An illustrated overview of the proof of Claim 2 from the proof of Claim 1 from the proof of Haxell's Theorem on independent transversals

Shagnik Das

In this note I hope the emphasise the main points in the proof of Claim 2, which were probably not very clear in lecture. To begin with, we recall the statement of the claim.

Claim 2. *Let S be a d-simplex with a subdivided boundary. It is possible to obtain a subdivision of S with the same boundary by adding vertices in the interior, in some order, such that each interior vertex is adjacent to at most $2d$ previous vertices*[1]*.*

Proof. The proof of Claim 2 is by induction on the dimension, d. The base case, $d = 1$, is straightforward, as we need only add one internal vertex adjacent to the two boundary vertices of the 1-simplex.

Now we proceed with the induction step, with $d \geq 2$. The first observation is that one can obtain a subdivision of S by adding a single internal vertex and making it adjacent to every boundary vertex. While this is certainly a subdivision, this does not prove Claim 2, because we are potentially introducing too many new edges at once.

Hence what we do is *imagine* we added all these edges, thus giving a subdivision with a lot of imaginary edges. We then will slowly change this subdivision into something acceptable by removing one imaginary edge at a time. This will require the introduction of many new internal vertices, but, using the induction hypothesis and a little topological wizardry, each of these new internal vertices introduces at most $2d$ edges at a time.

Step 0 Starting with our imaginary subdivision (with only one internal vertex, with an imaginary edge to every boundary vertex), order the imaginary edges arbitrarily.

Step 1 In our current subdivision (all real edges + all surviving imaginary edges), consider the next imaginary edge. Call this edge e. In steps 2 to 6, we will modify the subdivision to remove e.

Step 2 In order to remove e, we will have to replace all the subsimplices containing e. Consider the union of these subsimplices — these must account for some volume of space surrounding the edge e. In other words, the union of these simplices looks like[2] a closed d-dimensional ball[3], which we will call B_e, of which e is a diameter.

In three dimensions, you might imagine this union of simplices containing e looks something like a mandarin[4], with each segment of the mandarin representing one of

[1] That is, either boundary vertices or interior vertices that were added earlier.

[2] If you smooth out the surface a little, which doesn't really change anything.

[3] The closed unit d-ball B^d is the set of points $\vec{x} \in \mathbb{R}^d$ such that $\sum_{i=1}^d x_i^2 \leq 1$. Its boundary is the $(d-1)$-dimensional sphere S^{d-1}, consisting of all points $\vec{x} \in \mathbb{R}^d$ with $\sum_{i=1}^d x_i^2 = 1$.

[4] Or tangerine, or clementine, or whatever you prefer.

the subsimplices, and the edge e being the stringy white fibre[5] that runs from top to bottom.

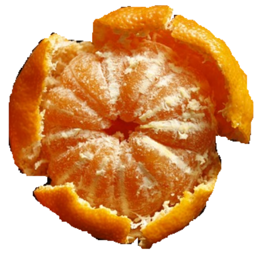

Figure 1: A half-peeled mandarin, or the union of subsimplices containing e?

Step 3 What we need to do is find a new subdivision of B_e that does not use the edge e.[6] We can then use the new subdivision for B_e, and use the rest of the old subdivision for $S \setminus B_e$, and we will have a subdivision of S that does not use e.[7]

Step 4 It is a topological fact that if we consider the union of the subsimplices that make up B_e, and remove the edge e, what we end up with is essentially the $(d-2)$-dimensional sphere S^{d-2}. In our three-dimensional example, we get an equatorial circle on the surface of the mandarin that separates the endpoints of the edge e.[8]

Step 5 We now consider the $(d-1)$-dimensional closed ball we obtain by 'filling in' this sphere S^{d-2}, which we call B'_e. That is, consider the surface of the ball that is exposed when we cut along the equatorial circle. B'_e is topologically equivalent to a $(d-1)$-simplex, and has a subdivided boundary (the remaining faces of the subsimplices surrounding e).

Figure 2: The cut mandarin, exposing the 2-simplex B'_e inside the equator.

[5]I don't know what the proper term for this is.

[6]In our fruity three-dimensional example, we want to find a new way to divide up the interior of the mandarin, rather than splitting it into its segments.

[7]Since our new subdivision for B_e only changes the interior, and not the boundary, we can still combine it with the old subdivision outside B_e.

[8]You can think of the endpoints of the edge e as the little green stub of the mandarin and its antipode, situated at the north and south poles with respect to this equatorial circle.

By induction, we can add interior vertices to B'_e, with each new vertex being adjacent to at most $2(d-1)$ previous vertices, and obtain a subdivision of B'_e.

Note that in Figure 2, if one looks closely, one can make out a triangulation of the surface of the mandarin, with each of the segments of the mandarin contributing a triangle. This is not a valid subdivision of B'_e, because the central vertex has too many edges to boundary vertices. However, this represents the initial imaginary subdivision of B'_e in Step 0 of our induction hypothesis. What we will instead use is the final subdivision the induction hypothesis gives us.

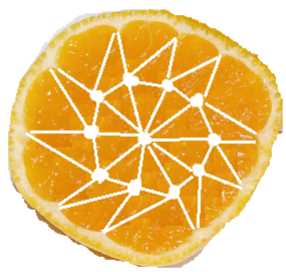

Figure 3: The good final subdivision of B'_e.

Step 6 However, this only gives a subdivision of B'_e, a lower-dimensional simplex inside B_e. What we need is a subdivision of all of B_e.

We do this by lifting each $(d-1)$-simplex in the subdivision of B'_e into two d-simplices — one above the equator and one below. This can be done by adding to each new vertex in B'_e two edges: one to each endpoint of the edge e. This means that instead of being adjacent to at most $2(d-1)$ previous vertices, the new vertices are now adjacent to at most $2d$ previous vertices, but this is still okay. This then provides a subdivision of the whole d-ball B_e into d-simplices, none of which use the edge e.

For our three-dimensional mandarin, this means we find a different way to cut the mandarin. Rather than splitting it into its natural segments, we break it up into different pieces, none of which has a edge that runs the whole length of the mandarin.[9]

Step 7 This new subdivision of B_e, together with the old subdivision of $S \setminus B_e$, gives a subdivision of S that does not use the edge e. Repeating this process for every imaginary edge e, we will obtain a subdivision of S that does not use any imaginary edges.

In this subdivision, every new vertex added was adjacent to at most $2d$ previous vertices (and the original internal vertex is not adjacent to *any* previous vertices, since all of those edges have been removed), completing the induction step and the proof of Claim 2. □

In closing, I concede that some of my references to mandarins / tangerines / clementines / oranges may not have been botanically correct, but if that is your biggest concern with this proof, then this note has served its purpose.

[9]Which is a terrible way to eat a mandarin, if you ask me.