A proof of the (1,2g)-self-duality of branchwidth*

Georgios Kontogeorgiou† Dimitrios M. Thilikos†‡§
Dimitris Zoros†

January 11, 2020

Abstract

We prove that the branchwidth of a 2-edge-connected hypergraph embedded on an oriented 2-manifold of genus g is an $(1, 2g)$-self-dual parameter, partially resolving a conjecture by Sau and Thilikos.

1 Background

In [5], Sau and Thilikos proved that branchwidth is a $(6, 2eg - 4)$-self-dual parameter over any surface S of Euler genus eg. They conjectured that this result is not tight, and that instead $(1, eg)$ is optimal. In this paper, we provide a proof that branchwidth is $(1, 2g)$-self-dual over *oriented* surfaces of genus g, thus proving the Sau-Thilikos conjecture in that particular case. This is the first proof that branchwidth is an additively self-dual width parameter.

2 Definitions and preliminary results

In this paper we consider *multigraphs*, i.e., graphs that may contain multiple edges and/or loops, and, in the second part of this paper, *hypergraphs*, i.e., graphs in which an edge can join any number of vertices instead of only two vertices. We call such edges *hyperedges* (or just edges) .

A *surface* is a connected compact 2-manifold without boundaries. We denote the *genus* of the surface (i.e., the "number of handles" present in the surface) by g.

An *open arc* A on a surface S is a subset of S homeomorphic to $(0, 1)$. The set of its endpoints is defined to be the set $\bar{A} \setminus A$.

---

*Emails of authors: zpentomino@gmail.com, sedthilk@thilikos.info, dzoros@math.uoa.gr.

†Department of Mathematics, National and Kapodistrian University of Athens, Athens, Greece.

‡AlGCo project-team, CNRS, LIRMM, France.

§The work of the second author has been supported by the project DEMOGRAPH (ANR-16-CE40-0028).

An *embedding* of a graph G on a surface S is a representation Λ of G on S with points of S associated bijectively with vertices and open arcs associated bijectively with edges in such a way that:

i) the endpoints of an arc are associated with the ends of the edge that is associated with said arc,

ii) the points of the arcs are not associated with vertices, and

iii) two arcs never intersect over a point of either.

Since we will be studying graphs that are embedded on surfaces, we will be using a common symbol for both the graph and its embedding, extending this liberty to vertices/points and edges/open arcs.

A connected subset X of a surface S is a *star* if it contains a point v_x called its *centre* such that $X \setminus \{v_x\}$ is a union of pairwise disjoint open arcs called *half-edges*. The set of endpoints of a star is the set $\bar{X} \setminus X$.

We may now define an *embedding* Λ *of a hypergraph* G on a surface S to be a representation of G on S with points of S associated bijectively with vertices and stars associated bijectively with hyperedges in such a way that:

i) the endpoints of a star are associated with the ends of the hyperedge that is associated with said star,

ii) the points of the stars are not associated with vertices, and

iii) two stars never intersect over a point of either.

Again we will be using a common symbol for both the hypergraph and its embedding. A *face* of an embedded hypergraph G is a connected component of $S \setminus G$.

The *dual graph* of a graph G embedded on a surface S is the graph G^* embedded on S that has exactly one vertex lying in the interior of each face of G and exactly one edge crossing each edge of G and adjoining the vertices (which may coincide) of G^* that rest inside the faces (which may coincide) of G on either side of said edge of G. An edge e^* of G^* crossing an edge e of G is called the *dual edge* of e. Similarly, the *dual of an hypergraph* G embedded on S is a hypergraph G^* which has exactly one vertex lying in the interrior of each face of G and exactly one star s^* aligned over each star s of G such that their centres coincide and the endpoints of s^* (which may in pairs coincide) are the vertices of G^* lying in those faces of G that are incident to s (these also may in pairs coincide).

A *cycle basis* of a graph G is a minimal set C of cycles such that every Eulerian subgraph of G (i.e., a subgraph containing a cycle which uses each edge of the subgraph exactly once) can be expressed as the symmetric difference of cycles in C. The number of cycles in every cycle basis of G is constant and equal to $|E(G)| - |V(G)| + c$, wherein c is the number of connected components of G. This number is called the *circuit rank* of G.

An *edge partition* of a hypergraph graph G is a partition of its edges to two disjoint sets, say (B, R). We will be referring to B as the *blue set*, whereas R will be the *red set* (throughout this paper, we will follow this terminology in all aspects of the partition, using phrases like *red edge*, *blue degree of a vertex*, *red cut set of* G, etc.). The *border* $\delta(B, R)$ of an edge partition of G is the largest subset of $V(G)$ whose each element is the end of both an edge of B and an edge of R. More generally, for two subsets U, V of S, we define $\delta(U, V) = \bar{V} \cap \bar{U}$. Since we have defined graph embeddings in such a way that edges can only intersect each other at their ends, the graph theoretic and the set theoretic definitions of the border are equivalent for graphs and will be used interchangeably.

A *branch decomposition* T of a hypergraph G is an (unrooted) binary tree whose leaves are bijectively associated with the edges of G. It is easy to notice that each edge e of T induces an edge partition (E_1^e, E_2^e) of G, wherein the parts are the sets of edges that label the leaves of each tree of $T \setminus e$. A branch decomposition is *connected* when for every $e \in E(T)$, both E_1^e and E_2^e induce a connected subgraph of G. The function $\delta_T(e) = |\delta(E_1^e, E_2^e)|$ is called *the evaluating function.* The *width* of a branch decomposition T is defined as $w(T) = \max\{\delta_T(e) \mid e \in\}$. Among all the branch decompositions of G, there are some with the least width. That is the *branchwidth* of G, and is denoted $\mathbf{bw}(G)$. A branch decomposition T of G with $w(T) = \mathbf{bw}(G)$ is called *optimal.*

A *daisy* is the graph containing a single vertex with multiple loops. Notice that the dual of a daisy is a tree and that both a daisy and a tree has branchwidth 1.

If T is a tree whose leaves are labeled, we denote by $L(T)$ the set of the labels of the leaves of T.

The *dual of a branch decomposition* T of G is a branch decomposition T^* of G^* whose underlying tree, say H, is the same as that of T and if a leaf of H is labeled by the edge a in T, then it is labeled by the edge a^* in T^*.

Let e be an edge of the underlying tree of a branch decomposition T defining the edge partition (E_1, E_2), then we denote by (E_1^*, E_2^*) the edge partition that e defines on T^*, and we further may independently refer to the respective parts E_1^*, E_2^*. It is important to note that E_1^*, E_2^* are not at all the dual graphs of E_1, E_2 respectively considered as graphs embedded on S, but rather they are the sets of edges dual to the edges of E_1, E_2 respectively.

When we *delete* an hyperedge e from a hypergraph $G = (V, E)$ we remove e from E but keep its endpoints in V. We denote the hypergraph obtained by $G \setminus e$.

The *edge contraction* of an edge $e = \{u, v\}$ of a graph G is the operation that deletes e, adds a new vertex x_{uv} to V and connects x_{uv} to all the neighbours of u and v (notice that, as we consider multigraphs, this operation may create multiple edges, and that a contraction of an edge makes a triangle into a *dipole*, and a dipole into a loop). We denote the graph obtained by G/e.

Definition 1. *Let S be a surface and Λ_S the set of all finite, connected, bridgeless hypergraphs embedded on S. A function $f : \Lambda_S \to \mathbb{N}$ is called (a, b)-self-dual if for every $G \in \Lambda_S$ it holds that $f(G^*) \leq a \cdot f(G) + b$. Moreover, if $a = 1$, then*

f is called additively self-dual.

We aim to show that the branchwidth function $\mathbf{bw}(G)$, defined over connected, *bridgeless* graphs (*bridge* is an edge whose deletion disconnects the graph) embedded on an oriented 2-manifold is $(1, 2g)$-self-dual, and then extend this result to hypergraphs.

2.1 The planar case

We first illustrate our modus operanti in the planar case. In [1], a proof is provided ensuring that the branchwidth of a connected, bridgeless graph is indeed equal to its *matroid branchwidth.* This is defined to be a version of branchwidth with $w(E_1, E_2) = r(E_1) + r(E_2) - r(E(G)) + 1$ as the evaluating function for a partition (E_1, E_2), wherein for a set of edges E, $r(E)$ is the cardinality of a maximal acyclic subset of E (that is, a spanning forest - in matroid language, $r(E(G))$ is the rank of the graphic matroid of G). This is well-defined and naturally coincides with the branchwidth of the cycle matroid of G. It also provides an immediate proof for the planar case, since the branchwidth of a planar matroid is equal to that of its dual. However, this cannot be extended to a general proof of our hypothesis in 2-manifolds, as the only matroids that are at the same time graphic and co-graphic are planar matroids. Still, as observed in [1], $w(E_1, E_2) = n_1 - c(E_1) + n_2 - c(E_2) - n + c(E(G)) + 1$, wherein n_1, n_2, n are the vertices spanned by $E_1, E_2, E(G)$ respectively and c is the function over graphs that counts the connected components of its arguments. Since $n_1 + n_2 - n = |\delta(E_1, E_2)|$, we obtain $w(E_1, E_2) = |\delta(E_1, E_2)| + c(E(G)) - c(E_1) - c(E_2) + 1$.

Theorem 1. *The branchwidth of a connected, bridgeless, embedded planar graph is equal to that of its dual.*

Proof. Let G be a planar, connected, bridgeless graph and let (B, R) be an edge partition of G.

By a well-known theorem [2], according to which there is always an optimal connected branch decomposition, we may consider B, R to be connected edge-sets of G. However, in this case, and since G is a connected graph, it holds that $w(B, R) = |\delta(B, R)|$. It is then convenient to use w as the evaluating function, since we need only consider connected branch decompositions of G. We aim to calculate the difference $\mathbf{bw}(G) - \mathbf{bw}(G^*)$. Let us begin by calculating the difference $|w(B, R)| - |w(B^*, R^*)|$. Without loss of generality, we can choose a blue edge b and contract it to study how this difference will change. Suppose b is not a loop. The contraction of b reduces $r(E(G))$ by 1, since G/e is connected and has one fewer vertex than G, and for every connected graph H it holds that $r(E(H)) = V(H) - 1$. Furthermore, it reduces $r(B)$ by 1 for the exact same reason. Finally, if both ends of b are contained in $w(B, R)$, $r(R)$ is reduced by one (to see this, consider the graph $R' = R \cup \{e\}$; a spanning tree of R is also a spanning tree of R', and we can use the same argument as in the two previous cases). If that is not the case, then $r(R)$ remains unaltered. That means that

$|w(B, R)|$ is reduced by either 1, if b lies between red edges, or 0, otherwise. If all blue edges are loops, then we do nothing.

We now study the change in $|w(B^*, R^*)|$. Again, let us assume that b is not a loop. We observe that the contraction of b in G entails the subtraction of b^* in G^*. This leaves $r(R^*)$ unchanged, and $r(E(G^*))$ as well. Indeed, if $G^* \setminus b^*$ does not have a common spanning tree with G^*, then b^* is a bridge of G^*, and thus b is a loop, contrary to our assumption. Similarly, $r(B^*)$ is reduced by 1 iff b^* is a bridge of B^*. Again we remind that, if all blue edges are loops, we enact no changes. Unitl this point, we have made no use of the fact that G is a planar graph.

We now notice the following fact: that the contraction of a non-loop blue edge of G reduces $r(R)$ iff it reduces $r(B^*)$. Indeed, suppose that b^* is not a bridge of B^*. Then b^* is contained in a cycle of B^*, thus it is contained in a blue cycle of G^*, thus, by cut-cycle duality, b is contained in a blue minimal cut set of G. Since R is connected, we deduce that one of the ends of b ought to have a red degree equal to 0. For the inverse, if one of the ends of b, say v, has a red degree of 0, then let $b, e_1, ..., e_k$ be the (blue) edges that have v as endpoint, with the exception of any loops, presented clockwise (if v belongs only in loops and non-separating edges, then b is a bridge of G, a contradiction). Then $b^*, e_1^*, ..., e_k^*$ is a cycle of B^*, which means that b^* is not a bridge.

The above implies that $|w(B, R)| - |w(B^*, R^*)|$ remains unchanged by contractions of blue edges that are not loops. The exact same argument applies for red edges. Now we can start contracting non-loop edges indiscriminately. Let G_k be the result of applying some k edge contractions on G, and similarly define B_k, R_k. Contraction leaves connected graphs connected, so, as long as G_k has more than one vertices, it also has a non-loop edge. Let $n = |V(G)|$. Notice that G_{n-1} is in fact a *daisy* and therefore $w(B_{n-1}, R_{n-1}) = w(B^*_{n-1}, R^*_{n-1}) = 1$. Consequently, $|w(B, R)| - |w(B^*, R^*)| = |w(B_{n-1}, R_{n-1})| - |w(B^*_{n-1}, R^*_{n-1})| = 0$ for every connected edge partition (B, R).

Thus, if T is an optimal branch decomposition of G, it holds that $w(T^*) = w(T) = \mathbf{bw}(G) \Rightarrow \mathbf{bw}(G^*) \leq w(T^*) = \mathbf{bw}(G)$. □

2.2 $(1, 2g)$-self-duality of branchwidth on graphs

Observe that in the previous proof we used the cut-cycle duality Theorem. Cut-cycle duality appears to lie in the vey heart of our approach, and yet it does not extend to general 2-manifolds, because Jordan's Theorem doesn't - on an oriented 2-manifold M of positive genus, it can be true that a simple closed curve C is such that $M \setminus C$ is connected. However, the following holds:

Proposition 1 ([3])**.** *Consider an oriented manifold M of genus g and a set $\{C_1, ..., C_k\}$ of distinct closed curves embedded on M. If $M \setminus \cup_i^k C_i$ is connected, then $k \leq 2g$.*

Using this proposition we can prove the following.

Theorem 2. *The branchwidth of graphs embedded on an oriented 2-manifold of genus g is a $(1, 2g)$-self-dual width parameter.*

Proof. Let M be some oriented 2-manifold of genus g. Consider a connected, bridgeless graph G edge-partitioned to the connected sets (B, R), much like before, though this time embedded on M. Again, we care to calculate the difference $|w(B,R)| - |w(B^*,R^*)|$. Once more, we begin contracting non-loop blue edges. Following the exact same arguments as in the planar case, we deduce that $|w(B,R)|$ is reduced by either 1 if b lies between red edges or 0 otherwise, and that $|w(B^*,R^*)|$ is reduced by 1 if b^* is a bridge of B^* or 0 otherwise. Moreover, if b^* is a bridge of B^*, then b lies between red edges. Indeed, let v be an endpoint of b with red degree equal to 0 and let $b, e_1, ...e_k$ be the (blue) edges that having v as endpoint, excluding contractible loops and non-separating edges, presented clockwise. This description applies to at least one edge - indeed, if b^* is a bridge rather than a loop, then b is a surface-separating edge, thus its side surfaces are distinct, which would not be the case if there was a line starting at one of its side surfaces, ending at the other, and crossing only non-separating edges. Then $b^*, e_1^*, .., e_k^*$ is a blue cycle containing b^*, therefore the latter is not a bridge of B^*, a contradiction. Thus, we may contract all non-loop edges of B whose duals are bridges of B^* without altering $|w(B,R)| - |w(B^*,R^*)|$. Here the similarities with the planar case end.

Suppose that b^* is not a bridge of B^*, and thus is contained in a cycle of B^*. That does not necessarily mean that b is in a cut set of G, hence it might well be that both of the ends of b have positive red degree, without that meaning that R has somehow become disconnected.

According to Proposition 1, as long as there is a cycle basis with more than $2g$ cycles in B^*, it divides M to at least 2 connected components. Then there is an edge b^* belonging to a cycle C of said cycle basis of B^* that bounds two distinct faces (that cycle might be b^* itself, if the latter is a loop). So, instead of any arbitrary edge of B whose dual is not a bridge of B^*, we choose and contract b. Since R is connected, the red degree of one of the ends of b is 0. Indeed, were it otherwise, there would be a red edge in either side of C, that is, inscribed in distinct faces of B^*. The edges of B^* are only traversed by edges of B. Thus, if r_1 and r_2 are red edges contained in distinct faces of B^*, they cannot be connected by a red path, meaning that R is disconnected, a contradiction.

It is safe then to contract b (hence delete b^*), in the sense that the count $|w(B,R)| - |w(B^*,R^*)|$ will not be altered. Furthermore, deleting b^* diminishes the circuit rank of B^* by 1. We consider again a cycle basis of B^* and repeat. Applying this technique, we are left in the end with a cycle basis of $2g$ (or fewer) cycles in B^*. Notice that these remaining cycles may not separate M to distinct connected components, in which case the above argument fails and it is quite possible that any contractions of edges whose dual is not a bridge lead to a reduction of $|w(B,R)|$ by 1 without affecting $|w(B^*,R^*)|$. With every such occurence, the difference $|w(B,R)| - |w(B^*,R^*)|$ is reduced by 1, and so does the circuit rank of B^*.

From the above observations, we deduce that we can contract all non-loop edges of B, inducing a reduction of $|w(B,R)| - |w(B^*,R^*)|$ by some number $k \leq 2g$. We do exactly that, reducing B to a daisy. Notice that the blue daisy is indeed a set of cycles (loops), so, if these happen to exceed $2g$ in number, they divide M to distinct connected components. In that case, we choose a loop b of B that bounds two distinct components of $M \setminus B$, and we subtract it. As a result, its dual will contract. Let us check whether $|w(B^*, R^*)|$ changes. Of course, both $r(E(G^*))$ and $r(B^*)$ are reduced by 1, inducing no change - so it is a matter of whether $r(R^*)$ changes. If either end of b^* has a red degree of 0, R^* remains as it was, and so does $r(R^*)$. Otherwise, b^* connects two distinct connected components of R^* (since they lie inside distinct faces of B, whose boundaries are only traversed by edges of B^*, thus there cannot be a red path of G^* connecting them). Joining any two graphs from a common vertex produces a graph whose spanning tree has an equal number of edges to the spanning forest of their disjoint union. Thus, $r(R^*)$ does not change in that case either.

Each such loop-subtraction of B leaves one fewer blue loop in G, all the while maintaining the difference $|w(B, R)| - |w(B^*, R^*)|$. At some point, this process shall leave B with $2g$ loops or fewer (if there were fewer to begin with). Then our above argument may not be likewise implemented, and every subtraction of a loop of B may lead to its dual contracting so as to adjoin two red edges of G^* that already belong to the same connected component of R^*, reducing $r(R^*)$ (and thus $|w(B^*, R^*)|$) by 1. When the last loop of B goes, $|w(B, R)| - |w(B^*, R^*)|$ might have grown for as much as $l \leq 2g$.

What now remains is a fully red graph (the subgraph of G that we have been calling R), along with its dual embedding on M. Thus, $|w(\emptyset, R)| = |w(\emptyset, R^*)| = 1$. However, $||w(B,R)| - |w(B^*,R^*)| - k + l| = |w(\emptyset,R)| - |w(\emptyset,R^*)| = 0$, therefore $||w(B,R)| - |w(B^*,R^*)|| \leq |l - k| = 2g$. Hence, the widths of dual branch decompositions lie in an interval of $2g$ from each other. Indeed, if $w(E_1^*, E_2^*) = w(T^*)$ for a branch decomposition T of G, then $w(T^*) = w(E_1^*, E_2^*) \leq w(E_1, E_2) + 2g \leq w(T) + 2g$. So if T is an optimal branch decomposition for G, then $\mathbf{bw}(G) = w(T) \geq w(T^*) - 2g \geq \mathbf{bw}(G^*) - 2g \Rightarrow \mathbf{bw}(G^*) \leq \mathbf{bw}(G) + 2g$, which concludes the proof. □

3 $(1, 2g)$-self-duality of branchwidth

Before we proceed to the proof of our main result we need to establish some more notation.

For a given unlabeled hypergraph G we define its *incidence graph* $I(G)$ so that every vertex and hyperedge of G is associated with a unique vertex of $I(G)$ and two vertices of $I(G)$ are incident if and only if one is associated with a vertex v of G, the other is associated with a hyperedge e of G, and v is an end of e. We will refer to the set of edges of $I(G)$ that are incident to a vertex of $I(G)$ associated with the hyperedge e of G with the symbol $E(e)$. Every edge of $I(G)$ has only one end associated with a hyperedge, so it only belongs to one such set, that is, $\{E(e) \mid e \text{ is a hyperedge of } G\}$ is a disjoint family of sets.

Now, for this final part, consider G to be a hypergraph embedded on a surface S. Notice that any given embedding of G on S naturally defines an embedding of $I(G)$ on S, with the vertices of $I(G)$ that are associated with vertices of G being represented as the respective vertices of G, the vertices of $I(G)$ that are associated with hyperedges of G being represented as the centres of the respective stars of G, and the edges of $I(G)$ incident to a vertex of $I(G)$ that is associated with the vertex v of G and a vertex of $I(G)$ that is associated with the hyperedge e of G being represented as the half edge that connects v with the centre of e. In future, when referencing $I(G)$, we will always consider it to be embedded on S in this manner, and will conveniently identify the edges of $I(G)$ with the half-edges of G to which they are associated.

Let T be a branch decomposition of G. We will denote by $I(T)$ the branch decomposition of $I(G)$ that is obtained from T as follows: first, for each edge e of G, we construct a rooted binary tree T_e whose leaves are bijectively labeled with the edges in $E(e)$ in some order. We will name these trees "half-edge trees". Observe that there are several trees that match the description of T_e, but the exact choice we make upon the creation of $I(T)$ is of no consequence and will not come up in the proof. Consequently, by T_e we will denote any specific representative of the set of possible half-edge trees of a star of G. After constructing them, we conjoin the half-edge trees with T, identifying the root of T_e with the leaf of T labeled by e for each $e \in E(G)$. Finally, in the ensuing labeled tree, we unlabel the labeled internal nodes. Notice that the underlying tree of T and the half-edge trees are subtrees of $I(T)$, and we will be referring to them as such.

To prove Theorem 3 we need first prove the following lemmata.

Lemma 1. *Let T be a branch decomposition of a hypergraph G. Then $w(T) = w(I(T))$.*

Proof. Let e be an edge of $I(T)$. If e does not belong to a half-edge tree but rather to T, then $\delta_{I(T)}(e) = \delta(E_1, E_2)$ for two disjoint half-edge sets E_1, E_2 such that if for some hyperedge y of G we have $x \in E(y)$ and $x \in E_i$ then $E(y) \subseteq E_i$, that is, half-edges of a common hyperedge are in the same part of the partition. Now let $E'_i, i \in \{1, 2\}$ be the set of hyperedges whose half-edges belong to E_i. Observe that E_i and E'_i define the same subsets of S, and so, by definition, $\delta_{I(T)}(e) = \delta(E_1, E_2) = \delta(E'_1, E'_2) = \delta_T(e)$, so that such a partition of $I(G)$ defines a partition of G with the same border. For the other case, if e belongs to T_h for some $h \in E(G)$, then $\delta_{I(T)}(e) = \delta(E_1, E_2)$, wherein E_1 only contains some half-edges from the same star h of G. This border then contains only certain ends of said star, plus its centre. So we obtain that $\delta(E_1, E_2) < \delta(h, E(G) \setminus h) + 1 \leq (\delta(h, E(G) \setminus h) - 1) + 1 = \delta(h, E(G) \setminus h) = \delta_{I(T)}(e')$ for $e' \in T$ being the edge of $I(T)$ that connects to the root of T_h. Hence, $w(I(T))$ is obtained for a partition of $I(T)$ defined by an edge $e \in T$, and, for every $e \in T$, such a partition of $I(G)$ defines a partition of G with the same border. The result is imminent. □

Lemma 2. *Let T be a branch decomposition of a hypergraph G. Then $w(I(T^*)) = w(I(T)^*)$.*

Proof. We have shown above that $w(I(T^*))$ is obtained for a partition of $I(T^*)$ defined by an edge $e \in T^*$. We will first demonstrate the respective result for

$w(I(T)^*)$. Since the underlying tree of $I(T)^*$ is the same as that of $I(T^*)$, we can naturally refer to the underlying trees of T^* and of the half-edge subtrees of $I(T^*)$ as subtrees of $I(T)^*$. We observe that, for a star h of G, $E(h)^*$ is a cycle of $I(G)^*$ with as many edges as $E(h)$. Let e be an edge of $I(T)^*$ that belongs to a subtree T_h^*. Then $\delta(e) = \delta(E_1, E_2)$, wherein E_1 contains only certain edges of $E(h)^*$, that is, $\delta(E_1, E_2)$ contains only the ends of some of the edges in $E(h)^*$. Thus, $\delta(E_1, E_2) \le \delta_{I(T)^*}(e')$, wherein e' is the edge of $I(T)^*$ that connects to the root of T_h^*, and hence the respective border contains the ends of all of the edges in $E(h)^*$. Now we need only show that for $e \in T$ the partitions of $I(G^*)$ and $I(G)^*$ defined by e on $I(T^*)$ and on $I(T)^*$, say (E_1, E_2) and (E_1', E_2'), have the same border. Indeed, recall that, for every $h \in G$ and each of the two partitions, $E(h)$ will be contained wholly in one of the parts, so that the centre of h cannot be part of the border. So $E(h^*)$ and $E(h)^*$ contribute to the borders of the two partitions, respectively, the same set U_h of vertices of G^* corresponding to the faces of G that are incident to both h and another star whose half-edges all belong to a different part of the partition. So then

$$\begin{aligned}
\delta_{I(T^*)}(e) &= \delta(E_1, E_2) \\
&= \delta(\cup_{\{h\in G | E(h^*)\in E_1\}} E(h^*), \cup_{\{h\in G | E(h^*)\in E_2\}} E(h^*)) \\
&= \delta(\cup_{\{h\in G | E(h^*)\in E_1\}} U_h, \cup_{\{h\in G | E(h^*)\in E_2\}} U_h) \\
&= \delta(\cup_{\{h\in G | E(h)^*\in E_1'\}} E(h)^*, \cup_{\{h\in G | E(h)^*\in E_2'\}} E(h)^*) \\
&= \delta(E_1', E_2') = \delta_{I(T)^*}(e)
\end{aligned}$$

Since this equality is true for every $e \in T$ and both branch decompositions manifest their width for a partition defined by such an e, we have $w(I(T^*)) = w(I(T)^*)$. □

Lemma 3. *Let G be a hypergraph. There exist a connected branch decomposition of G of width* $\mathbf{bw}(G)$.

Proof. To prove this fact, we will have to delve somewhat deeper into the specifics of [2]. In general, the authors of [2] prove that, for every branch decomposition T of a bridgeless, connected graph, there is another which is connected and its width does not exceed $w(T)$. The corollary that we have used, namely that for 2-edge-connected graphs one can always find a connected optimal branch decomposition, is simply obtained by applying this result to any optimal branch decomposition of a given graph. To reach this conclusion, the authors of [2] first constructed an algorithm over branch decompositions named *Make-It-Connected*. This algorithm searches the input for *quartets*. A quartet in a branch decomposition T is an ordered set (A_1, A_2, B_1, B_2) of four mutually disjoint subtrees of T satisfying the following: 1. there is an edge $e = \{x, y\}$ of T such that the roots a_1 and a_2 of A_1 and A_2 are both adjacent to x in T , and the roots b_1 and b_2 of B_1 and B_2 are both adjacent to y in T; 2. $\delta(L(A_1), L(B_1)) \neq \emptyset$ and $\delta(L(A_2), L(B_2)) \neq \emptyset$; 3. $\delta(L(A_1), L(A_2)) = \emptyset$. While there exists a quartet in T, the algorithm Make-It-Connected considers it and and interchanges the positions of A_2 and B_1 in T, or the positions of A_2 and B_2

if (A_1, A_2, B_2, B_1) is also a quartet, with a preference for the change that yields the branch decomposition of the least width among the two, and a bias for the former in case of equality. In [2], it is proven that this algorithm halts, producing a branch decomposition that is connected and has width no larger than that of the input. To prove our claim, we simply consider an optimal branch decomposition T of G and perform the Make-It-Connected process on $I(T)$. Since $\mathbf{bw}(G) = w(T) = w(I(T))$, this will produce a connected branch decomposition R of $I(G)$ with width at most $\mathbf{bw}$(G). Now observe that Make-It-Connected preserves the half edge trees. Indeed, for $h \in G$, a step of Make-It-Connected may consider a quartet that is fully contained within T_h, so that the altered T_h remains a half edge tree for h. Otherwise, it may consider a quartet such that T_h is fully contained in one of its four parts. This leaves T_h in itself unaltered. Finally, it may not consider a quartet (A_1, A_2, B_1, B_2) such that T_h is the smallest tree that contains A_1, A_2 because this won't be a quartet, as condition (3) is violated - the centre of h is contained in $\delta(L(A_1), L(A_2))$. So R is in fact of the form $I(T')$ for a branch decomposition T' of G. We can obtain T' from R simply by labeling h the root of T_h for every $h \in G$ and then removing all the half edge trees. Now, since $I(T')$ is connected, T' is connected as well. Indeed, if $e \in T'$, then we have shown at the proof of claim 1 that the partitions of G and $I(G)$ defined by e on T' and $I(T')$ respectively have the same two subsets U, V of S as parts. If $I(T')$ is connected, that means that, for every $e \in I(T')$, U^e and V^e are connected, so the same applies to all $e \in T'$, so T' is connected. Also, $\mathbf{bw}(G) \leq w(T') = w(R) \leq w(I(T) = w(T) = \mathbf{bw}(G)$, so that T' is also optimal. The proof is complete. □

We can now prove the following.

Theorem 3. *The branchwidth of hypergraphs embedded on an oriented* 2*-manifold of genus* g *is a* $(1, 2g)$*-self-dual width parameter.*

Proof. According to Lemma 3 there exists a connected branch decomposition of G, say T, such that $\mathbf{bw}(G) = w(T)$. From Lemma 1 it follows that $\mathbf{bw}(G) = w(I(T))$. As $w(I(T)) \geq w(I(T)^*) - 2g$ it holds that $\mathbf{bw}(G) \geq w(I(T)^*) - 2g$. According to Lemma 2 we conclude that

$$\mathbf{bw}(G) \geq w(I(T^*)) - 2g = w(T^*) - 2g \geq \mathbf{bw}(G^*) - 2g$$

Thus $\mathbf{bw}(G^*) \leq \mathbf{bw}(G) + 2g$. □

References

[1] *Branchwidth of graphic matroids*, Frederic Mazoit and Stephan Thomasse.

[2] *Connected graph searching*, Lali Barrière, Paola Flocchini, Fedor V. Fomin, Pierre Fraigniaud, Nicolas Nisse, Nicola Santoro, Dimitrios M. Thilikos, Information and Computation 219, 2012.

[3] *Knots and links*, Peter R. Cromwell, Cambridge University Press, 2004.

[4] *Tree-width of hypergraphs and surface duality*, Frederic Mazoit, Journal of Combinatorial Theory, Series B, 102, 2012.

[5] *On self-duality of branchwidth in graphs of bounded genus*, Ignasi Sau, Dimitrios M. Thilikos, Discrete Applied Mathematics, Volume 159, Issue 17, 2011, Pages 2184-2186