Closed Riemann Surfaces and Metric Graphs
0. Introduction.

A **Riemann surface** S is a one-dimensional complex manifold, i.e. S is a (Hausdorff second countable topological) space with a holomorphic system of one-complex variable local coordinates. Holomorphic means that a function $f : U \to \mathbb{C}$ defined on an open subset of S is holomorphic as a function of one local coordinate if and only if it is holomorphic as a function of the others. Equivalently, if z and w are local coordinates then $z = h(w)$ is holomorphic as a function of w whenever both are defined. The topology of a closed (compact) Riemann surface is either the sphere or a g-holed torus.

There is a single system of holomorphic local coordinates on a sphere, but there are *moduli* (families) of Riemann surfaces with the topology of each g-holed torus. When $g = 1$, these are the elliptic curves, whose moduli have one complex dimension; they may be seen as points of a fundamental domain in the upper half plane. When $g > 1$, the moduli of Riemann surfaces has $3g - 3$ complex dimensions.

By the maximum principle, the only holomorphic functions on a closed Riemann surface S are the constants, but the meromorphic functions on S determine the Riemann surface S in a precise, algebraic manner. This ties the study of Riemann surfaces to algebra, number theory and group theory. A meromorphic function is a "partially defined" function:

$$\phi : S - - > \mathbb{C}$$

that is well-defined and holomorphic away from finitely many points of S, at which ϕ has poles as isolated singularities. It is a very reasonable question, answered remarkably precisely by the Riemann-Roch theorem, whether a given Riemann surface has **any** non-constant meromorphic functions!

Lemma 0.1. If ϕ is a meromorphic function on a closed Riemann surface S, the total number of zeroes (with multiplicity) is equal to the total number of poles (with multiplicity).

Proof. The residue of the meromorphic *differential*

$$\frac{1}{2\pi i} \frac{d\phi}{\phi}$$

at each point of S records the multiplicity of zero at the point (positively) and the multiplicity of pole (negatively) and the residues sum to zero when the Riemann surface is closed.

A general meromorphic *differential* ω is an object of the form $\omega = \psi(z)dz$ in each local coordinate z. But unlike meromorphic functions, a differential "transforms" into a nearby local coordinate by the rule:

$$\psi(z)dz = \psi(h(w))h'(w)dw$$

designed so that if $\phi(z)$ is a (meromorphic) function, then:

$$d\phi =: \phi'(z)dz = \phi'(h(w))h'(w)dw = \phi'(w)dw$$

defines a meromorphic differential, by the chain rule.

Even though the coefficients $\psi(z)$ of the meromorphic differential ω do not determine a well-defined function on S, the multiplicities of their zeroes and poles are well-defined since $h'(w) \neq 0$ by virtue of the fact that h has a holomorphic inverse. On the other hand, if $\omega_1 = \psi_1(z)dz$ and $\omega_2 = \psi_2(z)dz$ are meromorphic differentials, then the *ratio* of their coefficients:

$$\phi(z) = \psi_1(z)/\psi_2(z) = \psi(_1h(w))/\psi_2(h(w)) = \phi(h(w))$$

does patch to a meromorphic function. Conversely, if ϕ is a meromorphic function, then $\phi\omega$ is a meromorphic differential.

Example. The **Riemann Sphere** is the topological space $\mathbb{CP}^1 := \mathbb{C} \cup \{\infty\}$ (the one-point compactification of $\mathbb{C} = \mathbb{R}^2$) covered by $U = \mathbb{CP}^1 - \{\infty\}$ with local coordinate z and $V = \mathbb{CP}^1 - \{0\}$ with local coordinate w and $z = h(w)$ for $h(w) = 1/w$ on $\mathbb{C} - \{0\} \subset \mathbb{C}$. Thus, for example,

$$z - 2 = \frac{1}{w} - 2 = \frac{1 - 2w}{w}$$

in each of the coordinates, giving a meromorphic function with a zero at $z = 2$ (equivalently at $w = \frac{1}{2}$) and a pole at $w = 0$ (the point at infinity).

The meromorphic differential:

$$\omega = dz = -\frac{1}{w^2}dw$$

has a pole of order 2 at $w = 0$ and no zeroes. The differential:

$$\frac{1}{2\pi i}\frac{dz}{z} = -\frac{1}{2\pi i}\frac{dw}{w}$$

as in Lemma 0.1 has poles at $z = 0$ and $w = 0$ whose residues (1 and -1) reflect the fact that z has a simple zero at $z = 0$ and a simple pole at $w = 0$.

Let S be a closed Riemann Surface of genus g. Then the *holomorphic* differentials (meromorphic differentials with no poles) are a vector space of dimension g. This is in contrast with holomorphic functions on S, which are always a vector space of dimension 1 (the constants). In the example of the sphere above, non-zero meromorphic differentials always have at least two poles, and so zero is the only holomorphic differential.

For a set X, let:

$$X^d = X \times \cdots \times X = \{(x_1, ..., x_d) \mid x_i \in X\}$$

be the set of d-tuples of points X, and let: $X_d = \{x_1 + \cdots + x_d \mid x_i \in X\}$ be the set of *unordered* d-tuples of points of X. This can be thought of as a subset of the *free abelian group* on the points of X:

$$X_d \subset \mathbb{Z}[X] = \left\{ \sum_{i=1}^{n} m_i x_i \mid m_i \in \mathbb{Z} \right\}$$

There is a **degree map**

$$\deg : \mathbb{Z}[X] \to \mathbb{Z}; \ \deg \left(\sum_{i=1}^{n} m_i x_i \right) = \sum_{i=1}^{n} m_i$$

and we will let $\mathbb{Z}[X]_0 = \ker(\deg)$.

When $X = S$ is a Riemann surface, we call elements of $\mathbb{Z}[S]$ **divisors**, and elements of S_d are **effective divisors** of degree d. While the object $\mathbb{Z}[S]$ is very large and formal, the parameter spaces S_d for effective divisors of degree d are complex manifolds of dimension d (see §5).

Definition 0.2. The divisor $div(\phi)$ of a meromorphic function ϕ on S is:

$$div(\phi) = \sum_{p \in S} \mathrm{ord}_p(\phi) p \in \mathbb{Z}[S]_0$$

where $\mathrm{ord}_p(\phi)$ is the *order of vanishing* of ϕ at $p \in S$, which is non-zero only at finitely many points of S.

Note that the meromorphic functions on S are a **field**, denoted $\mathbb{C}(S)$ and the div function is a group homomorphism:

$$div : \mathbb{C}(S)^* \to \mathbb{Z}[S]_0$$

converting multiplication (of functions) to addition (of divisors).

Similarly, each meromorphic differential has an associated divisor:

$$\mathrm{div}(\omega) = \sum_{p \in S} \mathrm{ord}_p(\psi) \cdot p \in \mathbb{Z}[S]$$

with a well-defined degree (independent of ω). We will see that this degree is $2g - 2$.

Soundbite. Meromorphic functions on S satisfy *polynomial* relations.

We will explore this in detail, and see analogous behavior in graphs.

A (finite, combinatorial) **graph** Γ consists of the following data:

$$V \text{ a finite set of } \textit{vertices} \text{ of } \Gamma$$

$$E \text{ a finite set of } \textit{edges} \text{ of } \Gamma$$

$$\epsilon : E \to V_2, \text{ the ends of each edge}$$

which has a very coarse topology in which the closed sets are finite sets of vertices as well as finite sets of edges with all their ends. Note that in this topology, an individual edge $e \in E$ is an **open** set and that the smallest open neighborhood of a vertex $v \in \Gamma$ is the vertex v together with all "adjacent" edges e for which v is an end of e (the "star" of v).

Definition 0.3. The *valence* $val(v)$ of $v \in V$ is given by the equality:

$$\sum_e \epsilon(e) = \sum_v \mathrm{val}(v) \cdot v \in \mathbb{Z}[V]$$

i.e. $val(v)$ is the number of times v is an end of an edge adjacent to v.

Definition 0.4. A *path* in a graph Γ is a series of vertices and edges:

$$\{v_1, e_1, v_2, e_2, ..., e_n, v_{n+1}\}$$

with the property that $\epsilon(e_i) = v_i + v_{i+1}$. A path of the form vev is a *loop*. In general, a *circuit* is a path with $v_1 = v_{n+1}$ and no other repeated vertices.

Definition 0.5. The **genus** of a connected graph Γ is given by:

$$|V| - |E| = 1 - g \text{ (Euler's Formula)}$$

Lemma 0.6. $g(\Gamma) = 0$ if and only if Γ has no (non-trivial) circuits.

Proof. Suppose $\Lambda = \{v_1, e_1, ..., e_n v_1\}$ is a circuit in a graph Γ.

This is a genus one *subgraph* of Γ. The graph Γ is built by adding edges to Λ, with each added edge coming with one or zero new vertices (depending on whether the other end of the edge attaches to the subgraph). Thus the genus only *increases* as edges are added. So graphs with circuits have positive genus. On the other hand, suppose Γ is circuit-free and let $v \in V(\Gamma)$. Then either Γ is a singleton vertex (which has genus zero), or else Γ has an edge e adjacent to v. This edge is not a loop, so contracting the path vew gives a new graph Γ' with vew replaced by a single vertex v which is **also** circuit free and of the same genus as Γ. Repeating such contractions eventually contract all the edges of Γ to the singleton v, so $g(\Gamma) = 0$. $\square$

Graphs with no circuits are called *trees.*

Remark. If we contract non-loops adjacent to v as in the proof above on a graph of genus g until v is only adjacent to loops, the result is a **bouquet** of g loops with the single vertex v, and the subgraph of Γ that is contracted is a **spanning tree** of Γ, i.e. a connected tree inside Γ that has the same vertex set as Γ.

Chip Firing from v is a zero-sum game (later adapted to metric graphs) in which a divisor $D = \sum_v d_v \cdot v \in \mathbb{Z}[V]$ is transformed by reducing d_v by the valence of v and increasing d_w by one for all the neighbors of v. Given a divisor of positive degree but with some vertices "in debt" (with $d_v < 0$), the game is to transform D into an effective divisor $E \in V_d$ by firing from a sequence of vertices. Playing this game is analogous on a Riemann surface to adding $div(\phi)$ to a divisor D of degree d, to try to similarly create an effective divisor. In both contexts, one has:

Jacobi Inversion. The game is winnable for all divisors D of degree $\geq g$.

This is the first of many remarkable coincidences between graphs and Riemann surfaces that we will explore in these notes.