

The Search for a Global Primitive

Čech Cohomology with Coefficients in a Sheaf

Michael Weiss

This note tries to motivate the definition of Čech cohomology with coefficients in a sheaf. We will begin by looking at the first deRham cohomology group of a surface: this is a very intuitive mathematical object, with a transparent topological significance. We will see how Čech cohomology emerges from a natural attempt to calculate deRham cohomology. Then we will extract a more general framework from our constructions, and see how this approach finds other applications.

The purpose of this note is to foster intuition, not to provide a fully rigorous or general treatment.

1 deRham cohomology

Let X be a connected 2-dimensional differential manifold. Recall that a 1-form ω on X is **closed** if $d\omega = 0$, and is **exact** if $\omega = df$ for some differentiable function f . Also, f is called a **primitive** of df . Since $ddf = 0$, $\text{Exact}^1(X) \subseteq \text{Closed}^1(X)$, where of course $\text{Exact}^1(X)$ is the set of exact 1-forms and $\text{Closed}^1(X)$ is the set of closed 1-forms. Both $\text{Exact}^1(X)$ and

$\text{Closed}^1(X)$ are vector spaces over $\mathbb{R}$, so we can form the quotient:

$$\text{Rh}^1(X) = \text{Closed}^1(X)/\text{Exact}^1(X)$$

known as the first deRham cohomology group¹.

Stokes's theorem tells us that for any region R and any 1-form ω , $\int_{\partial R} \omega = \int_R d\omega$. In particular, if a loop l in X is contractible to a point and if $d\omega = 0$, then $\int_l \omega = 0$. So if $a, b \in X$ and p_1 and p_2 are homotopic paths from a to b , then $\int_{p_1} \omega = \int_{p_2} \omega$ for any closed 1-form ω .

So in particular, if X is simply connected and ω is closed and $a, b \in X$, then $\int_a^b \omega$ is well-defined, independent of path. If we fix a and let b vary, we have a function $f(b) = \int_a^b \omega$, and $df = \omega$. Conclusion: if X is simply connected, then $\text{Rh}^1(X)$ is trivial.

More generally, we see that a 1-form ω is exact precisely when its integrals over all loops are zero. For in that case, $\int_a^b \omega$ is well-defined, independent of the path from a to b , and $f(b) = \int_a^b \omega$ has $df = \omega$.

For a simple example of a non-trivial deRham group, we look at $X = \mathbb{R}^2 - \{0\}$. Points in X have polar coordinates (r, θ) ; although θ is well-defined only modulo 2π , $d\theta$ is a well-defined 1-form, given explicitly in (x, y) coordinates by $d\theta = (x dy - y dx)/(x^2 + y^2)$ (although we won't need this formula for anything). Now, $d\theta$ is not exact, as we see by integrating it over a loop around the origin: $\int_l d\theta = 2\pi$, but the integral of an exact 1-form over a loop is always zero.

In fact, if ω is any closed 1-form, then $\omega = c d\theta + df$ for some constant c and some function f . It is easy to find c and f : integrate the desired equation over a loop l around the origin, getting $\int_l \omega = 2\pi c$. With c thus determined, we have $\int_l (\omega - c d\theta) = 0$. In fact, $\omega - c d\theta$ has a zero integral along *any* loop in X (since l generates the fundamental group of X), so $\omega - c d\theta$ is exact.

¹ Really a vector space, as we've just seen, but for some reason it is traditional in algebraic topology to refer to homology and cohomology vector spaces as *groups*.

So $\text{Rh}^1(\mathbb{R}^2 - \{0\})$ is one-dimensional in this case, with $d\theta$ serving as a basis. Note that the fundamental group of $\mathbb{R}^2 - \{0\}$ is $\mathbb{Z}$, generated by a loop around the origin.

As the simplest compact example, we look at the torus $T = S^1 \times S^1$. Points on T have coordinates (φ, θ) , with φ and θ being well-defined only modulo constants, but again giving rise to well-defined 1-forms $d\varphi$ and $d\theta$. The fundamental group of T is $\mathbb{Z} \times \mathbb{Z}$, generated by an “equatorial” loop and a “meridional” loop. Integrating over these loops, we quickly find that $d\varphi$ and $d\theta$ are not exact, and that an arbitrary closed 1-form ω on T is uniquely expressible as $c_\varphi d\varphi + c_\theta d\theta + df$. So $\text{Rh}^1(T)$ is two-dimensional.

We see that $\text{Rh}^1(X)$ is in some sense “dual” to the fundamental group of X . 1-forms without primitives have their lack of exactness revealed by having non-zero integrals over loops. Actually, we’re better off working with the first homology group instead of the fundamental group. The homology group, like the deRham group, is always abelian; the fundamental group may not be. If we use real coefficients, then $H_1(X, \mathbb{R})$, like $\text{Rh}^1(X)$, is a vector space over $\mathbb{R}$.

A 1-cycle l can be regarded as an $\mathbb{R}$ -linear combination of loops, and we can integrate any 1-form over a 1-cycle. Let us define, for any 1-cycle l and closed 1-form ω

$$\langle \omega, l \rangle = \int_l \omega$$

i.e., a pairing between closed 1-forms and 1-cycles. Stokes’s theorem tells us that $\langle \omega, \partial R \rangle = 0$ whenever ω is a closed 1-form and ∂R is a 1-boundary. It is obvious that $\langle df, l \rangle = 0$ for all f and all 1-cycles l . Therefore:

$$\langle \omega + df, l + \partial R \rangle = \langle \omega, l \rangle$$

i.e., the pairing is well-defined on $\text{Rh}^1(X) \times H_1(X, \mathbb{R})$.

The 2-dimensional version of the deRham theorem says that if X is a paracompact surface, then this pairing establishes $\text{Rh}^1(X)$ as the dual space

to $H_1(X, \mathbb{R})$. The deRham theorem also generalizes to n dimensions. I won't pursue this further though, as we are now ready to turn to the Čech construction.

2 Čech cohomology

In the two examples of the previous section, we could easily find closed 1-forms that were not exact. Already for a surface of genus 2 (a “sphere with two handles”), it is not so clear how to proceed.

In one sense, $\text{Rh}^1(X)$ classifies the obstructions to solving equations of the form $df = \omega$, where we regard ω as given and f as unknown. Notice that the obstructions are always global in nature: if U is a simply connected open subset of X and ω is a closed 1-form on U , then $\omega = df$ for some f defined on U . Also, f is uniquely determined up to a constant².

The Čech idea is to try to patch together these local solutions to give a global solution, and see what, if anything, gets in the way.

Let $\mathfrak{U}$ be a covering of X with open sets. Assume all the open sets in $\mathfrak{U}$ are simply connected, and assume the same for the intersection of any finite sub-collection of $\mathfrak{U}$. Intuitively speaking, picture a tiling of X with convex polygons; now enlarge each polygon just a little, causing overlaps, and let $\mathfrak{U}$ be the set of interiors of these polygons. We can always choose the tiling so that no more than three polygons meet at any point. In keeping with this choice, we will assume that the intersection of any four distinct elements of $\mathfrak{U}$ is empty.

Let ω be a closed 1-form on X , fixed for the time being. For any $U \in \mathfrak{U}$, let F_U be the set of solutions to $df = \omega$ on U . Note that if $f, f' \in F_U$, then $f - f'$ is constant.

² We follow the convention that the term “simply connected” implies connected.

Let's say we pick an $f_U \in F_U$ for each $U \in \mathfrak{U}$. The f_U 's patch together to form a global solution if and only if for all $U, V \in \mathfrak{U}$, we have $f_U = f_V$ on the intersection $U \cap V$.

Define $f_{UV} = f_U - f_V$; of course, f_{UV} is defined only on $U \cap V$. Since $df_U = \omega = df_V$, we see that f_{UV} is constant. So the question is, can we choose the f_U 's so that all the f_{UV} are 0?

Choose a constant c_U for each $U \in \mathfrak{U}$. Then $f'_U = f_U + c_U$ is another set of local solutions, and every set of local solutions is of this form, for some set of c_U 's. Also:

$$f'_{UV} = f'_U - f'_V = (f_U - f_V) + (c_U - c_V) = f_{UV} + c_{UV}$$

Some definitions will bring the import of this equation into sharper focus.

For $q = 0, 1, 2$, let $N^q(\mathfrak{U})$ be the set of ordered $(q+1)$ -tuples $(U_0, \dots, U_q) \in \mathfrak{U}^{q+1}$ with non-empty intersection, i.e., $U_0 \cap \dots \cap U_q \neq \emptyset$. We stop at $q = 2$ because of our assumption that the intersection of any four distinct elements of $\mathfrak{U}$ is empty. $N^q(\mathfrak{U})$ is called the **q -nerve** of $\mathfrak{U}$.

Let $C^q(\mathfrak{U}, \mathbb{R})$, be the set of functions $c : N^q(\mathfrak{U}) \rightarrow \mathbb{R}$. In other words, $c \in C^0(\mathfrak{U}, \mathbb{R})$ assigns a constant c_U to each $U \in \mathfrak{U}$; $c \in C^1(\mathfrak{U}, \mathbb{R})$ assigns a constant c_{UV} to each $U, V \in \mathfrak{U}$ with $U \cap V \neq \emptyset$; and $c \in C^2(\mathfrak{U}, \mathbb{R})$ assigns a constant c_{UVW} to each $U, V, W \in \mathfrak{U}$ with $U \cap V \cap W \neq \emptyset$. Elements of $C^q(\mathfrak{U}, \mathbb{R})$ are called **q -cochains in $\mathfrak{U}$ with constant coefficients**.

We define $C^q(\mathfrak{U}, \mathcal{E})$ in a similar fashion, with these changes: $f \in C^0(\mathfrak{U}, \mathcal{E})$ assigns a smooth function f_U with domain U ; $f \in C^1(\mathfrak{U}, \mathcal{E})$ assigns a smooth function f_{UV} with domain $U \cap V \neq \emptyset$; and $C^2(\mathfrak{U}, \mathcal{E})$ assigns a smooth function f_{UVW} with domain $U \cap V \cap W \neq \emptyset$. Elements of $C^q(\mathfrak{U}, \mathcal{E})$ are called **q -cochains in $\mathfrak{U}$ with coefficients in $\mathcal{E}$** . At the moment, $\mathcal{E}$ is just a "marker", but in the next section we will give meaning to $\mathcal{E}$ all by itself.

We define maps $\check{d} : C^q(\mathfrak{U}, \cdot) \rightarrow C^{q+1}(\mathfrak{U}, \cdot)$ as follows:

$$\begin{aligned} q = 0 : \quad (\check{d}c)_{UV} &= c_U - c_V \\ q = 1 : \quad (\check{d}c)_{UVW} &= c_{UV} - c_{UW} + c_{VW} \\ q = 2 : \quad \check{d}c &= 0 \end{aligned}$$

For example, when $q = 0$, we have $c \in C^0(\mathfrak{U}, \cdot)$ assigning constants or functions c_U to U 's in $\mathfrak{U}$, and $\check{d}c \in C^1(\mathfrak{U}, \cdot)$ assigning constants or functions $(\check{d}c)_{UV}$ to (U, V) 's in $N^1(\mathfrak{U})$.

Finally, to preserve the supply of the letter ‘ $\mathfrak{U}$ ’ in our typesetter’s storeroom, we will let them be implicit in $C^q(\mathfrak{U}, \cdot)$, writing just $C^q(\cdot)$ from now on.

Let’s apply these definitions to our previous work. A set of local solutions $\bar{f} = \{f_U\}$ to $df = \omega$ is an element of $C^0(\mathcal{E})$. The mapping $\check{d} : C^0(\mathcal{E}) \rightarrow C^1(\mathcal{E})$ sends this to the “mismatch” set $\check{d}\bar{f} = \{f_{UV}\}$; as we’ve seen, this is actually an element of $C^1(\mathbb{R})$.

We are interested in whether $\bar{f} \in C^0(\mathcal{E})$ can be adjusted, by adding an element $\bar{c} \in C^0(\mathbb{R})$, to produce an element $\bar{f} + \bar{c}$ with $\check{d}(\bar{f} + \bar{c}) = 0$. In other words, is $\check{d}\bar{f}$ actually an element of $\check{d}C^0(\mathbb{R})$? If so, then $\bar{f}$ can be adjusted to produce a global solution to $df = \omega$.

Let us capture this question in a commutative diagram:

$$\begin{array}{ccc} C^0(\mathbb{R}) & \xrightarrow{\check{d}} & C^1(\mathbb{R}) \\ \downarrow & & \downarrow \\ C^0(\mathcal{E}) & \xrightarrow{\check{d}} & C^1(\mathcal{E}) \end{array}$$

The downward arrows are inclusion maps. We begin with an element on the bottom left, namely $\bar{f} \in C^0(\mathcal{E})$. We follow the $\check{d}$ arrow across to $\check{d}\bar{f} \in C^1(\mathcal{E})$, and then go up the down-arrow to $\check{d}\bar{f} \in C^1(\mathbb{R})$. If $\check{d}\bar{f}$ is the image of an element in $C^0(\mathbb{R})$, then $\bar{f}$ can be adjusted to produce a global solution to $df = \omega$.

The next step is to embed our diagram in a larger commutative diagram. Define $C^q(\mathcal{Z})$ just like $C^q(\mathcal{E})$, but with closed 1-forms taking the place of functions. For example, an element $\omega \in C^0(\mathcal{Z})$ assigns a closed 1-form ω_U with domain U to each $U \in \mathfrak{U}$. Here is our diagram:

$$\begin{array}{ccccc}
 C^0(\mathbb{R}) & \xrightarrow{\check{d}} & C^1(\mathbb{R}) & \xrightarrow{\check{d}} & C^2(\mathbb{R}) \\
 \downarrow & & \downarrow & & \downarrow \\
 C^0(\mathcal{E}) & \xrightarrow{\check{d}} & C^1(\mathcal{E}) & \xrightarrow{\check{d}} & C^2(\mathcal{E}) \\
 \downarrow & & \downarrow & & \downarrow \\
 C^0(\mathcal{Z}) & \xrightarrow{\check{d}} & C^1(\mathcal{Z}) & \xrightarrow{\check{d}} & C^2(\mathcal{Z})
 \end{array}$$

The second row of down-arrows are differentiations: we go from f_U to df_U , for example.

The following lemma and definition must come as no surprise.

Lemma: $\check{d}\check{d} = 0$.

Proof: Exercise.

Definition: $Z^q(\cdot) = \ker \check{d}$; elements of $Z^q(\cdot)$ are called **cocycles**. $B^q(\cdot) = \text{im } \check{d}$; elements of $B^q(\cdot)$ are called **coboundaries**. $H^q(\cdot) = Z^q(\cdot)/B^q(\cdot)$; $H^q(\cdot)$ is the q -th Čech cohomology group with coefficients in $\cdot$ for the cover $\mathfrak{U}$.

So far, we've shown how to start with a closed 1-form ω defined on the entire surface, and use it to get a “mismatch set” $\check{d}\check{f} \in C^1(\mathbb{R})$. Next we will locate ω in the diagram, “in disguise”. Now, ω is defined on the entire surface, but it determines a 0-cochain in $C^0(\mathcal{Z})$ by restriction: set $\omega_U = \omega|_U$. This cochain is a cocycle since $\omega_U = \omega_V$ on $U \cap V$, both being $\omega|(U \cap V)$.

Conversely, if $\bar{\omega} \in C^0(\mathcal{Z})$ is a cocycle, then the condition $\check{d}\bar{\omega} = 0$ simply says that the $\bar{\omega}_U$'s agree on all overlaps and so piece together to determine a closed 1-form on the whole surface. More generally, we see that cocycles in $Z^0(\cdot)$ are just globally defined (constants/functions/1-forms) in disguise.

From now on we will abuse notation and just write ω for both the global 1-form and for the cocycle (likewise for constants and functions).

So our journey starts with a cocycle ω in the bottom left corner of the diagram. We go up the down-arrow (i.e., integrate) to get the local primitives $\bar{f}_U$; of course, $\bar{f}$ might not be a cocycle. Then it's one step right with $\check{d}$, to a coboundary in $C^1(\mathcal{E})$. This coboundary is also a cocycle, and remains a cocycle when we go up one step to $C^1(\mathbb{R})$. But it won't be a coboundary unless it has a pre-image $\bar{c}$ in $C^0(\mathbb{R})$.

Let's pause and take stock. Our goal was to compute the deRham cohomology group $\text{Rh}^1(X) = \text{Closed}^1(X)/\text{Exact}^1(X)$. The previous paragraph suggests a way to map elements of $\text{Closed}^1(X)$ to elements of $Z^1(\mathbb{R})$; moreover, $\omega \in \text{Exact}^1(X)$ if and only if $\check{d}\bar{f} \in B^1(\mathbb{R})$. In other words, we suspect there is a map $\delta : \text{Rh}^1(X) \rightarrow H^1(\mathbb{R})$.

To show that the map δ really exists, we need to show two things: (a) Our procedure is independent of the choices made in defining $\bar{f} \in C^0(\mathcal{E})$; (b) If ω is exact, then $\check{d}\bar{f} \in Z^1(\mathbb{R})$ is a coboundary, i.e., an element of $B^1(\mathbb{R})$.

Proof of (a): any two choices for the first step, say $\bar{f}$ and $\bar{f}'$, differ by constants: $\bar{f}' = \bar{f} + \bar{c}$, where $\bar{c} \in C^0(\mathbb{R})$. Thus $\check{d}\bar{f}' = \check{d}\bar{f} + \check{d}\bar{c}$, so $\check{d}\bar{f}'$ and $\check{d}\bar{f}$ determine the same element of $H^1(\mathbb{R})$.

Proof of (b): say ω is exact, so $\omega = df$ where f may be regarded as an element of $Z^0(\mathcal{E})$. By part (a), we can make f our choice for $\bar{f} \in C^0(\mathcal{E})$. And $\check{d}f = 0$, giving us an element of $B^1(\mathbb{R})$, as required.

So we have a map $\delta : \text{Rh}^1(X) \rightarrow H^1(\mathbb{R})$. As it happens, δ is an isomorphism³. One half of this is easy: if $\delta[\omega] = 0$, then $\check{d}\bar{f} = \check{d}\bar{c}$ for some $c \in C^0(\mathbb{R})$, so $\check{d}(\bar{f} - \bar{c}) = 0$. So $\bar{f} - \bar{c}$ determines a function $f = \bar{f} - \bar{c}$ with $df = d(\bar{f} - \bar{c}) = \omega$. So $[\omega] = 0$ and $\ker \delta = 0$.

³ provided our surface X has a countable topology, a technical point we won't dwell on.

It's a bit harder to show that $\text{im } \delta = H^1(\mathbb{R})$. Let $\{c_{UV}\} \in Z^1(\mathbb{R})$. So:

$$c_{UV} - c_{UW} + c_{VW} = 0$$

for all $U, V, W \in \mathfrak{U}$. We need to find smooth functions f_U with

$$f_U - f_V = c_{UV}$$

for all $U, V \in \mathfrak{U}$. For if we can do that, then $\omega_U = df_U$ defines a cocycle: $\omega_U - \omega_V = d(f_U - f_V) = dc_{UV} = 0$. Thus the ω_U 's patch together to define an element $\omega \in \text{Closed}^1(X)$, and it is easy to see that δ maps $[\omega]$ to the cohomology class of $\{c_{UV}\}$.

Something more general is true: we can always solve the system of equations

$$f_U - f_V = f_{UV} \tag{1}$$

provided that

$$f_{UV} - f_{UW} + f_{VW} = 0 \tag{2}$$

where the f_{UV} 's are now just smooth functions, not necessarily constants. There is a slick proof using partitions of unity; see Foster, Theorem 12.6 (p.100). Here is a low-tech "proof" that offers some insight into why the result holds. Assume that $\mathfrak{U}$ is countable; assume we've defined f_V for all $V \in \mathfrak{V}$ where $\mathfrak{V}$ is a finite subset of $\mathfrak{U}$. Assume that (1) holds over $\mathfrak{V}$, i.e., $f_{VW} = f_V - f_W$ for all $V, W \in \mathfrak{V}$. Now let U be any element of $\mathfrak{U}$ not in $\mathfrak{V}$. We want to define f_U so that (1) continues to hold over $\mathfrak{V} \cup \{U\}$. If we can do that, we're golden.

The idea is quite simple: for any $V \in \mathfrak{V}$, we rewrite $f_{UV} = f_U - f_V$ as $f_U = f_V + f_{UV}$, defining the unknown f_U in terms of the known f_V and f_{UV} . Of course, this equation only makes sense on $U \cap V$. Very well, define f_U using this equation on $U \cap \bigcup_{V \in \mathfrak{V}} V$, and extend f_U arbitrarily on the rest of U . Two possible obstacles: (a) consistency; (b) smoothness.

The consistency problem is this: if $V, W \in \mathfrak{V}$, how do we know that $f_V + f_{UV}$ and $f_W + f_{UW}$ agree on $U \cap V \cap W$? Well:

$$\begin{aligned}
f_V + f_{UV} &\stackrel{?}{=} f_W + f_{UW} \\
f_V - f_W &\stackrel{?}{=} f_{UW} - f_{UV} \\
f_{VW} &\stackrel{?}{=} f_{UW} - f_{UV} \\
f_{UV} - f_{UW} + f_{VW} &\stackrel{?}{=} 0
\end{aligned}$$

But the last equation is just (2), the cocycle equation for $\{f_{UV}\}$. So if $\{f_{UV}\}$ is a cocycle, then we can define functions f_U satisfying (1).

How do we know we can make f_U smooth? I offer no real proof of this. But recall our intuitive picture of $\mathfrak{U}$ as coming from a tiling of X with convex polygons. The overlaps $U \cap \bigcup_{V \in \mathfrak{V}} V$ amount to a thin “fringe” around the edge of U (or perhaps just part of the edge). Assume inductively that all the f_V with $V \in \mathfrak{V}$ are smooth, so the equations $f_U = f_V + f_{UV}$ define a smooth function on the fringe. It should seem quite plausible that this function can be extended smoothly to all of U .

Summing up, we have shown that $\text{Rh}^1(X)$ is isomorphic to $H^1(\mathbb{R})$. But if we are given a finite covering $\mathfrak{U}$ of X , computing $H^1(\mathbb{R})$ is a (tedious) linear algebra exercise: the $C^q(\mathbb{R})$'s are finite dimensional vector spaces, and the d 's are linear maps.

3 Kirchhoff's Laws

In this section I give two more interpretations of $H^1(\mathbb{R})$. These interpretations will not be needed in the sequel, so you can skip to the next section if you like.

First, an observation. The cocycle equation says $f_{UV} - f_{UW} + f_{VW} = 0$. If we set $U = V = W$, we get $f_{UU} = 0$. If we set $U = W$, we then get

$f_{UV} = -f_{VU}$. Using this, we can rewrite the cocycle equation in the more symmetrical form $f_{UV} + f_{VW} + f_{WU} = 0$.

Let's assume our surface X is oriented, compact, and $\mathfrak{U}$ is finite. Let T be the tiling that gave rise to $\mathfrak{U}$, as described above. We will use the same letter U to refer both to a polygon in T , and to the open set in $\mathfrak{U}$ that comes from the polygon. From T we obtain another tiling T^* , known as the **dual tiling**, by picking a point p_U in the middle of each polygon U in T , and drawing an edge between two of these points when their polygons are adjacent. We assumed earlier that each vertex in T has degree at most 3; it follows that the polygons of T^* are all triangles.

Say now that $\{c_{UV}\}$ is a cocycle. Interpret c_{UV} as a **voltage drop** along the edge from p_U to p_V . The cocycle equation then says that the sum of the voltage drops around each triangle is zero.

Say now that $\{c_{UV}\}$ is a coboundary, coming from $\{c_U\}$. Interpret each c_U as the **voltage at** p_U , i.e., the electrical potential energy at p_U . So in this case, all the voltage drops are true differences of voltages.

Kirchhoff's 2nd law states that the sum of the voltage drops around any circuit is zero. Kirchhoff's 2nd law is equivalent to saying that the voltage drops are all differences of voltages; this is well-known, and not hard to prove in any case.

If we have an electrical network, Kirchhoff's 2nd law can be expressed as a system of linear equations in the c_{UV} 's. We will use a little homology theory to help analyze this system. First we consider a general electrical voltage network N . N consists of nodes and directed edges, with a voltage drop $c_e \in \mathbb{R}$ assigned to each edge e . Let $C_0(N)$ (respectively $C_1(N)$) be the abelian group of integral sums of nodes (respectively edges) of N . Define $de = \text{head}(e) - \text{tail}(e)$ for any edge e , and extend linearly to get $d : C_1(N) \rightarrow C_0(N)$. (Note that $-e$ thus looks like e reversed.) As usual, $Z_1(N) = \ker d$, and elements of $Z_1(N)$ are called **cycles**. Every cycle C has

an associated voltage drop equation, namely $\sum_{e \in C} c_e = 0$.

The next two facts are not hard to prove, but we won't take the time. (1) $Z_1(N)$ is a direct sum of copies of $\mathbb{Z}$. (2) Given a basis for $Z_1(N)$, the associated voltage drop equations are linearly independent, and every other voltage drop equation is a linear combination of them. So the rank of the Kirchhoff voltage equation system is equal to the rank of $Z_1(N)$.

Now, our electrical network $N = T^*$ is drawn on a surface; the edges are all edges of triangles. This enables us to introduce another level to the homology, namely $C_2(N)$, integral sums of triangles. We define $d : C_2(N) \rightarrow C_1(N)$ in the usual way, with $d(\text{triangle}) =$ the sum of the edges taken counterclockwise. (Here we use the orientation of the surface.) Then we define $B_1(N) = \text{im } d$, and $H_1(N) = Z_1(N)/B_1(N)$.

If we add up *all* the triangle cycles, we get 0. (We're assuming our surface is without a border.) Suppose we take all the triangles but one. Do we get a basis for $Z_1(N)$? Not necessarily. For example, if N is drawn on the surface of a torus, then an equatorial or a meridional cycle is not a linear combination of the triangle cycles.

Let's say we have a basis for $Z_1(N)$ consisting of all but one of the triangles, plus r additional cycles. So $H_1(N)$ has rank r . But also, a basis for the Kirchhoff voltage drop equations consists of the triangle equations, plus r additional equations.

An element of $Z^1(\mathbb{R})$ is a set of voltage drops that satisfies all the triangle equations. An element of $B^1(\mathbb{R})$ also satisfies the r additional equations. Basic linear algebra tells us that $B^1(\mathbb{R})$ is a subspace of $Z^1(\mathbb{R})$, and $H^1(\mathbb{R})$ has dimension r (since the r additional equations are independent).

Our second interpretation uses the tiling T directly. We will interpret T as an electrical current network: this has nodes and directed edges, and a **current flow** of strength c_e along each edge e . Kirchhoff's 1st law states that the net flow out of any node is 0.

Let $\{c_{UV}\}$ be a cocycle for T . We want to interpret c_{UV} as the strength of a current flow along the edge $U \cap V$ (i.e., the edge shared by the polygons U and V). There is a sign issue to settle: do we use c_{UV} or c_{VU} , and which direction is positive flow?

We adopt the following conventions. First, arbitrarily assign a directional arrow to each edge in the tiling. If $\overrightarrow{pq}$ is one of these directed edges, let U be the polygon on your left as you traverse the edge from p to q , and let V be the polygon on your right. (Here we use the orientation of the surface to determine left and right.) Interpret c_{UV} as the strength of the flow from p to q ; a negative current from p to q is regarded as a positive current from q to p .

Now let's say $U \cap V \cap W \neq 0$, so the polygons U, V, W meet at a point p . It is not hard to see that the cocycle equation for $U \cap V \cap W$ says that the *net flow out of p is zero*.

Total current in the network is neither created nor destroyed; this follows by summing Kirchhoff's 1st law over all nodes.

Say now that $\{c_{UV}\}$ is a coboundary, coming from $\{c_U\}$. Interpret each c_U as the strength of a counterclockwise circular flow of current around the boundary of U . An edge between two polygons U and V is subjected to two of the flows, in opposite directions; hence has net flow $c_U - c_V$.

Each circular flow has zero net flow at each of nodes it passes through. Thus the total flow obtained by summing all the circular flows satisfies Kirchhoff's 1st law.

As in the first interpretation, we have a relation to the homology group $H_1(T)$. $C_0(T)$ and $C_1(T)$ are as before; $C_2(T)$ consists of sums of the polygons of T . If we add up all the polygonal cycles, we get 0. Say a basis for $Z_1(T)$ consists of all the polygonal cycles but one, plus r additional cycles. With each of the r additional cycles we can associate a flow along that cycle.

An element of $Z^1(\mathbb{R})$ is a current flow satisfying Kirchhoff's 1st law. An element of $B^1(\mathbb{R})$ is a current flow that is a linear combination of the flows around the boundaries of the polygons. So $B^1(\mathbb{R})$ is a subspace of $Z^1(\mathbb{R})$, and $H^1(\mathbb{R})$ has dimension r (since the flows from the r additional cycles are all independent, and independent of the polygonal flows).

4 Bringing in the Sheaves

OK, I couldn't resist. But it's a snazzier section title than "Bringing in the homological algebra", which would be just as accurate.

I won't give the definition of a sheaf; see a textbook, e.g., Forster, *Lectures on Riemann Surfaces*. But to set the stage, I need to recall a few facts about sheaves:

1. If $\mathcal{F}$ is a sheaf of vector spaces on X , then for every open $U \subseteq X$, $\mathcal{F}(U)$ is a vector space.
2. If $V \subseteq U \subseteq X$, U and V open subsets of X , then there is a linear map $\rho_V^U : \mathcal{F}(U) \rightarrow \mathcal{F}(V)$ called the **restriction** map; if $f \in \mathcal{F}(U)$, then we denote $\rho_V^U(f)$ by $f|_V$.
3. For any $x \in X$, we have a vector space $\mathcal{F}_x$, the **stalk at x** , obtained by taking the direct limit of the $\mathcal{F}(U)$'s over the U 's containing x .
4. If $\mathcal{F}$ and $\mathcal{G}$ are sheaves of vector spaces, then a **sheaf homomorphism** $\alpha : \mathcal{F} \rightarrow \mathcal{G}$ assigns linear maps $\alpha(U) : \mathcal{F}(U) \rightarrow \mathcal{G}(U)$, commuting with the restriction maps in the obvious way. A sheaf homomorphism α induces linear maps on the stalks, $\alpha_x : \mathcal{F}_x \rightarrow \mathcal{G}_x$, by taking direct limits.

Three sheaves have played a prominent part in Section 2: the sheaf of constant functions (denoted $\mathbb{R}$, hopefully without confusion), the sheaf $\mathcal{E}$ of

smooth functions, and the sheaf $\mathcal{Z}$ of closed 1-forms. More precisely, $\mathbb{R}(U)$ is the set of constant real-valued functions with domain U ; $\mathcal{E}(U)$ is the set of smooth real-valued functions with domain U ; and $\mathcal{Z}(U)$ is the set of closed real-valued 1-forms with domain U .

The following fact played a key role: if U is simply connected, then

$$0 \rightarrow \mathbb{R}(U) \xrightarrow{i} \mathcal{E}(U) \xrightarrow{d} \mathcal{Z}(U) \rightarrow 0$$

is a short exact sequence (i = inclusion, d = differentiation). It is a routine exercise to show that for any $x \in X$,

$$0 \rightarrow \mathbb{R}_x \xrightarrow{i} \mathcal{E}_x \xrightarrow{d} \mathcal{Z}_x \rightarrow 0$$

is also exact.

All of Section 2 can be carried through in a much more general context. Say we have a sequence of sheaves $\mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H}$ such that

$$0 \rightarrow \mathcal{F}_x \xrightarrow{\alpha} \mathcal{G}_x \xrightarrow{\beta} \mathcal{H}_x \rightarrow 0$$

is exact for all $x \in X$. We say that such sheaves $\mathcal{F}, \mathcal{G}, \mathcal{H}$ form a **short exact sequence**.

If $0 \rightarrow \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \rightarrow 0$ is a short exact sequence of sheaves, then

$$0 \rightarrow \mathcal{F}(U) \xrightarrow{\alpha} \mathcal{G}(U) \xrightarrow{\beta} \mathcal{H}(U)$$

is exact for each open U in X ; this is a routine exercise⁴.

Note the missing 0 on the right end: $\mathcal{G}_x \rightarrow \mathcal{H}_x$ can be onto for each x , without $\mathcal{G}(U) \rightarrow \mathcal{H}(U)$ being onto for all open U . In a slogan: *locally onto does not imply globally onto*. We saw the prototypical case in Section 2: the

⁴ or lemmas 15.6, 15.8 in Forster, if you don't feel like doing it yourself

differentiation map d is locally onto from functions to 1-forms, but fails to be onto when U is not simply connected.

Say we have a cover $\mathfrak{U}$ of X with “nice” U ’s, i.e., $\beta : \mathcal{G}(U) \rightarrow \mathcal{H}(U)$ is onto for each $U \in \mathfrak{U}$. Just as in Section 2, we try to patch together local solutions to $\beta(g_U) = h_U$, where $g_U \in \mathcal{G}(U)$ and $h_U \in \mathcal{H}(U)$. We consider the mismatches $g_{UV} = g_U - g_V$, and try to make them zero by adjusting g_U : $g'_U = g_U + f_U$, $f_U \in \mathcal{F}(U)$.

We are led to define $C^q(\mathfrak{U}, \mathcal{S})$, the **q-cochains over $\mathfrak{U}$ with coefficients in a sheaf $\mathcal{S}$** , same as before. We define $\check{d}$, and then get into a big diagram chase over this:

$$\begin{array}{ccccccc}
 0 & \rightarrow & C^0(\mathcal{F}) & \xrightarrow{\check{d}} & C^1(\mathcal{F}) & \xrightarrow{\check{d}} & C^2(\mathcal{F}) \rightarrow \dots \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \rightarrow & C^0(\mathcal{G}) & \xrightarrow{\check{d}} & C^1(\mathcal{G}) & \xrightarrow{\check{d}} & C^2(\mathcal{G}) \rightarrow \dots \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \rightarrow & C^0(\mathcal{H}) & \xrightarrow{\check{d}} & C^1(\mathcal{H}) & \xrightarrow{\check{d}} & C^2(\mathcal{H}) \rightarrow \dots
 \end{array}$$

The upshot of the diagram chase is the so-called **long exact sequence**:

$$0 \rightarrow H^0(\mathcal{F}) \xrightarrow{\alpha} H^0(\mathcal{G}) \xrightarrow{\beta} H^0(\mathcal{H}) \xrightarrow{\delta} H^1(\mathcal{F}) \xrightarrow{\alpha} \dots$$

Of course, I’ve omitted lots of details (e.g., X should be paracompact Hausdorff); see the textbooks. It can be shown that the long exact sequence for an n -dimensional space stops after $q = n$ (i.e., $H^q(\cdot) = 0$ for all $q > n$). That’s why our discussion in Section 2 went up only to $q = 2$.

The interesting map here is $H^0(\mathcal{H}) \xrightarrow{\delta} H^1(\mathcal{F})$. This is called the **connecting homomorphism**. It’s what gave rise to the isomorphism $\text{Rh}^1(X) \cong H^1(\mathbb{R})$ in the deRham example. Let’s spell out the details. Recall that $\text{Closed}^1(X) \cong Z^0(\mathcal{Z})$ (which equals $H^0(\mathcal{Z})$), and $\text{im } \beta \cong \text{Exact}^1(X)$ (since $\beta = d$). So $\text{Closed}^1(X)/\text{Exact}^1(X) \cong H^0(\mathcal{Z})/dH^0(\mathcal{E}) \cong \delta H^0(\mathcal{Z})$.

That much tells us that $\text{Rh}^1(X)$ is isomorphic to $\text{im } \delta$. In the deRham case, δ was onto. If you go back and examine the argument we gave, you’ll see

that the key was showing that $H^1(\mathcal{E}) = 0$. Indeed, plugging this into the long exact sequence, we see it implies that δ is onto.

In general, if $H^1(\mathcal{G}) = 0$, then the long exact sequence tells us that

$$H^0(\mathcal{H})/\beta H^0(\mathcal{G}) \cong H^1(\mathcal{F})$$

$H^1(\mathcal{G}) = 0$ in turn suggests a certain “floppiness” to the sheaf $\mathcal{G}$: if we can define a sheaf element consistently on the overlaps, then we can extend it over the whole open set. This idea has been captured in the concept of a **fine sheaf**; consult the textbooks for the formal definition.

So if we have a short exact sequence $0 \rightarrow \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \rightarrow 0$, and $\mathcal{G}$ is a fine sheaf, then we get an isomorphism $H^1(\mathcal{F}) \cong H^0(\mathcal{H})/\beta H^0(\mathcal{G})$. But just as before, $H^0(\mathcal{S})$ is canonically isomorphic to $\mathcal{S}(X)$ for any sheaf $\mathcal{S}$, so we get the general result:

$$H^1(X, \mathcal{F}) \cong \mathcal{H}(X)/\beta \mathcal{G}(X) \tag{3}$$

where we’ve made the dependence on X explicit, in anticipation of the next paragraph.

It is time to mention an issue we left hanging in Section 2: the dependence of Čech cohomology on the cover $\mathfrak{U}$. This is dealt with by taking a directed limit over finer and finer open covers; it is not necessary to assume that the U ’s in the cover are simply connected or otherwise “nice”. The resulting cohomology groups are denoted $H^q(X, \mathcal{S})$, showing the dependence on the space X and the sheaf $\mathcal{S}$.

Let’s say U is **nice** if $H^q(U, \mathcal{S}) = 0$ for all $q > 0$. A **Leray cover** $\mathfrak{U}$ has the property that $\bigcap_i U_i$ is nice for all $(U_0, \dots, U_q) \in N^q(\mathfrak{U})$. Leray’s theorem states that for any Leray cover $\mathfrak{U}$, $H^q(X, \mathcal{S}) = H^q(\mathfrak{U}, \mathcal{S})$ for all q . We confined our discussion in Section 2 to Leray covers, to avoid dealing with limits over covers.

Modulo a forest of details, we’ve now seen how our discussion in Section 2 fits into a much more general framework. The real power of this framework

becomes apparent only when we look at other examples of short exact sequences of sheaves. Several are crucial for the theory of Riemann surfaces; we will look at three. From now on, all our functions and forms will be complex-valued rather than real-valued, as you'd expect when working with Riemann surfaces. (We could have done this in Sections 1 and 2, but the real-valued case is easier to visualize.)

The **Dolbeault sequence for functions** is:

$$0 \rightarrow \mathcal{O} \rightarrow \mathcal{E} \xrightarrow{d''} \mathcal{E}^{0,1} \rightarrow 0$$

Here, $\mathcal{O}$ is the sheaf of holomorphic functions on the Riemann surface X . $\mathcal{E}$ is the sheaf of smooth functions on X . $\mathcal{E}^{0,1}$ is the sheaf of 1-forms of the form $g d\bar{z}$ where g is a smooth function. Finally, d'' is defined as follows: given a smooth f , write

$$df = \frac{\partial f}{\partial z} dz + \frac{\partial f}{\partial \bar{z}} d\bar{z}$$

then set $d''f = (\partial f / \partial \bar{z}) d\bar{z}$.

If f is holomorphic, then $\partial f / \partial \bar{z} = 0$, and vice versa, by the Cauchy-Riemann equations. To show that $\mathcal{E} \xrightarrow{d''} \mathcal{E}^{0,1}$ is onto, we need to solve the equation $\frac{\partial f}{\partial \bar{z}} = g$ for any smooth g . However, since we're trying to show sheaf-exactness, we only need to do it locally. Dolbeault's lemma, a non-trivial result, says this is always possible when U is a disk; that's enough to give an exact sequence.

Equation (3) tells us that

$$H^1(X, \mathcal{O}) \cong \mathcal{E}^{0,1}(X) / d''\mathcal{E}(X)$$

Clearly this tells us *something* about holomorphic functions on X , but its intuitive meaning does not leap off the page.

The **Dolbeault sequence for forms** is:

$$0 \rightarrow \Omega \rightarrow \mathcal{E}^{1,0} \xrightarrow{d} \mathcal{E}^{(2)} \rightarrow 0$$

Here, Ω is the sheaf of holomorphic 1-forms, i.e., forms of the form $f dz$ with f holomorphic. $\mathcal{E}^{1,0}$ is the sheaf of 1-forms of the form $f dz$ with f smooth. $\mathcal{E}^{(2)}$ is the sheaf of 2-forms, i.e., forms of the form $g dz \wedge d\bar{z} = -2ig dx \wedge dy$, with g smooth. Again we need a non-trivial result, again called Dolbeault's lemma, to show this is exact.

Equation (3) tells us that

$$H^1(X, \Omega) = \mathcal{E}^{(2)}(X) / d\mathcal{E}^{1,0}(X)$$

Now suppose that X is compact. Any element $\xi \in \mathcal{E}^{(2)}(X)$ can be integrated over X to give us a complex number. Stokes's theorem tells us that if we integrate $d\zeta$ over X , with $\zeta \in \mathcal{E}^{1,0}(X)$, we get zero. So the isomorphism gives us a linear map $H^1(X, \Omega) \rightarrow \mathbb{C}$. This is called the **residue map**, and plays a key role in the theory of Riemann surfaces.

The **Riemann-Roch sequence** is:

$$0 \rightarrow \mathcal{O}_D \rightarrow \mathcal{O}_{D+P} \xrightarrow{\beta} \mathbb{C}_P \rightarrow 0$$

Here, D is a **divisor**, that is, a function $X \rightarrow \mathbb{Z}$ that is non-zero only on an isolated set of points. Let's assume X is compact, so this set is finite. So we can write a divisor, symbolically, as a sum of points. P is a point of X , and $D + P$ is thus also a divisor.

$\mathcal{O}_D$ is the sheaf of meromorphic functions f such that $\text{ord}_x(f) \geq -D(x)$ for all x . In other words, at places where $D(x) > 0$, f is allowed to have a pole of order at most $D(x)$; at places where $D(x) < 0$, f is required to have a zero of order at least $|D(x)|$; everywhere else, f has to be holomorphic. More precisely, these are the conditions imposed on an $f \in \mathcal{O}_D(X)$. For $f_U \in \mathcal{O}_D(U)$, we just ignore any points in D that are not in U , since f_U is not defined outside U .

Note that $\mathcal{O}_D \subseteq \mathcal{O}_{D+P}$. For example, if P does not appear in D , then $\mathcal{O}_{D+P}$ allows f to have a simple pole at P , but otherwise imposes the same conditions on f as $\mathcal{O}_D$ imposes.

If $f_U \in \mathcal{O}_{D+P}(U)$ and if $P \in U$, then we can expand f_U in a Laurent series about P . Say $D(P) = k$ and say we chose a chart with $z = 0$ at P . Then

$$f_U(z) = \sum_{i=-(k+1)}^{\infty} c_i z^i$$

All the terms except $c_{-(k+1)}z^{-(k+1)}$ add up to a function in $\mathcal{O}_D$. So if we let $\beta_U(f_U) = c_{-(k+1)}$, β_U is a homomorphism from $\mathcal{O}_{D+P}$ to $\mathbb{C}$ with kernel $\mathcal{O}_D$.

So long as $P \in U$, we can define β_U with image $\mathbb{C}$. If $P \notin U$, just define β_U to be the zero homomorphism. The β_U 's define a sheaf homomorphism β whose target is the so-called **skyscraper sheaf** $\mathbb{C}_P$: $\mathbb{C}_P(U) = \mathbb{C}$ when $P \in U$, and $\mathbb{C}_P(U) = 0$ otherwise.

$\mathcal{O}_{D+P}$ is not generally a fine sheaf, so we don't get to apply equation (3). However, it is not hard to show that $H^0(X, \mathbb{C}_P) = \mathbb{C}$ and $H^1(X, \mathbb{C}_P) = 0$, so the long exact sequence looks like this:

$$0 \rightarrow \mathcal{O}_D(X) \rightarrow \mathcal{O}_{D+P}(X) \xrightarrow{\beta} \mathbb{C} \xrightarrow{\delta} H^1(X, \mathcal{O}_D) \rightarrow H^1(X, \mathcal{O}_{D+P}) \rightarrow 0$$

So $H^1(X, \mathcal{O}_{D+P}) \cong H^1(X, \mathcal{O}_D)/\delta\mathbb{C}$.

No surprise, the connecting homomorphism δ has something important to tell us. If $\delta = 0$, then $H^1(X, \mathcal{O}_{D+P})$ and $H^1(X, \mathcal{O}_D)$ have the same dimension, but if δ is injective (the only other possibility), then $H^1(X, \mathcal{O}_{D+P})$ has dimension one less than $H^1(X, \mathcal{O}_D)$.

But $\ker \delta = \text{im } \beta$. And it is easy to see that $\text{im } \beta = 0$ if and only if $\mathcal{O}_D(X) = \mathcal{O}_{D+P}(X)$.

Consider the difference in dimensions $\dim \mathcal{O}_D(X) - \dim H^1(X, \mathcal{O}_D)$. From the previous paragraphs, we see that this difference increases by one when we replace D with $D + P$.

From here it is but a hop, skip, and jump to the Riemann-Roch theorem. Let D_0 be the 0 divisor, $D_0(x) = 0$ for all x . Let $\deg D = \sum_x D(x)$. If

D has only positive coefficients (i.e., $D(x) \geq 0$ for all x), we conclude by induction that

$$\dim \mathcal{O}_D(X) - \dim H^1(X, \mathcal{O}_D) = \deg D + \dim \mathcal{O}_{D_0}(X) - \dim H^1(X, \mathcal{O}_{D_0})$$

$\dim H^1(X, \mathcal{O}_{D_0})$ is called the **genus** of X , denoted g . $\dim \mathcal{O}_{D_0}(X) = 1$, since the only functions holomorphic on all of X are constant. $\dim H^1(X, \mathcal{O}_D)$ is called the **index of speciality** of D , denoted $i(D)$. So we have

$$\dim \mathcal{O}_D(X) = 1 - g + \deg D + i(D) \tag{4}$$

provided only that $D(x) \geq 0$ for all x . I leave it as an exercise to remove this assumption: equation (4), the **Riemann-Roch theorem**, holds for all divisors D .

All this hinges on a major assumption, unstated till now: $H^1(X, \mathcal{O}_D)$ is finite-dimensional, likewise $\mathcal{O}_D(X)$. Using the results of the previous paragraphs, it's not hard to reduce this to the finite-dimensionality of $H^1(X, \mathcal{O})$.

But the proof of *that* is deep, requiring power tools from analysis. I shall say no more.