A Infinite sets

This appendix gives a minimum-fuss summary of the set-theoretic notions and facts, such as Zorn's lemma and transfinite induction, that are used in Chapter 8.

Let A, B be sets. If there exists a bijective map between A and B, we write $|A| = |B|$ and say that A and B have *the same cardinality*. This is clearly an equivalence relation between sets, and we may think of the *cardinality* $|A|$ of A as the equivalence class containing A. We write $|A| \leqslant |B|$ if there exists an injective map $A \to B$. This is clearly well defined, and it is a partial ordering: if there are injective maps $A \to B$ and $B \to A$, there is also a bijection $A \to B$.[1] For every set there exists another that is bigger; for example, $|A| < |B|$ when B is the power set of A, the set of all its subsets.

cardinality

The natural numbers are defined inductively as $n := \{0, \ldots, n-1\}$, starting with $0 := \emptyset$. The usual expression of $|A| = n$ can then be read more formally as an abbreviation for $|A| = |n|$.

$\mathbb{N}$

A set A is *finite* if there is a natural number n such that $|A| = n$; otherwise it is *infinite*. A is *countable* if $|A| \leqslant |\mathbb{N}|$, and *countably infinite* if $|A| = |\mathbb{N}|$. A bijection $\mathbb{N} \to A$ is an *enumeration* of A. If A is infinite then $|\mathbb{N}| \leqslant |A|$. Thus, $|\mathbb{N}|$ is the smallest infinite cardinality; it is denoted by $\aleph_0$. There is also a smallest uncountable cardinality, denoted by $\aleph_1$. If $|A| = |\mathbb{R}|$ then A is uncountable, and we say that A has *continuum many* elements. For example, there are continuum many infinite 0–1 sequences. (Whether $|\mathbb{R}|$ is equal to $\aleph_1$ or greater depends on the axioms of set theory assumed; in our context, this question does not arise.) We remark that if A is infinite and its elements are countable sets, then the union of all these sets is no bigger than A itself: $|\bigcup A| \leqslant |A|$.

[1] This is the Cantor-Bernstein theorem; a simple graph-theoretic proof is given in Proposition 8.4.8.

An element x of a partially ordered set X is *minimal* in X if there is no $y \in X$ with $y < x$, and *maximal* if there is no $z \in X$ with $x < z$. A partially ordered set may have one or many elements that are maximal or minimal, or none at all. An *upper bound* in X of a subset $Y \subseteq X$ is any $x \in X$ such that $y \leqslant x$ for all $y \in Y$.

A *chain* is a partially ordered set in which every two elements are comparable. If $(C, \leqslant)$ is a chain, and if $x, y \in C$ satisfy $x < y$ but no element z of C is such that $x < z < y$, then x is called the *predecessor* of y in C, and y the *successor* of x. A set of the form $\{x \in C \mid x < z\}$, for a given $z \in C$, is a proper *initial segment* of C.

successor

A partially ordered set $(X, \leqslant)$ is *well-founded* if every non-empty subset of X has a minimal element, and a well-founded chain is said to be *well-ordered*. For example, $\mathbb{N}$, $\mathbb{Z}$ and $\mathbb{R}$ are all chains (with their usual orderings), but only $\mathbb{N}$ is well-ordered. Note that every element x of a well-ordered set X has a successor (unless x is maximal in X): the unique minimal element of $\{y \in X \mid x < y\} \subset X$. However, an element of a well-ordered set need not have a predecessor, even if it is not minimal. An element that has no predecessor is called a *limit*; for example, the number 1 is a limit in the well-ordered set

well-ordering

limit

$$A = \{1 - \tfrac{1}{n+1} \mid n \in \mathbb{N}\} \cup \{2 - \tfrac{1}{n+1} \mid n \in \mathbb{N}\}$$

of rationals.

One of the many statements equivalent to the axiom of choice (which we assume throughout) is that for every set X there exists a relation by which X is well-ordered:

Well-ordering theorem. *Every set can be well-ordered.*

Two well-ordered sets are said to have *the same order type* if there is a bijection between them which preserves their orders. Thus $\mathbb{N}$ and the set of even natural numbers have the same order type, but this differs from the order type of the set A defined above. Having the same order type is clearly an equivalence relation, which justifies the term if we think of those order types themselves as equivalence classes.

When one considers properties shared by all well-ordered sets of the same order type, it is convenient to represent each order type by a specially chosen set of that type, its *ordinal*. The ordinal representing the order type of $\mathbb{N}$, for instance, is by custom denoted as ω; our example above thus says that the set of even natural numbers has (the) order type (of) ω. Finite chains of the same cardinality always have the same order type; we choose n as the ordinal representing the chains of order n.

ordinals

ω

There is an *addition* of ordinals, defined by taking as the sum $\alpha + \beta$ the ordinal representing the order type of the concatenation of α with β (in this order); note that this is again a well-ordered set. For example, $\alpha + 1$ is the successor of α. Note that no inverse operation '$-$' is defined.

If an ordinal β has the same order type as a proper initial segment of another ordinal α, we write $\beta < \alpha$. For example, we have $0 \leqslant n < \omega$ for every natural number n. It can be shown that $<$ defines an ordering, even a well-ordering, on every set of ordinals. On $\mathbb{N}$, this ordering coincides with the usual one, so our notation is unambiguous.

Since a set S of ordinals is itself well-ordered, it has an order type – just like any other well-ordered set. If the ordinal α is a strict upper bound for S, then the order type of S is at most α; it is equal to α if S consists of all the ordinals up to (but excluding) α. In fact, just like the natural numbers, infinite ordinals are usually defined in such a way that α and $\{\beta \mid \beta < \alpha\}$ are actually identical; then our ordering $<$ for ordinals coincides with the relation $\in$.

This makes it natural to write a well-ordered set S, of order type α say, as a family $S = \{s_\beta \mid \beta < \alpha\}$ with $s_\gamma < s_\beta$ for all $\gamma < \beta < \alpha$. This is common practice when one proves statements about the elements of S by *transfinite induction*, which works as follows.

transfinite induction

Suppose we want to show that every $s \in S$ satisfies some proposition P; let us write $P(s)$ to express that it does. Just as in ordinary induction we prove, for every $\beta < \alpha$, that *if* P holds for every s_γ with $\gamma < \beta$ *then* P also holds for s_β. In practice, we usually have to distinguish the two cases of β being a limit ordinal or a successor. Checking $P(s_0)$ from first principles, as in ordinary induction, is part of the first case, because 0 counts as a limit and the premise of P_γ for all $\gamma < 0$ is void. The conclusion then is that $P(s_\beta)$ for every $\beta < \alpha$, that is, every $s \in S$ satisfies P.

This is certainly simple – but is it correct? Well, any proper justification of transfinite induction requires a formal treatment of set theory, but so does ordinary induction. Informally, what we have shown is that the set

$$\{\beta < \alpha \mid P(s_\beta) \text{ fails}\}$$

has no least element. Since it is well-ordered, it must therefore be empty, so $P(s_\beta)$ holds for all $\beta < \alpha$.

Similarly, we may define things inductively. Such a *recursive definition* specifies for each ordinal α some object x_α, in a way that may refer to the objects x_β with $\beta < \alpha$ (which we think of as 'having been defined earlier'). Our definition of the natural numbers at the start of this appendix is a simple example.

recursive definition

In practice, the definition of x_α often makes sense only for ordinals α less than some fixed ordinal α^*, although the smallest such α^* may not be known in advance. For example, if the x_α are to be distinct vertices picked recursively from a graph G according to some given rules, it is clear that we shall not be able to find such x_α for all $\alpha < \alpha^*$ when $|\alpha^*| > |G|$, because $\alpha \mapsto x_\alpha$ would be an injective map from α^* to $V(G)$ showing that $|\alpha^*| \leqslant |G|$. Since there exist ordinals larger than $|G|$, such as any ordinal equivalent to a well-ordering of the power set of $V(G)$,

this means that our recursion cannot go on indefinitely, i.e. we shall not be able to define x_α for all ordinals α. We may not know in advance when our recursion will get stuck, i.e., which is the smallest ordinal α for which x_α cannot be found in compliance with our rules. But this does not matter: we simply *define* α^* as the first ordinal α for which x_α cannot be found, content ourselves with having defined x_α for all $\alpha < \alpha^*$, and say that our recursion *terminates* at step α^*. (In fact, we usually *want* a recursive definition to terminate. In our example, we might wish to consider the set of all vertices $x \in G$ that got picked by our definition, and this will be the set $\{x_\alpha \mid \alpha < \alpha^*\}$.)

Note that our recursive definition for x_α may involve choices. In our example, x_α might be required to be a neighbour of some x_β with $\beta < \alpha$, but there may be several such x_β, each with several neighbours that have not yet been picked. This does not cause our recursion to get stuck at step α: we just pick one eligible vertex as x_α, and proceed. In other words, we accept $\{x_\alpha \mid \alpha < \alpha^*\}$ as a properly defined set even though we may not 'know' its elements x_α constructively.

Back to proving things, here is a formal statement of Zorn's lemma:

Zorn's Lemma. *Let $(X, \leqslant)$ be a partially ordered set such that every chain in X has an upper bound in X. Then X contains at least one maximal element.*

Note that, in applications of Zorn's lemma, the relation $\leqslant$ need not correspond to an intuitive notion of 'smaller than'. Applied to sets or to graphs, for example, it can stand for '$\supseteq$' just as much as for '$\subseteq$'. Then the 'upper bound' of a chain $\mathcal{C}$ is typically its overall intersection $\bigcap \mathcal{C}$.

Finally, compactness. The infinity lemma discussed in Chapter 8 generalizes as follows.[2] As before, we consider a collection $\{ X_p \mid p \in P \}$ of finite sets, but rather than indexing these by natural numbers we have one such set X_p for every element p of some partially ordered set $(P, \leqslant)$. All we assume about P is that every two elements have a common upper bound: for all p, q there exists an r such that $p \leqslant r$ and $q \leqslant r$. Furthermore, we have maps $f_{qp} \colon X_q \to X_p$ for all $q > p$, which are compatible in that $f_{qp} \circ f_{rq} = f_{rp}$ whenever $r > q > p$.

Generalized Infinity Lemma. *For every such family $\{ X_p \mid p \in P \}$ of finite sets there exists a family $\{ x_p \mid p \in P \}$ of representatives $x_p \in X_p$ such that $f_{qp}(x_q) = x_p$ whenever $q > p$.*

The infinity lemma is clearly a special case of this, with $P = \mathbb{N}$ and the f_{qp} defined by iterating the lemma's predecessor function f.

[2] In category theory, our 'generalized infinity lemma' is known as the fact that the inverse limit of any directed inverse system of finite sets is non-empty. Instead of finite sets one can take any other compact spaces; see Chapter 8.8 for an application.

The following more combinatorial encoding for compactness arguments brings out particularly well the choices involved, and their interdependence. Let X be any set, S a finite set, and $\mathcal{F}$ a set of finite subsets of X. Assume that every $Y \in \mathcal{F}$ comes with a fixed set $\mathcal{A}(Y)$ of $Y \to S$ functions, its *admissible functions*. Call $\mathcal{Y} \subseteq \mathcal{F}$ *compatible* if there exists a function $f\colon X \to S$ all whose restrictions to the sets in $\mathcal{Y}$ are admissible, i.e. which satisfies $f|_Y \in \mathcal{A}(Y)$ for all $Y \in \mathcal{Y}$.

Compactness Principle. *$\mathcal{F}$ is compatible if every finite $\mathcal{Y} \subseteq \mathcal{F}$ is compatible.*

The proofs of the generalized infinity lemma and of the compactness principle are each just a few lines based on Tychonoff's theorem that every product of compact spaces is compact. We illustrate this for the compactness principle.

Think of the set of all $X \to S$ functions as a product of $|X|$ copies of S, so they form a compact space. For every finite $Y \subseteq X$, the set of all functions $f\colon X \to S$ with $f|_Y \in \mathcal{A}(Y)$ is closed (as well as open) in this product space. The result now follows from the 'finite intersection property' of compact spaces, that a family of closed sets has a non-empty intersection as soon as all its finite subfamilies do.

B Surfaces

This appendix offers a summary of background information about surfaces, as needed for an understanding of their role in the proof of the graph minor theorem or the proof of the 'general Kuratowski theorem' for arbitrary surfaces given in Chapter 12.7. In order to be read at a rigorous level it requires familiarity with some basic definitions of general topology (such as of the product and the identification topology), but no more.

A *surface*, for the purpose of this book, is a compact connected[1] Hausdorff topological space S in which every point has a neighbourhood homeomorphic to the Euclidean plane $\mathbb{R}^2$. An *arc*, a *circle*, and a *disc* in S are subsets that are homeomorphic in the subspace topology to the real interval $[0,1]$, to the unit circle $S^1 = \{x \in \mathbb{R}^2 : \|x\| = 1\}$, and to the unit disc $\{x \in \mathbb{R}^2 : \|x\| \leqslant 1\}$ or $\{x \in \mathbb{R}^2 : \|x\| < 1\}$, respectively.

surface

arc

circle S^1

disc

The *components* of a subset X of S are the equivalence classes of points in X where two points are *equivalent* if they can be joined by an arc in X. The surface S itself, being connected, has only one component.

component

The *frontier* of X is the set of all points y in S such that every neighbourhood of y meets both X and $S \setminus X$. The frontier F of X *separates* $S \setminus X$ from X: since $X \cup F$ is closed, every arc from $S \setminus X$ to X has a first point in $X \cup F$, which must lie in F. A component of the frontier of X that is a circle in S is a *boundary circle* of X. A boundary circle of a disc in S is said to *bound* that disc.

frontier

boundary circle

There is a fundamental theorem about surfaces, their *classification*. This says that, up to homeomorphism, every surface can be obtained from the sphere $S^2 = \{x \in \mathbb{R}^3 : \|x\| = 1\}$ by 'adding finitely many handles or finitely many crosscaps', and that surfaces obtained by adding different numbers of handles or crosscaps are distinct. We shall not need the classification theorem, but to form a picture let us see what the

sphere S^2

[1] Throughout this appendix, 'connected' means 'arc-connected'.

handle above operations mean. To *add a handle* to a surface S, we remove two open discs whose closures in S are disjoint, and identify[2] their boundary circles with the circles $S^1 \times \{0\}$ and $S^1 \times \{1\}$ of a copy of $S^1 \times [0,1]$ disjoint from S. To *add a crosscap*, we remove one open disc, and then identify opposite points on its boundary circle in pairs.

crosscap

In order to see that these operations do indeed give new surfaces, we have to check that every identification point ends up with a neighbourhood homeomorphic to $\mathbb{R}^2$. To do this rigorously, let us first look at circles more generally.

cylinder A *cylinder* is the product space $S^1 \times [0,1]$, or any space homeomorphic to it. Its *middle circle* is the circle $S^1 \times \{\frac{1}{2}\}$. A *Möbius strip* is any space homeomorphic to the product space $[0,1] \times [0,1]$ after identification of $(1,y)$ with $(0,1-y)$ for all $y \in [0,1]$. Its *middle circle* is the set $\{(x,\frac{1}{2}) \mid 0 < x < 1\} \cup \{p\}$, where p is the point resulting from the identification of $(1,\frac{1}{2})$ with $(0,\frac{1}{2})$. It can be shown[3] that every circle C in a surface S is the middle circle of a suitable cylinder or Möbius strip N in S, which can be chosen small enough to avoid any given compact subset of $S \setminus C$. If this *strip neighbourhood* is a cylinder, then $N \setminus C$ has two components and we call C *two-sided*; if it is a Möbius strip, then $N \setminus C$ has only one component and we call C *one-sided*.

Möbius strip

strip neighbourhood

two-sided

one-sided

Using small neighbourhoods inside a strip neighbourhood of the (two-sided) boundary circle of the disc or discs we removed from S in order to attach a crosscap or handle, one can show easily that both operations do produce new surfaces.

separating circle Since S is connected, $S \setminus C$ cannot have more components than $N \setminus C$. If $S \setminus C$ has two components, we call C a *separating* circle in S; if it has only one, then C is *non-separating*. While one-sided circles are obviously non-separating, two-sided circles can be either separating or non-separating. For example, the middle circle of a cylinder added to S as a 'handle' is a two-sided non-separating circle in the new surface obtained. When S' is obtained from S by adding a crosscap in place of a disc D, then every arc in S that runs half-way round the boundary circle of D becomes a one-sided circle in S'.

The classification theorem thus has the following corollary:

Lemma B.1. *Every surface other than the sphere contains a non-separating circle.*

---

[2] This is made precise by the *identification topology*, whose formal definition can be found in any topology book. Since S^1 has two possible orientations, two copies of S^1 can be identified in two essentially different ways. The corresponding two ways of adding a handle yield different new surfaces. For the classification one only uses one of these, the way that preserves the orientability of the surface (as in Figure B.1).

[3] In principle, the strip neighbourhood N is constructed as in the proof of Lemma 4.2.2, using the compactness of C. However since we are not in a piecewise linear setting now, the construction is considerably more complicated.

We shall see below that, in a sense, our two examples of non-separating circles are all there are: cutting a surface along any non-separating circle (and patching up the holes) will always produce a surface with fewer handles or crosscaps.

An *embedding* $G \hookrightarrow S$ of a graph G in S is a map σ that maps the vertices of G to distinct points in S and its edges xy to $\sigma(x)$–$\sigma(y)$ arcs in S, so that no inner point of such an arc is the image of a vertex or lies on another arc. We then write $\sigma(G)$ for the union of all those points and arcs in S. A *face* of G in S is a component of $S \setminus \sigma(G)$, and the subgraph of G that σ maps to the frontier of this face is its *boundary*. Note that while faces in the sphere are always discs (if G is connected), in general they need not be.

embedding $\sigma\colon G \hookrightarrow S$

face

boundary

One can prove that in every surface one can embed a suitable graph so that every face becomes a disc. The following general version of Euler's theorem 4.2.9 therefore applies to all surfaces:

Theorem B.2. *For every surface S there exists an integer $\chi(S)$ such that whenever a graph G with n vertices and m edges is embedded in S so that there are ℓ faces and every face is a disc, we have*

$$n - m + \ell = \chi(S)\,.$$

This invariant χ of S is its *Euler characteristic*. For computational simplicity we usually work instead with the derived invariant

$$\varepsilon(S) := 2 - \chi(S)\,, \qquad \varepsilon(S)$$

the *Euler genus* of S, because χ is negative for most surfaces but ε takes its values in $\mathbb{N}$ (see below).

Euler genus

Perhaps the most striking feature of Euler's theorem is that it works with almost any graph embedded in S. This makes it easy to see how the Euler genus is affected by the addition of a handle or crosscap.

Indeed, let D and D' be two open discs in S that we wish to remove in order to attach a handle there. Let G be any graph embedded in S so that every face is a disc. If necessary, shift G on S so that D and D' each lie inside a face, f and f', say. Add cycles C and C' on the boundary circles of D and D', and join them by an edge to the old boundaries of f and f', respectively. Then every face of the resulting graph is again a disc, and D and D' are among these. Now remove D and D', and add a handle with an additional C–C' edge running along it. This operation makes the new handle into one new face, which is a disc. It thus reduces the total number of faces by 1 (since we lost D and D' but gained the new face on the handle) and increases the number of edges by 1, but leaves the number of vertices unchanged. As a result, ε grows by 2.

Similarly, replacing a disc D bounded by a cycle $C \subseteq G$ with a crosscap decreases the number of faces by 1 (since we lose D), but leaves $n - m$ unchanged if we arrange the cycle C in such a way that vertices get identified with vertices when we identify opposite points.

We have thus shown the following:

Lemma B.3.

(i) *Adding a handle to a surface raises its Euler genus by 2.*

(ii) *Adding a crosscap to a surface raises its Euler genus by 1.* □

Since the sphere has Euler genus 0 (Theorem 4.2.9), the classification theorem and Lemma B.3 tell us that ε has all its values in $\mathbb{N}$. We may thus try to prove theorems about surfaces by induction on ε. For the induction step, we could simply undo the addition of a handle or crosscap described earlier, cutting along the new non-separating circle it produced (which runs around the new handle or 'half-way' around the crosscap) and restoring the old surface by putting back the disc or discs we removed. A problem with this is that we do not normally know where on our surface this circle lies, say with respect to a given graph embedded in it.

However, the genus-reducing cut-and-paste operation can be carried out with any non-separating circle: we do not have to use one that we know came from a new handle or crosscap. This is an example of a more general technique known as *surgery*, and works as follows.

cutting

Let C be a non-separating circle in a surface $S \neq S^2$. To *cut* S *along* C, we form a new space S' from S by replacing every point $x \in C$ with two points x', x'' and defining the topology on the modified set as follows.[4] Let N be any strip neighbourhood of C in S, and put $X' := \{x' \mid x \in C\}$ and $X'' := \{x'' \mid x \in C\}$. If N is a cylinder, then $N \setminus C$ has two components N' and N'', and we choose the neighbourhoods of the new points x' and x'' in S' so that X' and X'' become boundary circles of N' and N'' in S', respectively, and $N' \cup X'$ and $N'' \cup X''$ become disjoint cylinders in S'. If N is a Möbius strip, we choose these neighbourhoods so that X' and X'' each form an arc in S' and $X' \cup X''$ is a boundary circle of $N \setminus C$ in S', with $(N \setminus C) \cup X' \cup X''$ forming one cylinder in S'. Finally, we turn S' into a surface by *capping its holes*: for each of the (two or one) boundary circles X' and X'' or $X' \cup X''$ of $S \setminus C$ in S' we take a disc disjoint from S' and identify its boundary circle with X', X'' or $X' \cup X''$, respectively, so that the space obtained is again a surface.

capping

[4] The description that follows may sound complicated, but it is not: working in our concrete models of the cylinder and the Möbius strip it is easy to write down explicit neighbourhood bases that define a topology with the properties stated. As all we want is to obtain some surface of smaller genus, we do not care about uniqueness (which will follow anyhow from Lemma B.4 and the classification).

Computing how these operations affect the Euler genus of S is again easy, assuming we can embed a graph in S so that every face is a disc and C is the image of a cycle. (This can always be done, but it is not easy to prove.[5]) Indeed, by doubling C we left $n-m$ unchanged, because a cycle has the same number of vertices as edges. So all we changed was ℓ, which increased by 2 in the first case and by 1 in the second.

Lemma B.4. *Let C be any non-separating circle in a surface S, and let S' be obtained from S by cutting along C and capping the hole or holes.*

(i) *If C is one-sided in S, then $\varepsilon(S') = \varepsilon(S) - 1$.*

(ii) *If C is two-sided in S, then $\varepsilon(S') = \varepsilon(S) - 2$.* □

Lemma B.4 gives us a large supply of circles to cut along in an induction on the Euler genus. Still, it is sometimes more convenient to cut along a separating circle, and many of these can be used too:

Lemma B.5. *Let C be a separating circle in a surface S, and let S' and S'' be the two surfaces obtained from S by cutting along C and capping the holes. Then*

$$\varepsilon(S) = \varepsilon(S') + \varepsilon(S'') .$$

In particular, if C does not bound a disc in S, both S' and S'' have smaller Euler genus than S.

Proof. As before, embed a graph G in S so that every face is a disc and C is the image of a cycle in G, and let $G' \hookrightarrow S'$ and $G'' \hookrightarrow S''$ be the two graphs obtained in the surgery. Thus, G' and G'' both contain a copy of the cycle on C, which we assume to have k vertices and edges. Then, with the obvious notation, we have

$$\begin{aligned} \varepsilon(S') + \varepsilon(S'') &= (2 - n' + m' - \ell') + (2 - n'' + m'' - \ell'') \\ &= 4 - (n+k) + (m+k) - (\ell+2) \\ &= 2 - n + m - \ell \\ &= \varepsilon(S) . \end{aligned}$$

Now if S' (say) is a sphere, then $S' \cap S$ was a disc in S bounded by C. Hence, if C does not bound a disc in S then $\varepsilon(S')$ and $\varepsilon(S'')$ are both non-zero, giving the second statement of the lemma. □

We now apply these techniques to prove a lemma for our direct proof in Chapter 12 of the 'Kuratowski theorem for arbitrary surfaces', Corollary 12.7.3.

[5] Perhaps the simplest proof was given by C. Thomassen, The Jordan-Schoenflies theorem and the classification of surfaces, *Amer. Math. Monthly* **99** (1992), 116–130.

[12.7.4] **Lemma B.6.** *Let S be a surface, and let $\mathcal{C}$ be a finite set of disjoint circles in S. Assume that $S \setminus \bigcup \mathcal{C}$ has a component D_0 whose closure in S meets every circle in $\mathcal{C}$, and that no circle in $\mathcal{C}$ bounds a disc in S that is disjoint from D_0. Then $\varepsilon(S) \geqslant |\mathcal{C}|$.*

Proof. We begin with the observation that the closure of D_0 not only meets but even contains every circle $C \in \mathcal{C}$. This is because C has a strip neighbourhood N disjoint from all the other circles in $\mathcal{C}$ (since their union is compact), and each of the (one or two) components of $N \setminus C$ has all of C in its closure. Since D_0 meets, and hence contains, at least one component of $N \setminus C$, its closure contains C.

$\mathcal{C}_1$, $\mathcal{C}_2^1$, $\mathcal{C}_2^2$ Let us partition $\mathcal{C}$ as $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2^1 \cup \mathcal{C}_2^2$, where the circles in $\mathcal{C}_1$ are one-sided, those in $\mathcal{C}_2^1$ are two-sided but non-separating, and those in $\mathcal{C}_2^2$ are separating. We shall, in turn, cut along all the circles in $\mathcal{C}_1$, some $|\mathcal{C}_2^2|$ non-separating circles not in $\mathcal{C}$, and at least half the circles in $\mathcal{C}_2^1$. This $S_0, \dots, S_n$ will give us a sequence $S_0, \dots, S_n$ of surfaces, where $S_0 = S$, and S_{i+1} C_i is obtained from S_i by cutting along a circle C_i and capping the hole(s). Our task will be to ensure that C_i is non-separating in S_i for every $i = 0, \dots, n-1$. Then Lemma B.4 will imply that $\varepsilon(S_{i+1}) \leqslant \varepsilon(S_i) - 1$ for all i and $\varepsilon(S_{i+1}) \leqslant \varepsilon(S_i) - 2$ whenever $C_i \in \mathcal{C}_2^1$, giving

$$\varepsilon(S) \geqslant \varepsilon(S_n) + |\mathcal{C}_1| + |\mathcal{C}_2^2| + 2\,|\mathcal{C}_2^1|/2 \geqslant |\mathcal{C}|$$

as desired.

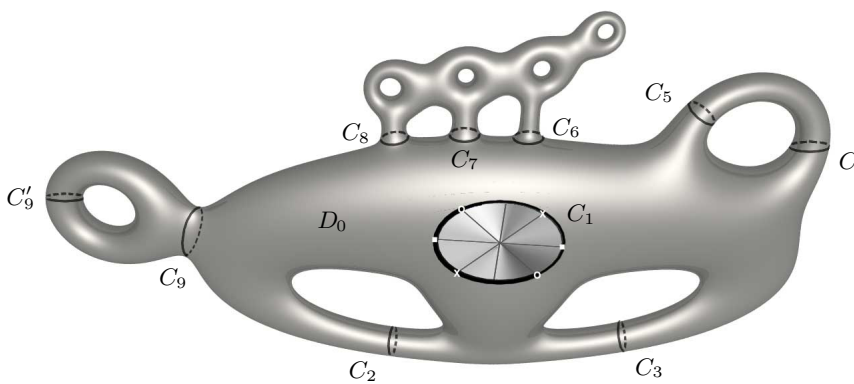

Fig. B.1. Cutting the 1-sided circle C_1 and the 2-sided circles C_2, C_3 and C_5, C_7, C_8 and C_9' does not separate S

Cutting along the circles in $\mathcal{C}_1$ (and capping the holes) is straightforward: since these circles are one-sided, they are always non-separating.

Next, we consider the circles in $\mathcal{C}_2^2$, such as C_9 in Figure B.1. For every $C \in \mathcal{C}_2^2$, denote by $D(C)$ the component of $S \setminus C$ that does not contain D_0. Since every circle in $\mathcal{C}$ lies in the closure of D_0 but no point

of $D(C)$ does, these $D(C)$ are also components of $S \setminus \bigcup \mathcal{C}$. In particular, they are disjoint for different C. Thus, each $D(C)$ will also be a component of $S_i \setminus C$, where S_i is the current surface after any surgery performed on the circles in $\mathcal{C}_1$ and inside $D(C')$ for some $C' \neq C$. Given a fixed circle $C \in \mathcal{C}_2^2$, let S' be the surface obtained from $D(C)$ by capping its hole. Since C does not bound a disc in S that is disjoint from D_0, we know that S' is not a sphere and hence contains a non-separating circle C' (Lemma B.1). We choose C' so that it avoids the cap we added to form S', i.e. so that $C' \subseteq S \setminus C$. Then C' is also non-separating in the current surface S_i (since every point of $S_i \setminus C'$ can be joined by an arc in $S_i \setminus C'$ to C, which is connected), and we may select C' as a circle C_i to cut along.

It remains to select at least half of the circles in $\mathcal{C}_2^1$ as circles C_i to cut along. We begin by selecting all those whose entire strip neighbourhoods (i.e., both their 'sides') lie in D_0. (In Figure B.1, these are the circles C_2 and C_3.) These circles C are non-separating also in the surface S_i current before they are cut, because D_0 will lie inside a component of $S_i \setminus C$. Every other $C \in \mathcal{C}_2^1$ lies in the closure also of a component $D(C) \neq D_0$ of $S \setminus \bigcup \mathcal{C}$. (In Figure B.1, these are the circles $C_4, \ldots, C_8$.) For every component D of $S \setminus \bigcup \mathcal{C}$ we select all but one of the circles $C \in \mathcal{C}_2^1$ with $D(C) = D$ as a cutting circle C_i. Clearly, each of these C_i will be non-separating also in its current surface S_i, and their total number at least $|\mathcal{C}_2^1|/2$. □