience Publishers, New York, 1962.

[9] A. Deitmar and S. Echterhoff, *Principles of Harmonic Analysis*, Springer, 2014.

[10] S. A. Morris, *Pontryagin Duality and the Structure of Locally Compact Abelian Groups*, Cambridge University Press, 1977.

[11] L. Pontrjagin, The Theory of Topological Commutative Groups, *The Annals of Mathematics*, Volume 35, Number 2, 1934.

[12] S. Mac Lane, *Categories for the Working Mathematician*, Springer, 1998.

[13] R. Diestel, *Graph Theory*, Springer, 2005.

[14] G. Grimmett, *Perlocation*, Springer, 1999.