APPLIED TOPOLOGY LECTURE NOTES

CHAD GIUSTI

UPDATED: SEPTEMBER 19, 2017

Δ-complexes

Simplicial complexes will be our fundamental method for specifying topological spaces and doing computations, particularly as we begin to focus on finite sets of data points. They have the advantage of being simple to specify combinatorially and easy to interpret. However, their rigidity makes them somewhat cumbersome to deal with by hand, or even computationally in many circumstances; for example, the smallest simplicial model of a torus has fourteen 2-simplices. Thus, for purposes of being able to get our hands dirty while building these tools, we need a model that allows for smaller constructions of topological spaces of interest.

Definition. A (finite) Δ-*complex*[1] is a topological space X constructed inductively as follows. For each $k = 0, 1, 2, \dots, N$:

i. Fix a finite set S_k. To each $a \in S_k$, associate a standard k-simplex Δ^k_a in $\mathbb{R}^{k+1}$ with basis $e^a_0, e^a_1, \dots e^a_k$.

ii. If $k > 0$, For each $a \in S_k$, for each $i = 0, 1, \dots k$, choose an element $b(a,i) \in S_{k-1}$ and a bijection between the basis $\{e^{b(a,i)}_0, \dots, e^{b(a,i)}_{k-1}\}$ and $\{e^a_0, \dots, \widehat{e^a_i}, \dots, e^a_k\}$. Extend this bijection to a continuous function $j_{a,i,b} : \Delta^{k-1}_{b(a,i)} \hookrightarrow \Delta^k_a$ by linearity[2].

Finally, let $X = \left(\coprod_k \coprod_{a \in S_k} \Delta^k_a\right) / \sim$, where for all pairs (a,i) and $x \in \Delta^{k-1}_{b(a,i)}$, we set $x \sim j_{a,i,b}(x)$. In a Δ-complex, we call the component spaces Δ^k_a *k-cells* and the functions $j_{a,i,b}$ *attaching maps*.

[1] This terminiology is not standardized – various people give these different names. Other choices include *complete semi-simplicial spaces* and *triangulated spaces*.

[2] As in our construction of the realization of a simplicial complex, we can write this linear map $j_{a,i,b}$ as a $(k \times (k-1))$-matrix with a unique 1 in each column, a unique 1 in each row but the $(i+1)$st, and all other entries zero. However, unlike the simplicial case, the choice of row isn't constrained by the column.

In spite of looking superficially similar, Δ-complexes are much more flexible than simplicial complexes. For each cell we add, we are still required to glue the boundary cells in a linear fashion to existing cells one dimension down. However, we can choose any attaching map, and we can glue *any* boundary cell onto *any* existing smaller cell.

Examples.

- Let $S_0 = \{v\}, S_1 = \{w\}$. For $k = 0$, we have only Δ^0_v and for $k = 1$, only Δ^1_w. Because $|S_0| = 1$, we have no choice of what 0-cells to attach our 1-cell to:

$$b(w,0) = b(w,1) = v.$$

Similarly, we only have one choice for the attaching maps:

$$j_{w,0,v} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \qquad j_{w,1,v} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Thus, we have attached both vertices of the 1-cell to the same point. A quick sketch should convince you that $X \simeq S^1$.

- Let $S_0 = \{v_1, v_2, v_3\}, S_1 = \{w_1, w_2, w_3\}, S_2 = \{x_1, x_2\}$. Let's glue the first 1-cell to a single vertex, as before, so

$$b(w_1, 0) = b(w_1, 1) = v_1.$$

However, we'll use the others to attach the remaining vertices to v_1:

$$b(w_2, 0) = v_2 \quad b(w_2, 1) = v_1 \quad b(w_3, 0) = v_3 \quad b(w_3, 1) = v_1.$$

In each case, the attaching map is automatic. Finally, we need to choose attachments for the 2-cells. We'll do this as follows:

$$b(x_1, 0) = w_1 \quad b(x_1, 1) = b(x_1, 2) = w_2$$

$$b(x_2, 0) = w_1 \quad b(x_2, 1) = b(x_2, 2) = w_3$$

Thus, we're going to glue *two* of the boundary edges in each 2-cell to a single 1-cell. Here are the maps we'll choose for the first one:

$$j_{x_1,0,w_1} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad j_{x_1,1,w_2} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}, \qquad j_{x_1,2,w_2} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

What do each of these do? In the first case, we know that w_1 has both boundary points glued to a single vertex, creating a circle. The attaching map $j_{x_1,0,w_1}$ gives us a linear gluing for the cell w_1 along the edge of $\Delta^2_{x_1}$ spanned by $e_1^{x_1}$ and $e_2^{x_1}$, which in the final complex travels along the circle created by w_1. The second and the third are more interesting: in the first, we glue the points $e_0^{w_2}$ to $e_0^{x_1}$, and then travel down the edge of the triangle to $e_1^{x_1}$, gluing linearly; in the second, we start at the same place, but move down the edge to $e_2^{x_1}$. This agreement means we "zip up" the two edges[3]. Thus, we've glued the cell onto the circle and one of the two remaining edges to create a cone.

We'll do exactly the same thing for the other 2-cell[4].

$$j_{x_2,0,w_1} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad j_{x_2,1,w_2} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}, \qquad j_{x_2,2,w_2} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

What did we end up with? Two cones joined along their bases – thus, $X \simeq S^2$.

[3] We could have done this differently. If we'd swapped the direction of gluing in both $j_{x_1,1,w_2}$ and $j_{x_1,2,w_2}$, we'd still have gotten a cone. If we swapped only one, we'd have introduced a "twist" in the simplex before gluing it. This can be useful, but is very hard to envision.

[4] It is worth asking: "Do I need to be careful about which way I choose to attach the two cones to the circle?" Unofficial exercise: convince yourself that it doesn't matter.

- Let $S_0 = \{v_1, v_2\}, S_1 = \{w_1, w_2, w_3, w_4\}, S_2 = \{x_1, x_2\}$. We'll glue the 1-cells on as follows:

$$b(w_1, 0) = b(w_1, 1) = v_1,$$
$$b(w_2, 0) = b(w_3, 0) = b(w_4, 0) = v_2,$$
$$b(w_2, 1) = b(w_3, 1) = b(w_4, 1) = v_1.$$

If we stop right here, we have a more general version of a graph than our original definition allows – self-loops, multiple edges between vertices, and a notion of direction are all built in. However, now we're going to add 2-cells:

$$b(x_1, 0) = w_1 \quad b(x_1, 1) = w_2 \quad b(x_1, 2) = w_3,$$
$$b(x_2, 0) = w_1 \quad b(x_2, 1) = w_2 \quad b(x_2, 2) = w_4.$$

So, these two triangles share a pair of edges, but each has an independent third edge. One of the edges they share is the circle attached to v_1 – if you draw two triangles joined at a circle, you get something that looks a liitle like a cannoli. However, it's easier to envision this situation if you imagine cutting that edge at the vertex and laying the whole thing flat, giving you a square with both pairs of opposite corners waiting to be identified as v_1 and v_2, and the 1-cell w_1 running through the center connecting v_1 to itself. Further, a pair of opposite edges of the square should be glued together. We haven't yet specified the attahing maps for the 2-cells, so we have choices here, but any choice which respects the fact that we have two vertices glues them with a "twist" – this is a Möbius band.

For concreteness, let's write down the maps. The easiest way to do so is to first record the direction each 1-cell is glued in by drawing an arrow along the edge. The three 1-cells that go from v_1 to v_2 are oriented in that direction (per our choice of b maps), so draw arrows on those edges. Since the edge w_1 goes from v_1 to v_1, we can put the arrow in any way we choose without issue. Now, we choose a way to map the pair of 2-cells in that is again consistent with our b maps. In both cases, the boundary component corresponding to dropping e_0 is w_1, so label the "outside" corner at vertex v_2 of each triangle by e_0. Similarly, use the b to put the e_1 and e_2 into each triangle. If the arrow on the edge points from the smaller numbered basis vector to the larger, we're gluing "consistently" and if not we're gluing "inconsistently" with the usual orientation on the 2-simplex. That corresponds to either swapping the rows in the matrix or leaving them be. The following are the proper choice of maps based on my version of this picture.

$$j_{x_1,0,w_1} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad j_{x_1,1,w_2} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \\ 1 & 0 \end{pmatrix}, \qquad j_{x_1,2,w_3} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

$$j_{x_2,0,w_1} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad j_{x_2,1,w_2} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}, \qquad j_{x_2,2,w_4} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

There is an appropriate notion of a homomorphism of Δ-complexes, which takes k-cells to k-cells and requires that for every cell, if we take the boundary of a cell and then apply the morphism, we get the same answer as if we'd applied the morphism to the cell and then taken the boundary – that is, that the attachment structure is maintained. However, since we built these complexes directly as topological spaces, we can just apply continuous functions should we need to.