ASIAN J. MATH.
Vol. 13, No. 3, pp. 385–404, September 2009
© 2009 International Press
008

THE TOPOLOGY OF MODULI SPACES OF TROPICAL CURVES WITH MARKED POINTS*

DMITRY N. KOZLOV†

Abstract. In this paper we study the topology of moduli spaces of tropical curves of genus g with n marked points. We view the moduli spaces as being embedded in a larger space, which we call the *moduli space of metric graphs with n marked points.* We describe the shrinking bridges strong deformation retraction, which leads to a substantial simplification of all these moduli spaces.

In the rest of the paper, this reduction is used to analyze the case of genus 1. The corresponding moduli space is presented as a quotient space of a torus with respect to the conjugation $\mathbb{Z}_2$-action; and furthermore, as a homotopy colimit over a simple diagram. The latter allows us to compute all Betti numbers of this moduli space with coefficients in $\mathbb{Z}_2$.

Key words. Tropical geometry, combinatorial algebraic topology, moduli spaces, complexes of trees, metric graphs.

AMS subject classifications. Primary 55U10; Secondary 14M99

1. Moduli spaces in tropical geometry.

In this paper we study the moduli spaces of tropical curves with marked points from the topological point of view. These spaces were recently introduced by Mikhalkin in [Mi07, Mi06] as important gadgets in tropical geometry. A nice general introduction to the subject of tropical geometry can be found in [RST05]. In addition, the tropical analog of Grassmannian, the parametrizing space of all tropical linear spaces, was introduced by Speyer and Sturmfels in [SS04]. The moduli spaces of tropical curves have also been studied by Gathmann, Kerber, and Markwig in [GKM07], using tropical fans.

Mikhalkin's investigation centers on the tropical geometry of these spaces, going in particular depth in the case of genus 0; here we complement his pioneering work by focusing exclusively on the topological properties. Accordingly, we define the moduli spaces as embedded in a larger space which we call the *moduli space of metric graphs with n marked points,* in particular, they inherit the natural topology from that larger space. We then prove that a simultaneous contraction of all the bridges is a strong deformation retraction. The rigorous proof of this fact is somewhat technical and requires corresponding precision in the definition of the topology on the set of isometry classes of metric graphs with n marked points. The main technical problem is to take care of the symmetries arising from the action of the automorphism group of the graph. All of this is done in Section 3.

Since all the edges in a tree are bridges, the shrinking bridges strong deformation retraction contracts the entire moduli space of tropical curves of genus 0 with n marked points to a point. Therefore, the corresponding space is not of much interest from the topological point of view, as far as our current study is concerned.[1]

The first topologically interesting case, that of genus 1, is dealt with in Section 4. In this framework the shrinking bridges strong deformation retraction simplifies the analysis of the space dramatically, reducing it to the quotient of the torus by the

*Received January 21, 2009; accepted for publication September 30, 2009.

†Department of Mathematics, University of Bremen, 28334 Bremen, Federal Republic of Germany (dfk@math.uni-bremen.de). This research was supported by University of Bremen, as part of AG CALTOP.

[1]It should be noted that the topology of the link of that vertex is interesting, and important in the study of phylogenetic trees, see [BHV01].

conjugation action of $\mathbb{Z}_2$. This can be done because a bridge-free graph of genus 1 is isomorphic to a cycle. At present it seems difficult to describe the homotopy type of the obtained space in simple terms. However, it is possible to view it as a homotopy colimit of a simple diagram (actually just a gluing of two identical mapping cylinders). We use this presentation to compute the Betti numbers of the space with coefficients in $\mathbb{Z}_2$, as well as to make a conjecture about the homology groups with integer coefficients. The case of higher genus is dealt with in [Ko09].

We refer to [Ko07] for the concepts and tools of Combinatorial Algebraic Topology which are used throughout this paper. The few notions of the category theory that we use here can also be found in [Ko07, Chapters 4 and 15] or in [McL98].

2. Metric graphs.

2.1. Graphs and graph homomorphisms. All the graphs considered in this paper will be finite (meaning that both the set of vertices and the set of edges are finite) and undirected, however loops and multiple edges are allowed, and will, in fact, be essential for our investigation. Let us now fix our notations.

Definition 2.1. *A* **graph** *G is a pair of finite sets $V(G)$ and $DE(G)$ equipped with set maps $\mathfrak{op} : DE(G) \to DE(G)$ and $\partial^\bullet : DE(G) \to V(G)$, such that*
- $\mathfrak{op} \circ \mathfrak{op} = id_{DE(G)}$,
- *the map $\mathfrak{op}$ has no fixed points.*

When more precise specification is needed, we also write $\mathfrak{op}_G$ instead of $\mathfrak{op}$. We shall call $V(G)$ the *set of vertices*. We think of the elements of $DE(G)$ as *directed edges*, where the map $\mathfrak{op}$ changes orientation of such a directed edge to the opposite one, and the map $\partial^\bullet$ takes a directed edge to its source vertex. Accordingly, we introduce notation $\partial_\bullet := \partial^\bullet \circ \mathfrak{op} : DE(G) \to V(G)$ for the target vertex map. One can think of a graph as a diagram of sets over the category with two objects and three non-identity morphisms, see Figure 2.1, with the compositions of the morphisms given by the rules $\mathfrak{op}^2 = \mathrm{id}$ and $\partial^\bullet \circ \mathfrak{op} = \partial_\bullet$.

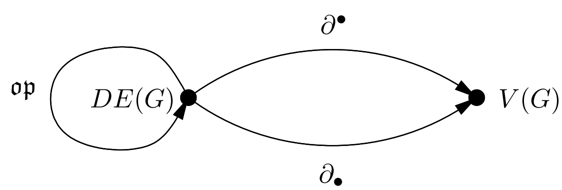

Fig. 2.1. *A graph G viewed as a diagram of sets.*

The condition that $\mathfrak{op}$ has no fixed points implies that elements of $DE(G)$ come in pairs $\{e, \mathfrak{op}(e)\}$. These pairs, or equivalence classes, are the *(undirected) edges* of G and we denote the corresponding set by $E(G) := DE(G)/(e \sim \mathfrak{op}(e))$. Elements $e \in DE(G)$, such that $\partial^\bullet e = \partial_\bullet e$ are called *directed loops*, and the corresponding equivalence classes in $E(G)$ are called *loops*. For arbitrary vertices $x, y \in V(G)$, we let

$$DE(x,y) := \{e \in DE(G) \,|\, \partial^\bullet e = x,\ \partial_\bullet e = y\}$$

denote the set of edges directed from x to y. Clearly $DE(G) = \bigcup_{x,y \in V(G)} DE(x,y)$, and the union is disjoint. The map $\mathfrak{op}$ is a bijection between $DE(x,y)$ and $DE(y,x)$, and we set

$$E(x,y) := (DE(x,y) \cup DE(y,x))/(e \sim \mathfrak{op}(e)) \subseteq E(G).$$

We have $E(x,y) = E(y,x)$, for all $x, y \in V(G)$. In particular $DE(x,x)$ denotes the set of directed loops at x, and $E(x,x) := DE(x,x)/(e \sim \mathfrak{op}(e))$ denotes the set of loops at x.

As an example, for a graph G with one vertex and one loop we have $V(G) = \{v\}$, $DE(v,v) = DE(G) = \{e_1, e_2\}$, with $\mathfrak{op}(e_1) = e_2$, $\mathfrak{op}(e_2) = e_1$, $\partial^\bullet e_1 = \partial_\bullet e_1 = \partial^\bullet e_2 = \partial_\bullet e_2 = v$, and $E(v,v) = E(G) = \{\{e_1, e_2\}\}$, $|E(G)| = 1$.

Definition 2.2. *For two graphs G and H, a* **graph homomorphism** *from G to H is simply a map between corresponding diagrams of sets. In concrete terms, it consists of two set maps $\varphi_V : V(G) \to V(H)$ and $\varphi_{DE} : DE(G) \to DE(H)$, such that $\varphi_{DE} \circ \mathfrak{op}_G = \mathfrak{op}_H \circ \varphi_{DE}$, and $\varphi_V \circ \partial^\bullet = \partial^\bullet \circ \varphi_{DE}$.*

A graph homomorphism is called a **graph isomorphism** *if the involved set maps φ_V and φ_{DE} are bijections.*

Since $\varphi_{DE} \circ \mathfrak{op}_G = \mathfrak{op}_H \circ \varphi_{DE}$, we obtain induced map $\varphi_E : E(G) \to E(H)$. For example, the graph with one vertex and one loop described above has two automorphisms, i.e., invertible graph homomorphisms to itself. Both are identity maps on the sets $V(G)$ and $E(G)$. However, on the set $DE(G)$, one is the identity map, and the other one swaps the directed edges e_1 and e_2.

2.2. The CW complex $\Delta(G)$ and the genus of a graph. We shall now associate a topological space $\Delta(G)$ to a graph G. To avoid making noncanonical choices, and to aid our further considerations, we would like to think of the space $\Delta(G)$ as obtained by gluing together closed intervals corresponding to elements of $DE(G)$. For this, let B_e denote the closed interval, a copy of $[0,1] \subset \mathbb{R}$, corresponding to the element $e \in DE(G)$, and let $\Omega(G)$ be the union of all disjoint closed intervals B_e. Let furthermore $V(G)$ also denote the discrete set of points indexed by elements of $V(G)$, and let $W(G)$ be the discrete set of points indexed by the union

$$\{(e, \partial^\bullet e) \mid e \in DE(G)\} \cup \{(e, \partial_\bullet e) \mid e \in DE(G)\} \subseteq DE(G) \times V(G).$$

Consider the following maps:

- a map $\alpha : W(G) \to V(G)$, defined by $\alpha(e,v) := v$;
- a map $\beta : W(G) \to \Omega(G)$, which takes $(e, \partial^\bullet e)$ to the point in B_e corresponding to 0, and takes $(e, \partial_\bullet e)$ to the point in B_e corresponding to 1;
- a map $\gamma : \Omega(G) \to \Omega(G)$ which takes a point in B_e corresponding to $x \in [0,1]$ to the point in $B_{\mathfrak{op}(e)}$ corresponding to $1 - x$, for all $e \in DE(G)$, and all $x \in [0,1]$.

Together with spaces $W(G)$, $V(G)$, and $\Omega(G)$ these maps form a diagram shown in Figure 2.2, which one can think of as a gluing data for $\Delta(G)$. This is made precise by the following definition.

Definition 2.3. *For an arbitrary graph G, we let the topological space $\Delta(G)$ be the colimit of the diagram shown in Figure 2.2.*

We let $q_G : V(G) \cup W(G) \cup \Omega(G) \to \Delta(G)$ denote the map induced by the structural maps from the spaces in the diagram to the colimit of that diagram. The map $q_G|_{\Omega(G)}$, and hence also the map q_G, is surjective.

Clearly, $\Delta(G)$ has a structure of a 1-dimensional CW complex, whose 0-cells are indexed by the vertices of G, 1-cells are indexed by the edges of G, i.e., by $\mathfrak{op}$-invariant

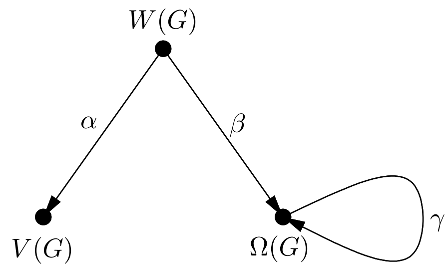

Fig. 2.2. *The gluing data for the space* $\Delta(G)$.

pairs of elements from $DE(G)$, and the attachment maps are given by the vertex-edge incidences. When appropriate we shall identify vertices of G with corresponding 0-cells of $\Delta(G)$, and edges of G with corresponding open 1-cells of $\Delta(G)$.

A graph homomorphism from G to H induces a natural diagram map from the gluing data of G to the gluing data of H, and therefore it also induces a natural CW map from $\Delta(G)$ to $\Delta(H)$, which is in fact a homeomorphism when restricted to any open 1-cell of $\Delta(G)$. We denote both maps by φ_Δ. When φ is a graph isomorphism, the map φ_Δ is a CW isomorphism.

As the last piece of terminology here, the first Betti number of $\Delta(G)$ will be called the *genus* of G, and denoted by $g(G)$. Clearly $g(G) = |E(G)| - |V(G)| + 1$.

2.3. Metric graphs with marked points. To introduce more structure we now vary the lengths of the edges.

Definition 2.4. *Let G be a graph. We say that G is a* **metric graph** *when we are given a function $l_G : E(G) \to (0, \infty)$, called the* **edge-length function**.

The index G will be skipped in l_G whenever it is clear which graph is considered. We shall also use l_G to denote the corresponding $\mathfrak{op}$-invariant function $l_G : DE(G) \to (0, \infty)$.

Given a metric graph (G, l_G), there is a standard way to use the function l_G to turn the topological space $\Delta(G)$ into a metric space, which we now describe. We identify each B_e with the metric space $[0, l_G(e)]$, instead of the topological space $[0, 1]$, with the standard distance function given by $\tilde{d}(x, y) = |x - y|$, for $x, y \in [0, l_G(e)]$.

We can now define a distance function d on $\Delta(G)$ as follows: $d(x, x) := 0$, for all $x \in \Delta(G)$, and for $x, y \in \Delta(G)$, $x \neq y$ we set

$$d(x, y) := \min \sum_{i=1}^{n} \tilde{d}(x_i, y_i),$$

where the minimum is taken over all $2n$-tuples $(x_1, y_1, x_2, y_2, \dots, x_n, y_n)$ of points from $\Omega(G)$, such that points x_i and y_i belong to the same closed interval B_e, for all $i = 1, \dots, n$, so the distance $\tilde{d}(x_i, y_i)$ is taken in this interval, and furthermore $q_G(x_1) = x$, $q_G(y_n) = y$, and $q_G(y_i) = q_G(x_{i+1})$ for all $i = 1, \dots, n-1$. From now on, whenever G is a metric graph, we think of $\Delta(G)$ as a metric space with the standard metric which we just described.

Given a graph homomorphism from G to H, we can adjust the induced diagram map from the gluing data of G to the gluing data of H to the metric setting, by taking the map from B_e to $B_{\varphi_E(e)}$ to be the dilation with the scaling factor $l_H(\varphi_E(e))/l_G(e)$.

In the colimit we get the induced map from $\Delta(G)$ to $\Delta(H)$, which we also denote by φ_Δ.

Definition 2.5. *Let G be a metric graph, and let n be a nonnegative integer. We say that G is a* **metric graph with n marked points***, when we are given a function $p_G : [n] \to \Delta(G)$ called the* **marking function**.

Here we use the convention $[n] := \{1, \dots, n\}$ for natural numbers n, and $[0] := \emptyset$. Formally, metric graphs with n marked points are given by triples (G, l_G, p_G). Clearly, these generalize metric graphs, which we can recover by setting $n := 0$. For $x \in \Delta(G)$, we say that *x is marked with $p_G^{-1}(x)$*, or simply that *x is marked*, in case that subset of $[n]$ is not empty. We call a point $x \in \Delta(G)$ *special* if it is vertex or a marked point (or both).

Definition 2.6. *Two metric graphs G and H with n marked points are said to be* **isometric** *if there exists a graph isomorphism consisting of the maps $\varphi_V : V(G) \to V(H)$ and $\varphi_E : DE(G) \to DE(H)$, such that we have $l_G = l_H \circ \varphi_E$, and the marked points are mapped appropriately by the corresponding isometries of the edges, i.e., $p_H = \varphi_\Delta \circ p_G$.*

A graph isomorphism $\varphi = (\varphi_V, \varphi_E)$ from G to H, which is also an isometry of metric graphs, induces an isometry φ_Δ of the corresponding metric spaces. Isometry of metric graphs with n marked points is clearly an equivalence relation. When G is a metric graph with n marked points we let $[G]$ denote the corresponding equivalence class (that is the set of all metric graphs with n marked points which are isometric to G). Likewise, for a set S of metric graphs with n marked points we set $[S] := \{[G] \mid G \in S\}$.

2.4. Shrinking edges. Given a metric graph G with n marked points and an edge $e \in E(G)$, which is not a loop, we can define a new metric graph $H = G/e$ as follows. Let $\{v, w\}$ be the set of the endpoints of e, $v \neq w$, and let z be a label which is not in $V(G)$; for reasons which will become clear shortly, we let the label z be the set $\{v, w\}$ itself. Let furthermore $e_1, e_2 \in DE(G)$ be the directed edges corresponding to e. We set $DE(H) := DE(G) \setminus \{e_1, e_2\}$ and $V(H) := (V(G) \setminus \{v, w\}) \cup \{z\}$. In particular, we see that $E(H) = E(G) \setminus \{e\}$. Let $c : V(G) \to V(H)$ be the map defined by $c(v) := c(w) := z$, and $c(x) := x$ for $x \neq v, w$. We now set $\mathfrak{op}_H := \mathfrak{op}_G|_{DE(H)}$ and $\partial_H^\bullet := c \circ \partial_G^\bullet$. In concrete terms, we have

$$\begin{aligned}
DE_H(x, y) &:= DE_G(x, y), \text{ if } z \notin \{x, y\};\\
DE_H(x, z) &:= DE_G(x, v) \cup DE_G(x, w), \text{ for } x \neq z;\\
DE_H(z, x) &:= DE_G(v, x) \cup DE_G(w, x), \text{ for } x \neq z;\\
DE_H(z, z) &:= DE_G(v, v) \cup DE_G(w, w) \cup DE_G(v, w) \cup DE_G(w, v) \setminus \{e_1, e_2\}.
\end{aligned}$$

The function l_H is set to be the restriction of l_G to $E(H)$. Clearly, we have a surjective map $\sigma : \Delta(G) \to \Delta(H)$ which "shrinks" the edge e, and the marking function p_H is taken to be the composition $\sigma \circ p_G$.

More generally, let $S \subseteq E(G)$ be a set of edges which forms a subforest of G, i.e., the induced graph *contains no cycles* (in particular, the set S contains no loops), and let $DS \subseteq DE(G)$ be the set of corresponding directed edges. One can then shrink the set S just like we shrunk a single edge. More precisely, let Σ be the graph whose set of vertices is $V(G)$ and whose set of edges is S. The new graph $H = G/S$ is now obtained by taking the connected components of Σ as vertices, and setting $DE(H) := DE(G) \setminus DS$, hence $E(H) = E(G) \setminus S$. Let $c : V(G) \to V(H)$ be the map taking every vertex of G to the connected component of Σ which contains it. We can then define $\mathfrak{op}_H := \mathfrak{op}_G|_{DE(H)}$ and $\partial^\bullet_H := c \circ \partial^\bullet_G$; just like for the case when S consists of a single edge. For $C, D \in V(H)$ we now have

$$\begin{aligned} DE_H(C,D) &:= \bigcup_{\substack{x\in V(C)\\ y\in V(D)}} DE_G(x,y), \text{ for } C \neq D; \\ DE_H(C,C) &:= \Big(\bigcup_{x,y\in V(C)} DE_G(x,y)\Big) \setminus DS. \end{aligned}$$

We again let l_H be the restriction of l_G to $E(G) \setminus S$, we have a surjective shrinking map of topological spaces $\sigma : \Delta(G) \to \Delta(H)$, and we set $p_H := \sigma \circ p_G$.

We want to point out a subtlety related to the edge shrinking. Given two disjoint sets of edges S_1 and S_2 we could shrink all these edges provided that $S_1 \cup S_2$ forms a subforest. If we shrink first S_1 and then S_2 we get a different graph from the one obtained by shrinking the set $S_1 \cup S_2$ right away. This is because the labels of the vertices will be different. For example, for a graph with 3 vertices and 2 edges G given by $V(G) := \{a, b\}$, $E_G(a, b) := \{e\}$, $E_G(b, c) := \{f\}$, shrinking first e and then f yields a graph with a single vertex labelled $\{\{a, b\}, c\}$ and no edges, while shrinking the entire set $\{e, f\}$ right away yields a graph with a single vertex labelled $\{a, b, c\}$ and no edges. However, it is easy to see that the isometry class $[G/S]$ of the obtained graph does not depend on the order in which we do the shrinking. We will implicitly use this fact in the future arguments.

3. The spaces MG_n and TM_n, and their deformation retracts.

3.1. The moduli space of metric graphs with n marked points and its modifications. Let n be a nonnegative integer, and let MG_n denote the set of all isometry classes of finite metric graphs with n marked points. We would like to turn this set into a topological space. For this we need to say when two isometry classes of metric graphs with n marked points "are close."

Let G be a metric graph with n marked points. We set $r(G) := \min d(x, y)$, where the minimum is taken over all pairs of special points $x, y \in \Delta(G)$. Note that since the number of special points is necessarily finite, the minimum is well-defined. We shall refer to the open interval $(0, r(G)/2)$ as the *admissible range* of G, this is the range from which the sizes of the neighborhoods of $[G]$ in MG_n shall be sampled, and depends only on the isometry class $[G]$, not on the choice of the representative G.

Let ε be a number from the admissible range of G. We now define a set $N_\varepsilon(G)$ as follows: a metric graph with n marked points H lies in $N_\varepsilon(G)$ if and only if

(1) the edges of H of length less than ε form a subforest;
(2) the metric graph G can be obtained from H by shrinking all the edges of lengths less than ε, as described in Subsection 2.4, and by subsequently varying lengths of remaining edges and positions of marked points by up to ε.

The latter can be formalized as follows. We say that a metric graph G with n marked points can be obtained from another metric graph H with n marked points by varying lengths of edges and the positions of marked points by up to ε if there exists a graph isomorphism φ from G to H such that

(1) for all $e \in E(G)$ we have $|l_G(e) - l_H(\varphi(e))| < \varepsilon$;
(2) for all $i \in [n]$, we have $d(\varphi_\Delta(p_G(i)), p_H(i)) < \varepsilon$.

Finally, we set $N_\varepsilon([G]) := [N_\varepsilon(G)] \subseteq MG_n$. Choosing different representatives of $[G]$ means simply changing labels of vertices and edges, therefore the set $N_\varepsilon([G])$ is independent on the particular choice of G.

It is easy to see that when ε_1 lies in the admissible range of G, and $\varepsilon_2 < \varepsilon_1$, we have $N_{\varepsilon_2}([G]) \subseteq N_{\varepsilon_1}([G])$. This is because for any $H \in N_{\varepsilon_2}(G)$, and for any edge e of H we cannot have $\varepsilon_2 \leq l_H(e) < \varepsilon_1$, as otherwise e would correspond to an edge of G, such that $|l_G(e) - l_H(e)| < \varepsilon_2$, which is impossible since the inequalities $l_G(e) \geq r(G)$ and $l_H(e) < \varepsilon_1 < r(G)/2$ imply

$$l_G(e) - l_H(e) > r(G) - r(G)/2 = r(G)/2 > \varepsilon_1 \geq \varepsilon_2.$$

We are now ready to topologize the set of all isometry classes of metric graphs with n marked points.

Definition 3.1. *Let n be a nonnegative integer. The* **moduli space of metric graphs with n marked points** *is the topological space whose set of points is given by MG_n, and whose topology is generated by the sets $N_\varepsilon([G])$ as follows: a subset $X \subseteq MG_n$ is open if and only if for every $[G] \in X$ there exists $\varepsilon > 0$, such that $N_\varepsilon([G]) \subseteq X$.*

We leave to the reader the verification of the fact that the spaces $N_\varepsilon([G])$ are themselves open. We extend the usage of MG_n to denote the corresponding topological space as well.

There are various natural modifications of MG_n. For example, one could require the metric graphs to be connected. We denote the corresponding subspace of MG_n by $MG_n^{\mathfrak{c}}$. Another, independent possibility is to require that the marked points are vertices of the graph. We denote the corresponding subspace of MG_n by $MG_n^{\mathfrak{v}}$. Combining, we let $MG_n^{\mathfrak{cv}} = MG_n^{\mathfrak{c}} \cap MG_n^{\mathfrak{v}}$ denote the subspace of MG_n consisting of the isometry classes of connected metric graphs with n marks on vertices (as we constructed it, multiple marks are allowed).

Given a metric graph G with n marked points, we let $G^{\mathfrak{v}}$ denote the metric graph obtained from G by turning all marked points into vertices (of course, in case they were not vertices already). Clearly, the isometry class $[G^{\mathfrak{v}}]$ depends on the isometry class $[G]$ only, hence the map $\varphi : MG_n \to MG_n^{\mathfrak{v}}$ given by $\varphi : [G] \to [G^{\mathfrak{v}}]$ is well-defined.

Proposition 3.2. *The map $\varphi : MG_n \to MG_n^{\mathfrak{v}}$ making all marked points into vertices is a retraction.*

Proof. By construction, we have $\varphi|_{MG_n^{\mathfrak{v}}} = \mathrm{id}_{MG_n^{\mathfrak{v}}}$. Furthermore, we have to see that the map φ is continuous. To see this, we show that for any metric graph G with n marked points, the neighborhood $N_{\varepsilon/2}([G])$ is mapped inside the neighborhood $N_\varepsilon([G^{\mathfrak{v}}])$, when ε is in the admissible range of G. Note that we can deform G, while staying inside the neighborhood, by the following operations: edge contractions of edges shorter than $\varepsilon/2$, changing edge lengths by up to $\varepsilon/2$, shifting marked points by at most $\varepsilon/2$. Each of these operations can be realized by edge contractions of edges shorter than ε, and changing edge lengths by at most ε in the graph $G^{\mathfrak{v}}$. We therefore conclude that the map φ is a retraction. □

3.2. Connected components of MG_n. Let us now describe the connected components of MG_n. Let G be a metric graph with n marked points, and let $G_1, \dots, G_t$ be its connected components. Assume that G_i has genus g_i, for $i = 1, \dots, t$, and let $A_1, \dots, A_t$ be disjoint, possibly empty sets whose union is $[n]$ (this is like a set partition, but with empty sets allowed). The set $\{(g_1, A_1), \dots, (g_t, A_t)\}$ is now the data which we associate to the graph G, where the sets A_i are the sets of marked points on each connected component G_i. It is not difficult to see that any graph with the same data lies in the connected component which contains $[G]$. Furthermore, the data of this type, meaning a set of tuples (g_i, A_i), such that $g_i \geq 0$ and $(A_1, \dots, A_t)$ is a set partition of $[n]$, possibly involving empty sets, index connected components of MG_n.

Consider now a connected component C_S indexed by the set

$$S = \{(g_1, A_1), \dots, (g_m, A_m), (0, \emptyset), \dots, (0, \emptyset), (1, \emptyset), \dots, (1, \emptyset), \dots\},$$

where the sets A_i are non-empty. Denote the number of appearances of the tuple $(k, \emptyset)$ in that set by n_k, for $k = 0, 1, \dots$. By our assumptions, only finitely many of these are different from 0. Then we have a homeomorphism

$$C_S \cong C_{\{(g_1,A_1)\}} \times \cdots \times C_{\{(g_1,A_1)\}} \times \mathrm{SP}^{n_0}(C_{\{(0,\emptyset)\}}) \times \cdots \times \mathrm{SP}^{n_i}(C_{\{(i,\emptyset)\}}) \times \dots,$$

where $\mathrm{SP}^t(X)$ denotes the t-fold symmetric product, i.e., the quotient space $(\underbrace{X \times \cdots \times X}_{t})/\mathcal{S}_t$, where the symmetric group $\mathcal{S}_t$ acts on the direct product by permutation of its factors.

We shall use $MG^{\mathfrak{c}}_{g,n}$ to denote the topological space whose points are the isomorphism classes of connected graphs of genus g with n marked points, which is the same as the connected component $C_{(g,[n])}$.

3.3. A stratification of MG_n. The moduli space MG_n has a natural stratification. To produce a stratum, fix a graph G, and for each $i \in [n]$, fix w_i, which is either a vertex or an edge of G. Now consider the set of all isometry classes of metric graphs with n marked points, which have a representative (H, l_H, p_H), such that there exists a graph isomorphism φ from G to H, for which the point $p_H(i)$ is equal to $\varphi(w_i) \in \Delta(H)$, if w_i is a vertex, or belongs the open edge $\varphi_{x,y}(w_i) \subset \Delta(H)$, if w_i is an edge with endpoints x and y, for all $i \in [n]$.

This stratum is indexed by the graph G together with the n-tuple $w = (w_1, \dots, w_n)$; we denote it by $\Sigma_{G,w}$. We shall call this stratification *standard*, and we shall its strata the *standard strata*. The stratum does not change if we replace G with an isomorphic graph, and change the n-tuple w accordingly. We shall implicitly use this fact in our discussion.

We let $\overline{\Sigma}_{G,w}$ denote the closure of the stratum $\Sigma_{G,w}$, and we let $\partial\Sigma_{G,w}$ denote $\overline{\Sigma}_{G,w} \setminus \Sigma_{G,w}$, which we shall call the *boundary* of the stratum. One can see that the boundary of an arbitrary stratum $\Sigma_{G,w}$ is a union of other strata. The indexing data of these strata can be obtained from (G, w) by a combination of the steps of the following two kinds

(1) replacing an edge by one of its endpoints in the n-tuple w;
(2) shrinking a non-loop edge in G and replacing this edge and its endpoints by the label of the thus obtained vertex in the n-tuple w.

In general the strata do not have to be manifolds. Consider, for example, the stratum $\Sigma_{G,w}$, where G is a graph with one vertex v and one edge e (which hence must be a loop), $n = 1$, and $w_1 = e$. We can vary the length of the edge in the open interval $(0, \infty)$, and we can slide the marked point along the edge. Thus potentially the point moves in the interval $(0, l_G(e))$, however, because of the symmetry which flips the loop, we need to identify coordinates x and $l_G(e) - x$, so we can choose a representative from the half-closed interval $(0, l_G(x)/2]$. These considerations show that the stratum $\Sigma_{G,w}$ is homeomorphic to the space $\{(l, x) \in \mathbb{R}^2 \mid l > 0,\ l/2 \geq x > 0\}$, which in turn is homeomorphic to the space $\{(x, y) \in \mathbb{R}^2 \mid x > 0,\ y \geq 0\}$.

On the positive side, as easily seen, the generic points of every stratum $\Sigma_{G,w}$ form an open manifold whose dimension is equal to the number of edges of G plus the number of labels which are edges in the n-tuple w. It follows, that there are infinitely many strata of dimension 0; these are indexed by graphs G with no edges, whose vertices are labeled by disjoint subsets of $[n]$, so that the union of all labels is $[n]$, which is the same as to index them by sets $\{A_1, \dots, A_t\}$, where we might have $A_i = \emptyset$, such that $[n] = \cup_{i=1}^t A_i$, and the union is disjoint.

In general, the space MG_n is somewhat technical to handle directly: it is infinite dimensional, and an arbitrarily small neighborhood of each point intersects infinitely many strata; for example, in the case with no marked points, an arbitrarily small neighborhood of the graph with one vertex and no edges intersects all strata indexed by trees. We shall therefore start by performing the shrinking bridges strong deformation retraction, in order to replace MG_n by a more manageable space.

3.4. The shrinking bridges strong deformation retraction. Recall, that an edge e of a graph G is called a *bridge* if deleting it from the graph G increases the number of connected components. Equivalently, e is a bridge if the endpoints of e belong to different connected components of $G - e$ (cf. [Di05, p. 11]). Clearly, shrinking a bridge will neither change the number of connected components of G, nor will it change the genera of these connected components. For example, every edge of a forest is a bridge. We shall call a graph which has no bridges *bridge-free*, and we shall denote the set of bridges of G by $\mathrm{Br}(G)$.

Let us now define a homotopy $\beta : MG_n \times [0,1] \to MG_n$. Let G be a metric graph with n marked points, and let $t \in (0,1]$. Let $\beta(G,t)$ denote the metric graph with n marked points obtained from G by scaling down all the edges in $\mathrm{Br}(G)$ by the factor t, and adjusting the marking function accordingly. For $t = 0$ we set $\beta(G,0)$ to be the graph obtained from G by shrinking all the bridges, this is allowed since the set of all bridges forms a subforest of G. Up to isomorphism, the metric graph $\beta(G,t)$ with n marked points is uniquely determined by the isomorphism class of G, and by the parameter t, hence the assignment $\beta([G],t) := [\beta(G,t)]$ is well-defined.

Let $MG_n^{\mathfrak{b}}$ denote the subspace of MG_n consisting of all the isometry classes of metric bridge-free graphs G with n marked points. Clearly, the space $MG_n^{\mathfrak{b}}$ is a union of standard strata, and the map $\beta(-,0)$ takes MG_n to $MG_n^{\mathfrak{b}}$ surjectively.

THEOREM 3.3. *The space $MG_n^{\mathfrak{b}}$ is a strong deformation retract of the space MG_n. The map $\beta : MG_n \times [0,1] \to MG_n$ provides a corresponding strong deformation retraction.*

Proof. As already mentioned, we have $\beta([G],0) \in MG_n^{\mathfrak{b}}$, for any metric graph G with n marked points. Furthermore, it follows directly from our definition of the map β, that $\beta([G],t) = [G]$, for all $[G] \in MG_n^{\mathfrak{b}}$, and that $\beta(-,1) = \mathrm{id}_{MG_n}$. Therefore, to prove that β is an appropriate strong deformation retraction, it is enough to show that it is continuous.

We shall now provide a direct, albeit somewhat tedious verification. We take a point $x \in MG_n$ and a point $y \in MG_n \times [0,1]$, such that $\beta(y) = x$. We then show that for every sufficiently small ε (how small it needs to be shall depend on x and y), there exists a neighborhood N of y, such that $\beta(N) \subseteq N_\varepsilon(x)$. We use the notations $y = ([G],t) \in MG_n \times [0,1]$, where G is a metric graph with n marked points, and accordingly $x = \beta([G],t) = [\beta(G,t)]$. It is technically easier to divide the argument into considering two separate cases.

Case 1. *We assume that $t > 0$.*

This is the easier one of the two cases. For sufficiently small $\varepsilon > 0$, we look for $\delta_1, \delta_2 > 0$, such that β maps the y-neighborhood $N_{\delta_1}([G]) \times [t-\delta_2, t+\delta_2]$ inside of $N_\varepsilon(x)$. Since $N_{\delta_1}([G]) = [N_{\delta_1}(G)]$, and $N_\varepsilon(x) = N_\varepsilon(\beta([G], t)) = N_\varepsilon([\beta(G,t)]) = [N_\varepsilon(\beta(G,t))]$, it is enough to find ε, δ_1, and δ_2, such that β maps $N_{\delta_1}(G) \times [t-\delta_2, t+\delta_2]$ inside of $N_\varepsilon(\beta(G,t))$, as long as the conditions on ε, δ_1, and δ_2 depend only on the isomorphism class of G, not on the specific representative. In any case, we assume that ε is sampled from the admissible range of G (which only depends on $[G]$).

Let now H be a metric graph with n marked points in $N_{\delta_1}(G)$, and let Σ denote the set of edges of H of length less than δ_1. By our construction, these must form a subforest. Let us fix some $\tilde{t}$ in the interval $[t-\delta_2, t+\delta_2]$. The graph $\beta(H, \tilde{t})$ is obtained from H by shrinking all the bridges by a factor $\tilde{t} \leq 1$, $\tilde{t} \neq 0$. Therefore, requesting that $\delta_1 \leq \varepsilon$ will ensure that the images of the edges from Σ will have length less than ε. On the other hand, we want the images of the edges from $E(H) \setminus \Sigma$, i.e., the original edges of G, to have lengths larger than ε. This can be ensured by requesting that

$$\tilde{t} \cdot \min_{e \in E(G)} l_G(e) > \varepsilon.$$

Therefore, the graph $\beta(G, \tilde{t})$ can be obtained from the graph $\beta(H, \tilde{t})$ by shrinking all the edges of length less than ε, and then varying the lengths of the remaining edges, as well as positions of marked points, by up to an ε. The latter follows from the fact that $\delta_1 \leq \varepsilon$, and shrinking some of the edges only decreases the edge length variation.

Passing on to the whole interval $[t-\delta_2, t+\delta_2]$, we choose ε so that

$$t \cdot \min_{e \in E(G)} l_G(e)/2 > \varepsilon,$$

and then take $\delta_1 \leq \varepsilon$ and $\delta_2 \leq t/2$. This choice of parameters verifies the continuity of the map β at the point y.

Case 2. *We assume that $t = 0$.*
We have $x = \beta([G], 0) \in MG_n^{\mathfrak{b}}$. This time, for sufficiently small $\varepsilon > 0$, we need to find $\delta_1, \delta_2 > 0$, such that β maps the y-neighborhood $N_{\delta_1}([G]) \times [0, \delta_2]$ inside of $N_\varepsilon(x)$. Just like in the first case, we can drop the isomorphism brackets, and search for ε, δ_1, and δ_2, such that β maps $N_{\delta_1}(G) \times [0, \delta_2]$ inside of $N_\varepsilon(\beta(G, 0))$.

Let again H be a metric graph with n marked points in $N_{\delta_1}(G)$, and let Σ denote the set of edges of H of length less than δ_1. It important to note that the set of edges $\Sigma \cup \mathrm{Br}(G)$ forms a subforest. This is because Σ is a forest, and adding bridges to any forest will not create cycles, since bridges cannot be a part of any cycle. Also, as noted before, we have $\mathrm{Br}(G) \subseteq \mathrm{Br}(H) \subseteq \mathrm{Br}(G) \cup \Sigma$.

The graph $\beta(H, 0)$ is obtained from H by shrinking all the bridges. This means shrinking all the bridges of G, and possibly some of the edges from Σ. Choosing ε smaller than $\min_{e \in E(G)} l_G(e)$, and then choosing $\delta_1 < \varepsilon$ ensures that the edges of $\beta(H, 0)$ whose lengths are less than ε are precisely the non-bridges from the set Σ. This means that the graph $\beta(G, 0)$ can be obtained from the graph $\beta(H, 0)$ by shrinking the edges of length less than ε, and varying the lengths of other edges, as well as positions of the marked points, by up to an ε.

Assume now δ_2 is chosen so that

$$\delta_2 \cdot \max_{e \in E(G)} l_G(e) < \varepsilon,$$

and choose $0 < t \leq \delta_2$. The graph $\beta(H,t)$ is obtained from H by scaling all the bridges down by the factor t. The way δ_2 is chosen, this will scale down all the bridges of G, so that they become shorter than ε. Furthermore, the conditions $\delta_1 < \varepsilon < \min_{e \in E(G)} l_G(e)$ imply that also all the edges from Σ will be shorter than ε (since they were shorter than ε to start with, and some were additionally shrunk), and that all the non-bridges of G will not be shorter than ε. Hence again the graph $\beta(G,0)$ can be obtained from the graph $\beta(H,t)$ by shrinking the edges of length less than ε, and varying the lengths of other edges, as well as positions of the marked points, by up to an ε.

This finishes the verification of the fact that the homotopy β is continuous at the point y in this case. □

3.5. Moduli space of tropical curves of genus g with n marked points. We now define a subspace of $MG_n^{\mathfrak{cv}}$ which is of special interest in tropical geometry and has been the starting point of the current investigation.

DEFINITION 3.4. *Let n be a nonnegative integer, and let d be a positive number. We define $TM_n(d)$ to be the subspace of $MG_n^{\mathfrak{cv}}$ consisting of the isomorphism classes of all metric graphs G with n marked points, such that*

(1) G has no vertices of valency 2;

(2) G has exactly n leaves[2], and these are marked 1 through n;

(3) the lengths of the edges leading to leaves are equal to d.

Simultaneous dilation of the edges leading to leaves gives a homeomorphism between spaces $TM_n(d_1)$ and $TM_n(d_2)$, for arbitrary positive d_1 and d_2. In fact, we might as well set the lengths of the edges leading to leaves to be equal ∞. The definitions from Section 2 and Subsection 3.1 generalize to this case in a straightforward manner. We call the obtained space $TM_n(\infty)$, and it is this space which is known as the *moduli space of tropical curves*. Since the spaces $TM_n(\infty)$ and $TM_n(d)$ are homeomorphic for any positive d, we restrict ourselves to considering $TM_n(1)$.

Letting d go to 0 we obtain yet another homeomorphic space, which it would be natural to denote by $TM_n(0)$, but for simplicity we just call it TM_n. This space can also be described directly, as is done in the next definition.

DEFINITION 3.5. *Let n be a nonnegative integer. We define TM_n to be the subspace of $MG_n^{\mathfrak{cv}}$ consisting of the isomorphism classes of all metric graphs G with n marked points, such that for every vertex of G the sum of its valency with the number of times it is marked should be at least 3.*

The condition in Definition 3.5 just means that every vertex of valency 2 should be marked, and that every leaf should be marked at least twice.

PROPOSITION 3.6. *For any nonnegative integer n, the spaces TM_n and $TM_n(1)$ are homeomorphic.*

Proof. A homeomorphism $\varphi : TM_n(1) \to TM_n$ is given by the simultaneous shrinking of all the edges leading to leaves, and shifting the labels of the leaves onto the internal vertices accordingly.

The inverse of φ is given by taking any graph G with $[G] \in TM_n$, and replacing every label l of a vertex v in G with a new vertex called v_l and labeled with l, and an edge connecting v with v_l. Since this operation "converts" labels on vertices to the

[2]Generalizing the terminology customary for trees, we use the word *leaves* to denote any vertex of valency 1, cf. [Di05, p. 13].

corresponding increase in their valencies, we see that the condition in $TM_n(1)$, that all non-leaves must have valency ≥ 3, corresponds precisely to the condition that for a graph G, such that $[G] \in TM_n$, and a vertex v of G, the sum of the valency of v with the cardinality of its label is ≥ 3. See Figure 3.1.

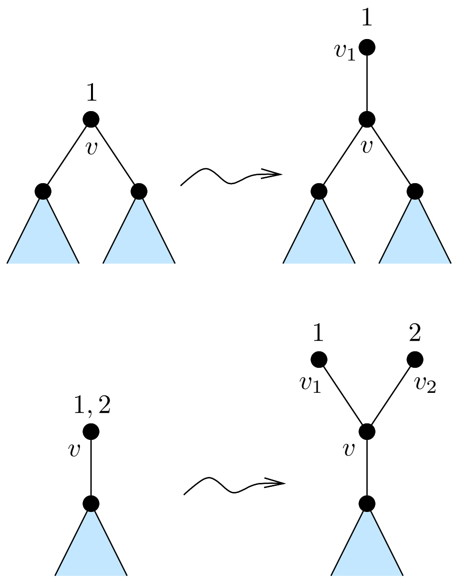

FIG. 3.1. *Examples of the inverse of φ.*

It is a straightforward verification that for the defined topologies, both the map φ and its inverse are continuous, implying that we have a homeomorphism. $\square$

It is not a difficult exercise to see that every two points of TM_n corresponding to metric graphs of the *same* genus, can be connected by a path inside TM_n, whereas obviously, every two points of TM_n corresponding to metric graphs of *different* genus, cannot be connected by such a path, not even inside of MG_n. Hence the connected components of TM_n are indexed by nonnegative integers g, corresponding to the genera of the involved graphs, and we call them $TM_{g,n}$.[3]

It is important to note that, unlike the spaces $TM_n(d)$, the space TM_n is a union of the standard strata of MG_n.

COROLLARY 3.7. *Let $TM_n^{\mathfrak{b}} := TM_n \cap MG_n^{\mathfrak{b}}$. Then the space $TM_n^{\mathfrak{b}}$ is a strong deformation retract of TM_n.*

Proof. The space TM_n is obviously closed under the shrinking bridges strong deformation retraction. Therefore, the Theorem 3.3 implies that TM_n strongly deformation retracts to $TM_n^{\mathfrak{b}}$. $\square$

The space $TM_n^{\mathfrak{b}}$ can be also described directly: it is the subspace of $MG_n^{\mathfrak{cv}}$ consisting of all isomorphism classes of metric graphs G with n marked points, such that

(1) every vertex of G of valency 2 is marked;

[3] It has been recently proved that the tropical moduli space $TM_{1,n}$ is a strong deformation retract of the graph moduli space $MG_{1,n}^{\mathfrak{v}} := MG_{1,n} \cap MG_n^{\mathfrak{v}}$, see [Ko08].

(2) the graph G has no bridges.

The connected components of $TM_n^{\mathfrak{b}}$ are again indexed by genera of the constituting graphs, and we shall use the notation $TM_{g,n}^{\mathfrak{b}} := TM_{g,n} \cap MG_n^{\mathfrak{b}}$.

4. The moduli space of the tropical curves of genus 1.

4.1. Presentation as a quotient space. Since all the edges of a tree are bridges, we see that $MG_{0,n}^{\mathfrak{b}}$ and $TM_{0,n}^{\mathfrak{b}}$ are just single points. Theorem 3.3 and Corollary 3.7 imply that the spaces $MG_{0,n}$ and $TM_{0,n}$ are contractible. In this section we shall focus on the next interesting case: namely the spaces of connected metric graphs of genus 1 with n marked points.

Let us start by analysing the space $\widetilde{X}_n = TM_{1,n}^{\mathfrak{b}}$. The bridge-free graphs of genus 1 are simply cycles, hence, since all their vertices have valency equal to 2, they should all be marked, and n should be at least 1. Reversely, all the marked points are vertices. We can therefore forget about the vertices and just record the marked points.

The points of the space $\widetilde{X}_n$ can thus be indexed by n-tuples of points on a circle, whose radius is an arbitrary positive real number, divided by the action of the orthogonal group on the circle. Factoring out the radius length of the circle, we have a homeomorphism

$$\widetilde{X}_n \cong X_n \times (0, \infty),$$

for $X_n := (\underbrace{S^1 \times \cdots \times S^1}_{n})/O(2)$, where S^1 denotes the unit circle, and the division is done with respect to the diagonal action of the orthogonal group $O(2)$, which acts on each term in the natural way.

Note, that the described action of $O(2)$ on the direct product $\underbrace{S^1 \times \cdots \times S^1}_{n}$ is transitive on the last coordinate. Let us fix the last coordinate to be $(1, 0) \in S^1$. Then the space X_n can be rewritten as

$$X_n \cong \underbrace{S^1 \times \cdots \times S^1}_{n-1}/\mathbb{Z}_2, \tag{4.1}$$

where $\mathbb{Z}_2$ is the subgroup of $O(2)$ which fixes the point $(1, 0)$. Clearly, the group $\mathbb{Z}_2$ consists of two elements, and the non-identity element is an involution which acts diagonally on the direct product of $n-1$ copies of S^1 by a reflection about the X-axis. One way to think of this action is to view the points of S^1 as complex numbers with absolute values equal to 1, in which case the action is just a simultaneous conjugation:

$$X_n \cong \{(z_1, \dots, z_{n-1}) \,|\, |z_i| = 1, \text{ for } i = 1, \dots, n-1\}/(z_1, \dots, z_{n-1}) \sim (\bar{z}_1, \dots, \bar{z}_{n-1}).$$

4.2. A cubical structure on X_n. Consider the CW structure on a unit circle S^1 consisting of two 0-cells $(1, 0)$ and $(-1, 0)$, and two 1-cells corresponding to upper and lower semicircles. This induces a CW structure on the direct product of n copies of S^1 (which is an n-torus), where the cells are simply direct products of cells of the factors. These are indexed by the ordered n-tuples of the cells of the factors and we shall use the following encoding: we write "+" for the 0-cell $(1, 0)$, "−" for the 0-cell $(-1, 0)$, "$i+$" for the upper semicircle, and "$i-$" for the lower semicircle, as shown on Figure 4.1. So, for example, $(i-, -, -, i+)$ would denote a 2-cell in the 4-torus. The open cells are actually cubes, so we can think of this decomposition as some sort of a cubical structure on $\underbrace{S^1 \times \cdots \times S^1}_{n}$. In this encoding, the boundary of a cell is obtained by replacing symbols $i+$ and $i-$ with symbols $+$ and $-$. The number of d-cells is $\binom{n}{d}2^n$, as we can put a plus or a minus in every coordinate, and then distribute d symbols "i" arbitrarily.

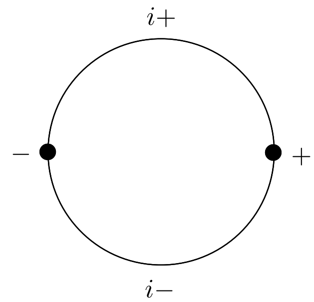

FIG. 4.1. *The cell encoding in Subsection 4.2.*

Clearly, this cubical structure on $\underbrace{S^1 \times \cdots \times S^1}_{n-1}$ is invariant under the $\mathbb{Z}_2$-action of simultaneous conjugation. In our symbolic notations, the action changes the signs assigned to i's and fixes the other coordinates. In particular, all 0-cells are fixed, and all higher dimensional cells come in pairs, the cells in each pair are being swapped by the involution. This means that in the quotient X_n we get the induced cubical structure, consisting of the orbits of the cells of the $(n-1)$-torus. These are also indexed by the n-tuples of the symbols from the set $\{+, -, i+, i-\}$, with an additional constraint that the first symbol $i+$ comes before the first symbol $i-$ (if at all). Hence the number of vertices in this cubical structure on X_n is 2^{n-1}, whereas the number of d-cubes, for $n-1 \geq d \geq 1$, is $\binom{n-1}{d}2^{n-2}$.

4.3. Small values of n. The presentation (4.1) can be used to analyze what happens for the small values of n. The space X_1 is just a point: the empty direct product is to be interpreted as a point here. The space $X_2 \cong S^1/\mathbb{Z}_2$ is homeomorphic to a closed interval, e.g., $[-1, 1]$ if we take the projection of S^1 onto the X-axis.

Let us consider $X_3 \cong (S^1 \times S^1)/\mathbb{Z}_2$. We have a described a cubical structure on this space which has 4 vertices, 4 edges and 2 squares. The two squares are indexed with $(i+, i+)$ and $(i+, i-)$ and it is easy to see that X_3 is obtained by gluing these 2 squares together along their entire boundaries. Hence we conclude that X_3 is homeomorphic to a 2-dimensional sphere.

As the last case, let us consider $X_4 \cong (S^1 \times S^1 \times S^1)/\mathbb{Z}_2$. This is a space glued together from 4 cubes; it has 8 vertices, 12 edges and 12 squares. All of the cubes share the same set of vertices, and it is easy to see that all points, except for these 8 vertices, have neighborhoods which are homeomorphic to open balls in $\mathbb{R}^3$. However, the space X_4 is not a manifold. To see this, let us take a look at the neighborhoods of these vertices. Clearly, it does not matter which one we take, so let us take the vertex $v = (+, +, +)$. As an open neighborhood of v we can take a cone with an apex in v over the link of v. The link of a vertex in a cubical complex is always a simplicial complex. The one we have here has 3 vertices, 6 edges, and 4 triangles. A presentation of this simplicial complex is given on Figure 4.2, from which it is clear that this link is homeomorphic to $\mathbb{RP}^2$. This shows that X_4 fails to be a manifold at these points.

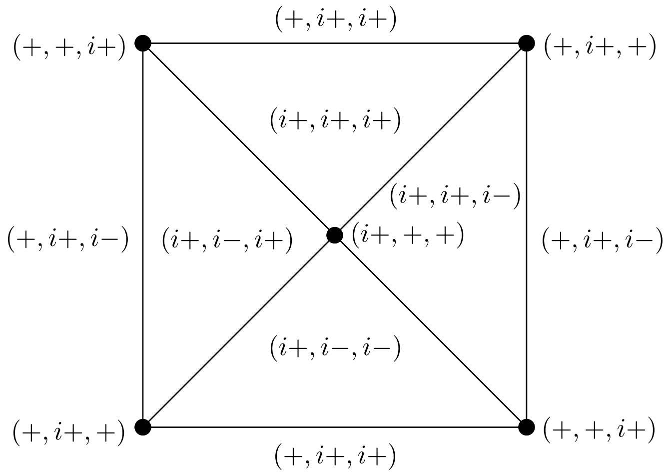

FIG. 4.2. *The link of the vertex* $(+,+,+)$ *in* X_4.

4.4. Surgery presentation. Let us now generalize our description of X_4 to the general case. In particular, we shall see that for $n \geq 4$ the vertices in the cubical structure on X_n are singularities, and if they are removed, we are left with an $(n-1)$-dimensional manifold.

We start with by viewing the n-torus T^n in the standard way as the quotient space of $\mathbb{R}^n$ divided by the group action of $\mathbb{Z}^n = \langle g_1, \dots, g_n \rangle$, given by

$$g_i : (x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n) \mapsto (x_1, \dots, x_{i-1}, x_i + 2, x_{i+1}, \dots, x_n).$$

Since $X_{n+1} \cong T^n/\mathbb{Z}_2$, as described above, we have $X_n \cong \mathbb{R}^n/G_n$, where G_n is a group defined by

$$G_n = \langle \gamma, g_1, \dots, g_n \,|\, g_i \circ \gamma = \gamma \circ g_i^{-1},\ g_i \circ g_j = g_j \circ g_i,\ \text{for } i,j \in [n] \rangle, \tag{4.2}$$

with the action of γ on $\mathbb{R}^n$ given by

$$\gamma : (x_1, \dots, x_i, \dots, x_n) \mapsto (-x_1, \dots, -x_i, \dots, -x_n).$$

It follows from the presentation (4.2) that the group G_n is the semidirect product $\mathbb{Z}^n \times_\phi \mathbb{Z}_2$, with the group homomorphism $\phi : \mathbb{Z}_2 \to \mathrm{Aut}\,(\mathbb{Z}^n)$ given by $\phi(\gamma)(g) = g^{-1}$, where γ is the nontrivial element of $\mathbb{Z}_2$ and $g \in \mathbb{Z}^n$ is arbitrary. In particular, the elements of G_n can be uniquely presented either as $\gamma g_1^{\alpha_1} \dots g_n^{\alpha_n}$ or as $g_1^{\alpha_1} \dots g_n^{\alpha_n}$, for $\alpha_1, \dots, \alpha_n \in \mathbb{Z}$. The action of the element $g_1^{\alpha_1} \dots g_n^{\alpha_n}$ is fixed-point-free, whereas the action of the element $\gamma g_1^{\alpha_1} \dots g_n^{\alpha_n}$ has a unique fixed point, whose coordinates are $(-\alpha_1, \dots, -\alpha_n)$.

We see that the action of G_n on $\mathbb{R}^n \setminus \mathbb{Z}^n$ is free, and that accordingly $(\mathbb{R}^n \setminus \mathbb{Z}^n)/G_n$ is an n-dimensional manifold. For the points in $\mathbb{Z}^n \subset \mathbb{R}^n$ we see that the stabilizers are of cardinality 2, and that given such a point v, the action of its stabilizer on an $(n-1)$-dimensional sphere centered at v is the standard antipodal action, hence the quotient of this sphere by that action is the projective space $\mathbb{RP}^{n-1}$.

We conclude that X_{n+1} can be obtained from an n-dimensional manifold, whose boundary consists of 2^n $(n-1)$-dimensional projective space $\mathbb{RP}^{n-1}$, by attaching 2^n cones - one at each boundary projective space.

4.5. The space X_n as a homotopy colimit. We would like to present the space X_n yet in another fashion. This also yields the opportunity to introduce the terminology for dealing with the case of higher genus. The following concept is an important construction in Combinatorial Algebraic Topology, see [Ko07, Chapter 15] for relevant background.

DEFINITION 4.1. *The* **homotopy colimit***, denoted* $\mathbf{hocolim}\,\mathcal{D}$*, of a diagram* $\mathcal{D}$ *of topological spaces over a triangulated space* Δ*, is the quotient space*

$$\mathbf{hocolim}\,\mathcal{D} = \coprod_{\sigma = v_0 \to \dots \to v_n} (\sigma \times \mathcal{D}(v_0))/\sim,$$

where the disjoint union is taken over all simplices in Δ*. The equivalence relation* $\sim$ *is generated by: for* $\tau \in \partial\sigma$, $\tau = v_0 \to \dots \to \hat{v}_i \to \dots \to v_n$*, let* $i : \tau \hookrightarrow \sigma$ *be the inclusion map, then*

- *for* $i > 0$, $\tau \times \mathcal{D}(v_0)$ *is identified with the subset of* $\sigma \times \mathcal{D}(v_0)$*, by the map induced by* i*;*
- *for* $\tau = v_1 \to \dots \to v_n$*, we have* $i(\alpha) \times x \sim \alpha \times \mathcal{D}(v_0 \to v_1)(x)$*, for any* $\alpha \in \tau$*, and* $x \in \mathcal{D}(v_0)$*.*

Examples of homotopy colimits include the mapping cone and the mapping cylinder.

For a diagram $\mathcal{D}$ over Δ we have a *base projection map*:

$$\mathrm{proj} = p_b : \mathbf{hocolim}\,\mathcal{D} \longrightarrow \Delta,$$

induced by the projections onto the first coordinate $\sigma \times \mathcal{D}(v_0) \to \sigma$, for every simplex $\sigma = v_0 \to \dots \to v_n$.

For $n \geq 1$, the space X_{n+2} can be considered as a homotopy colimit of a diagram $\mathcal{D}$ over an interval $[0,1]$, where the latter is viewed as a triangulated space with two vertices and one edge. Namely, set $\mathcal{D}(0) := \mathcal{D}(1) := X_{n+1} = T^n/\mathbb{Z}_2$, $\mathcal{D}((0,1)) := T^n$, and let both diagram maps $\mathcal{D}((0,1) \to 0)$ and $\mathcal{D}((0,1) \to 1)$ be the quotient maps $q : T^n \to T^n/\mathbb{Z}_2$. See Figure 4.3.

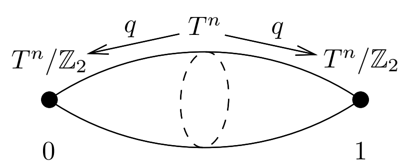

FIG. 4.3. *The space X_{n+2} presented as a homotopy colimit.*

We see that the base projection map $X_{n+2} \to [0,1]$ is induced by the projection of the first coordinate of T^{n+1} to the real axis:

$$(z_1, \dots, z_{n+1}) \mapsto \mathrm{Re}\, z_1.$$

By definition, this homotopy colimit is homeomorphic to the space obtained by taking the cylinder with the base T^n and then taking the quotient on both ends using the map q. The reader is welcome to compare this to the definition of Whitehead group, see [Co73].

4.6. The homology groups of X_n with coefficients in $\mathbb{Z}_2$. It follows from the presentation of X_{n+2} as a homotopy colimit that it can be seen as a union of two disjoint copies of a mapping cylinder of the map q, glued together along the topmost copies of T^n.

Let us slightly modify this picture. Let $b : X_{n+2} \to [0,1]$ be the base projection map corresponding to this diagram. We set $A := b^{-1}([0,1/2+\varepsilon))$ and $B := b^{-1}((1/2-\varepsilon,1])$, where ε is a small positive number, say $\varepsilon = 0.1$. Then we have $A \cap B = b^{-1}((1/2-\varepsilon, 1/2+\varepsilon))$. Since $A \cup B = X_{n+2}$, and the subspaces A, B, and $A \cap B$ are open in X_{n+2}, we can use the Mayer-Vietoris sequence to compute the homology groups of X_{n+2}:

$$\ldots \to \widetilde{H}_i(A \cap B) \to \widetilde{H}_i(A) \oplus \widetilde{H}_i(B) \to \widetilde{H}_i(X_{n+2}) \to \widetilde{H}_{i-1}(A \cap B) \to \ldots. \tag{4.3}$$

Clearly, both A and B are homotopy equivalent to $T^n/\mathbb{Z}_2$, whereas $A \cap B$ is homotopy equivalent to T^n. The inclusion maps $A \cap B \hookrightarrow A$, and $A \cap B \hookrightarrow B$ induce the same maps on the homology as the quotient map $q : T^n \to T^n/\mathbb{Z}_2$. It follows that the Mayer-Vietoris sequence (4.3) translates to the long exact sequence

$$\ldots \to \widetilde{H}_i(T^n) \xrightarrow{(q_*,q_*)} \widetilde{H}_i(T^n/\mathbb{Z}_2) \oplus \widetilde{H}_i(T^n/\mathbb{Z}_2) \to \widetilde{H}_i(X_{n+2}) \to \widetilde{H}_{i-1}(T^n) \xrightarrow{(q_*,q_*)} \ldots. \tag{4.4}$$

This long exact sequence splits into short exact sequences as follows from the next proposition.

Proposition 4.2. *For arbitrary $n \geq i \geq 1$, the induced map $q_* : \widetilde{H}_i(T^n;\mathbb{Z}_2) \to \widetilde{H}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2)$ is a 0-map.*

Proof. We know that $\widetilde{\beta}_i(T^n;\mathbb{Z}_2) = \binom{n}{i}$. Choose a subset S of $[n]$, such that $|S| = i$. Let σ_S to be the sum of all i-cells, indexed by n-tuples $(a_1,\ldots,a_n)$, such that

$$a_j = \begin{cases} +, & \text{for } j \notin S, \\ i+ \ \text{ or } i-, & \text{for } j \in S, \end{cases}$$

for all $j = 1,\ldots,n$. There are 2^i cells like that, and their sum is a cycle. If we let S run over all cardinality i subsets of $[n]$, we shall obtain precisely the representatives of $\binom{n}{i}$ generators of the group $\widetilde{H}_i(T^n;\mathbb{Z}_2)$.

All the cycles σ_S are $\mathbb{Z}_2$-invariant, with all the i-cells in each sum coming in $\mathbb{Z}_2$-invariant pairs. In every pair, both cells are mapped to the same cell in $T^n/\mathbb{Z}_2$, hence their contributions under the map q_* cancel out each other, when we work with $\mathbb{Z}_2$-coefficients. It follows that $q_*(\sigma_S) = 0$ already on the level of chains. Since the σ_S's generate $\widetilde{H}_i(T^n;\mathbb{Z}_2)$, we conclude that q_* is a 0-map on the homology groups. □

Certainly, if q_* is a 0-map, then so is the map (q_*,q_*) in the long exact sequence (4.4). Hence, the latter splits into the short ones of the type

$$0 \to \widetilde{H}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2) \oplus \widetilde{H}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2) \to \widetilde{H}_i(X_{n+2};\mathbb{Z}_2) \to \widetilde{H}_{i-1}(T^n;\mathbb{Z}_2) \to 0.$$

Since we are working over field coefficients we conclude that we have an isomorphism

$$\widetilde{H}_i(X_{n+2};\mathbb{Z}_2) \cong \widetilde{H}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2) \oplus \widetilde{H}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2) \oplus \widetilde{H}_{i-1}(T^n;\mathbb{Z}_2),$$

for all $i \geq 1$ and $n \geq 1$. Accordingly, we get

$$\widetilde{\beta}_i(T^{n+1}/\mathbb{Z}_2;\mathbb{Z}_2) = 2\widetilde{\beta}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2) + \widetilde{\beta}_{i-1}(T^n;\mathbb{Z}_2). \tag{4.5}$$

Using that formula we arrive at the following statement.

THEOREM 4.3. *The Betti numbers over $\mathbb{Z}_2$ of the space X_{n+2} are given by the following formula:*

$$\widetilde{\beta}_i(X_{n+2};\mathbb{Z}_2) = \begin{cases} \sum_{j=0}^{n-i+1} 2^j \binom{n-j}{i-1}, & \text{if } n+1 \geq i \geq 2; \\ 0, & \text{otherwise.} \end{cases} \tag{4.6}$$

Proof. Since $\widetilde{\beta}_0(T^n;\mathbb{Z}_2) = 0$, the equation (4.5) implies $\widetilde{\beta}_1(T^{n+1}/\mathbb{Z}_2;\mathbb{Z}_2) = 2\widetilde{\beta}_1(T^n/\mathbb{Z}_2;\mathbb{Z}_2)$, for $n \geq 1$. Hence $\widetilde{\beta}_1(T^{n+1}/\mathbb{Z}_2;\mathbb{Z}_2) = 2^n\widetilde{\beta}_1(T^1/\mathbb{Z}_2;\mathbb{Z}_2) = 0$.

Assuming now that $i \geq 2$ we can unfold the recursive formula (4.5) as follows:

$$\begin{aligned}
\widetilde{\beta}_i(T^{n+1}/\mathbb{Z}_2;\mathbb{Z}_2) &= 2\widetilde{\beta}_i(T^n/\mathbb{Z}_2;\mathbb{Z}_2) + \binom{n}{i-1} = \\
&= 2\left(2\widetilde{\beta}_i(T^{n-1}/\mathbb{Z}_2;\mathbb{Z}_2) + \binom{n-1}{i-1}\right) + \binom{n}{i-1} = \\
&= 4\widetilde{\beta}_i(T^{n-1}/\mathbb{Z}_2;\mathbb{Z}_2) + 2\binom{n-1}{i-1} + \binom{n}{i-1} = \\
&= \cdots = \binom{n}{i-1} + 2\binom{n-1}{i-1} + 4\binom{n-2}{i-1} + \cdots + 2^{n-i+1}\binom{i-1}{i-1}.
\end{aligned}$$

$\square$

For the integer coefficients the situation is somewhat more complicated, as the map q_* is not necessarily a 0-map anymore. Based on the examples for the small values of n and the general intuition for what this map is, we make the following conjecture.

CONJECTURE 4.4. *For $n \geq 1$ we have*

$$\widetilde{H}_{2i}(T^n/\mathbb{Z}_2;\mathbb{Z}) = \mathbb{Z}_2^{a(i,n)} \oplus \mathbb{Z}^{b(i,n)}, \text{ for } n \geq 2i \geq 2, \tag{4.7}$$

where $a(i,n) = \widetilde{\beta}_{2i+1}(T^n/\mathbb{Z}_2;\mathbb{Z}_2)$, and $a(i,n) + b(i,n) = \widetilde{\beta}_{2i}(T^n/\mathbb{Z}_2;\mathbb{Z}_2)$; whereas $\widetilde{H}_j(T^n/\mathbb{Z}_2;\mathbb{Z}) = 0$ for all other values of j.

Acknowledgments. The author would like to thank Bernd Sturmfels for introducing him into the world of tropical geometry and Grisha Mikhalkin for explaining the importance of the tropical moduli spaces. Furthermore, he thanks Karen Vogtmann for pointing to connections with outer spaces, see [CV86, Vo06]. He would also like to thank Alek Vainshtein and Alex Suciu for discussions of torus quotients, as well as Eva-Maria Feichtner for debating various formal approaches to the topology of the space of metric graphs. Finally, the author expresses his gratitude to the Mathematical Institute at Oberwolfach for the hospitality during the time a part of this research was done.

REFERENCES

[BHV01] L. BILLERA, S. HOLMES AND K. VOGTMANN, *Geometry of the space of phylogenetic trees*, Advances in Applied Math., 27 (2001), pp. 733–767.

[Co73] M. COHEN, *A course in simple-homotopy theory*, Graduate Texts in Mathematics, 10, Springer-Verlag, New York-Berlin, 1973.

[CV86] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math., 84:1 (1986), pp. 91–119.

[Di05] R. Diestel, *Graph theory*, Third edition, Graduate Texts in Mathematics, 173, Springer-Verlag, Berlin, 2005, xvi+411 pp.

[GKM07] A. Gathmann, M. Kerber and H. Markwig, *Tropical fans and moduli space of rational tropical curves*, Compositio Mathematica, to appear. `arXiv:0708.2268`

[Ko08] D.N. Kozlov, *Moduli spaces of metric graphs of genus 1 with marks on vertices*, Topology Appl., 156:2 (2008), pp. 433–437.

[Ko07] D.N. Kozlov, *Combinatorial Algebraic Topology*, Algorithms and Computation in Mathematics, 21, Springer-Verlag Berlin Heidelberg, 2008, xx+390 pp. 115 illus.

[Ko09] D.N. Kozlov, *Moduli spaces of tropical curves of higher genus with marked points and homotopy colimits*, to appear in Israel J. Math.

[McL98] S. Mac Lane, *Categories for the Working Mathematician*, Second edition, Graduate Texts in Mathematics, 5, Springer-Verlag, New York, 1998.

[Mi07] G. Mikhalkin, *Moduli spaces of rational tropical curves*, preprint. `arXiv:0704.0839`

[Mi06] G. Mikhalkin, *Tropical geometry and its applications*, International Congress of Mathematicians, Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 827–852.

[RST05] J. Richter-Gebert, B. Sturmfels and T. Theobald, *First steps in tropical geometry*, Idempotent mathematics and mathematical physics, Contemp. Math., 377, Amer. Math. Soc., Providence, RI, 2005, pp. 289–317.

[SS04] D. Speyer and B. Sturmfels, *The tropical Grassmannian*, Adv. Geom., 4:3 (2004), pp. 389–411.

[Vo06] K. Vogtmann, *The cohomology of automorphism groups of free groups*, International Congress of Mathematicians, Vol. II, pp. 1101–1117, Eur. Math. Soc., Zürich, 2006.