The Merino–Welsh Conjecture holds for Series–Parallel Graphs

Steven D. Noble[1] and Gordon F. Royle[2]

[1]Department of Mathematics, Brunel University, Uxbridge, UB8 3PH, United Kingdom
[2]Centre for the Mathematics of Symmetry and Computation, School of Mathematics & Statistics, University of Western Australia, 35 Stirling Highway, Nedlands WA 6009, Australia

November 13, 2014

Abstract

The Merino-Welsh conjecture asserts that the number of spanning trees of a graph is no greater than the maximum of the numbers of totally cyclic orientations and acyclic orientations of that graph. We prove this conjecture for the class of series-parallel graphs.

1 Introduction

In this paper, we are concerned with the relationship between three graph parameters, namely the number of spanning trees, the number of acyclic orientations, and the number of totally cyclic orientations of a graph. In principle, our graphs may have loops, bridges and multiple edges, although only the last will play any eventual role. Recapping some basic terminology, recall that a *spanning tree* of a graph G is a set of edges inducing a connected spanning subgraph of G, but containing no cycle of G; we use $\tau(G)$ to denote the number of spanning trees of G. An *orientation* of a graph is an assignment of a direction to each edge of the graph — the orientation is called *acyclic* if the resulting directed graph contains no directed cycles, and it is called *totally cyclic* if *every* edge is contained in a directed cycle. We use $\alpha(G)$ to

1

denote the number of acyclic orientations of G, and $\alpha^*(G)$ to denote the number of totally cyclic orientations of G. It is immediate that if G contains a loop, then $\alpha(G) = 0$ and that if G contains a bridge, then $\alpha^*(G) = 0$.

Examination of empirical evidence led Merino and Welsh to make the following conjecture relating these three graphical parameters:

Conjecture 1.1 (Merino-Welsh Conjecture [5]). *For any bridgeless, loopless graph G,*

$$\max\{\alpha(G), \alpha^*(G)\} \geq \tau(G). \tag{1}$$

The $\max(\cdot)$ function is awkward to carry through any sort of inductive proof, but there are two very natural generalisations of this conjecture involving the arithmetic and geometric means of α and α^*. It is easy to see that either of these conjectures, both of which first appeared in Conde & Merino [2], imply the original Merino-Welsh conjecture.

Conjecture 1.2 (Additive Merino-Welsh Conjecture). *For any bridgeless, loopless graph G,*

$$\alpha(G) + \alpha^*(G) \geq 2\tau(G). \tag{2}$$

Conjecture 1.3 (Multiplicative Merino-Welsh Conjecture). *For any bridgeless, loopless graph G,*

$$\alpha(G)\alpha^*(G) \geq \tau(G)^2. \tag{3}$$

As an example, we can consider the self-dual family of *wheels* where W_n consists of a single central vertex connected to every vertex of an n-cycle. The graph W_3 is the complete graph K_4, which has 16 spanning trees, 24 acyclic orientations and 24 totally cyclic orientations. In general, for $n \geq 3$, the wheel W_n has (the closest integer to) $(1+\tau)^n - 2$ spanning trees, where τ is the golden ratio, and $3^n - 3$ acyclic orientations and totally cyclic orientations, and so easily satisfies all versions of the conjecture. (An easy induction provides the number of acyclic orientation for W_n, while the number of spanning trees was originally determined by Sedláček [7].)

On the other hand, the *digon*, which is the 2-vertex graph with a double edge connecting its two vertices has 2 spanning trees, 2 acyclic orientations and 2 totally cyclic orientations and thus achieves equality in all three variants of the conjecture.

Intuitively, for any fixed number of vertices, the (connected) graphs with few edges tend to be "tree-like" and have more acyclic orientations than spanning trees, whereas for graphs with many edges, the number of totally cyclic orientations tends to dominate the number of spanning trees. This qualitative statement was made more precise by Thomassen.

Theorem 1.4 (Thomassen [8]). *If G is a simple graph on n vertices with $m \leq 16n/15$ edges, then*

$$\alpha(G) > \tau(G),$$

and if G is a bridgeless graph on n vertices with $m \geq 4n - 4$ edges then
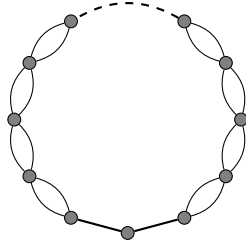
$$\alpha^*(G) > \tau(G).$$

Figure 1: Graph answering Thomassen's question in the negative

This result raises the possibility that the Merino-Welsh conjecture might be resolved by extending *upwards* the value of m for which acyclic orientations are guaranteed to dominate, and *downwards* the value of m for which totally cyclic orientations are guaranteed to dominate, until the two bounds meet and therefore cover every possibility. More precisely, Thomassen asked the following question:

Question 1.5 (Thomassen [8]). If G is a bridgeless, loopless graph with n vertices and m edges, then is it true that

$$\alpha(G) \geq \tau(G) \text{ if } m \leq 2n - 2, \text{ and}$$
$$\alpha^*(G) \geq \tau(G) \text{ if } m \geq 2n - 2?$$

However, the answer to Thomassen's question is "No"; if G is the graph of Figure 1 consisting of $n - 2$ digons and two edges arranged in a cycle (in arbitrary order), then

$$\tau(G) = 2^{n-1} + (n-2)2^{n-3},$$
$$\alpha(G) = 2^n - 2,$$
$$\alpha^*(G) = 2 \cdot 3^{n-2}.$$

Therefore for $n \geq 6$, it follows that $\alpha(G) < \tau(G)$. The dual graphs $H = G^*$ provides examples of n-vertex graphs with $2n - 2$ edges where $\alpha^*(H) < \tau(H)$. This example can easily be generalised to provide examples with fewer than $2n - 2$ edges with the same property — if $G_{n,k}$ is the graph obtained from an n-cycle by replacing any $n - k$ of the edges by digons, then simple counting shows that

$$\tau(G_{n,k}) = k2^{n-k} + (n-k)2^{n-k-1}$$

while $\alpha(G_{n,k}) = 2^n - 2$. If $n = 2^{k+1} - k$, then $\tau(G_{n,k}) = 2^n$ and so $\alpha(G_{n,k}) < \tau(G_{n,k})$. For larger n, the discrepancy is even greater.

Therefore any proof of the conjecture must either consider the two parameters simultaneously, or relate them in a more precise way to the graph structure. In the remainder of this paper, we give a proof that the Merino-Welsh conjecture holds

for *series-parallel graphs*, which are the graphs that arise from a single edge by repeatedly duplicating or subdividing edges in any fashion.

Although series-parallel graphs form only a tiny subset of all graphs, the class is closed under duality and includes numerous graphs with $m = 2n - 2$ edges (including the graphs described above).

2 Deletion and Contraction

The *Tutte polynomial* of a graph G is the 2-variable polynomial $T_G(x, y)$ that is equal to 1 if the graph has no edges, and is otherwise defined recursively by

$$T_G(x, y) = \begin{cases} x\, T_{G/e}(x, y), & \text{if } e \text{ is a bridge;} \\ y\, T_{G-e}(x, y), & \text{if } e \text{ is a loop;} \\ T_{G-e}(x, y) + T_{G/e}(x, y), & \text{otherwise.} \end{cases}$$

Here $G-e$ and G/e are the graphs obtained from G by either *deleting* or *contracting* e respectively. Our three parameters $\tau(G)$, $\alpha(G)$ and $\alpha^*(G)$ are all *evaluations* of the Tutte polynomial

$$\tau(G) = T_G(1, 1), \qquad \alpha(G) = T_G(2, 0), \qquad \alpha^*(G) = T_G(0, 2),$$

and so they individually satisfy the "deletion-contraction" identity. The fundamental problem preventing the Merino-Welsh conjecture from being a trivial exercise in induction is that *deletion* of an edge can create bridges, and *contraction* of an edge can create loops, thus creating graphs to which the inductive hypothesis does not apply.

Two edges e, f are in *series* in a graph if neither is a bridge and any cycle containing e also contains f, and are in *parallel* if neither is a loop and any edge-cutset containing e also contains f. The *series class* of e is the set $\sigma(e)$ of all edges in series with e and is denoted $\sigma(e)$, while the *parallel class* of e is the set $\pi(e)$ of all edges in parallel with e. If $|\sigma(e)| > 1$ then deleting e creates a bridge, while if $|\pi(e)| > 1$, then contracting e creates a loop. However if $|\sigma(e)| = |\pi(e)| = 1$, then e can be safely deleted and contracted, and it is immediate that the additive Merino-Welsh conjecture (Conjecture 1.2) holds for G if it holds for both $G\backslash e$ and G/e. To show that this is also true for the *multiplicative* Merino-Welsh conjecture (Conjecture 1.3) we first need some lemmas due to Jackson [4].

Lemma 2.1. *The condition $\alpha(G)\alpha^*(G) \geq \tau(G)^2$ is equivalent to the statement that for all $\lambda \in \mathbb{R}$,*

$$\alpha(G)\lambda^2 - 2\tau(G)\lambda + \alpha^*(G) \geq 0. \tag{4}$$

Proof. The left-hand side of (4), viewed as a polynomial in λ, is non-negative everywhere if and only if its discriminant is non-positive. As the discriminant is

$$4\tau(G)^2 - 4\alpha(G)\alpha^*(G)$$

the lemma follows. $\qquad\square$

4

Lemma 2.2. *Suppose that e is an edge of a graph G, neither a loop nor a bridge, such that $|\sigma(e)| = |\pi(e)| = 1$. Furthermore, suppose that $G-e$ and G/e, both satisfy the multiplicative Merino-Welsh conjecture (Conjecture 1.3). Then G also satisfies the multiplicative Merino-Welsh conjecture.*

Proof. By the deletion-contraction identity, we have

$$\alpha(G)\lambda^2 - 2\tau(G)\lambda + \alpha^*(G) = (\alpha(G-e)\lambda^2 - 2\tau(G-e)\lambda + \alpha^*(G-e)) \\ + (\alpha(G/e)\lambda^2 - 2\tau(G/e)\lambda + \alpha^*(G/e)).$$

Neither $G-e$ nor G/e has a loop or a bridge and so by Lemma 2.1, both of the bracketed terms are positive. By applying Lemma 2.1 again, the result follows. $\square$

Therefore, in any minimal counterexample to the Merino-Welsh conjecture, every edge must lie in a non-trivial series or parallel class. On the other hand, none of the series or parallel classes can be very large.

Lemma 2.3. *Suppose that G is a graph that satisfies the multiplicative Merino-Welsh conjecture (Conjecture 1.3), and that e is an edge of G with $|\pi(e)| \geq 2$. Let G^+ be formed from G by adding two edges f and g in parallel with e. Then G^+ also satisfies the multiplicative Merino-Welsh conjecture.*

Proof. Given a totally cyclic orientation of G, we may obtain a totally cyclic orientation of G^+ by orienting f and g in any way we choose. Consequently $\alpha^*(G^+) \geq 4\alpha^*(G)$. On the other hand, $\alpha(G^+) = \alpha(G)$, because the addition of parallel edges does not alter the number of acyclic orientations. The number of spanning trees of G^+ which do not contain a member of $\pi_{G^+}(e)$ is the same as the number of spanning trees of G which do not contain a member of $\pi_G(e)$. However the number of spanning trees of G^+ which do contain a member of $\pi_{G^+}(e)$ is equal to the number of spanning trees of G which contain a member of $\pi_G(e)$ multiplied by $|\pi_{G^+}(e)|/|\pi_G(e)|$. Consequently

$$\tau(G^+) \leq |\pi_{G^+}(e)|/|\pi_G(e)|\tau(G) \leq 2\tau(G).$$

Therefore $\alpha(G^+)\alpha^*(G^+) \geq 4\alpha(G)\alpha^*(G) \geq (2\tau(G))^2 \geq \tau(G^+)^2$. $\square$

The following is the analogous version of the previous lemma with parallel replaced by series. The proof is similar but the roles of α and α^* are interchanged.

Lemma 2.4. *Suppose that G is a graph that satisfies the multiplicative Merino-Welsh conjecture, and that e is an edge of G with $|\sigma(e)| \geq 2$. Let G^+ be formed from G by adding two edges f and g in series with e. Then G^+ also satisfies the multiplicative Merino-Welsh conjecture.* $\square$

The consequence of these two results is that in any minimal counterexample to the multiplicative Merino-Welsh conjecture, every edge lies either in a series class of size 2 or 3, or a parallel class of size 2 or 3.

5

3   Series–Parallel Graphs

A *two–terminal* graph (G, s, t) is a graph G with two distinct distinguished vertices, s and t, called the *terminals*, with s designated as the *source* and t as the *sink*. A *two–terminal series–parallel (TTSP)* graph is a two–terminal graph that can be constructed from the two–terminal graph K_2, with the sole edge connecting the source to the sink, by a sequence of the following operations:

1. (parallel connection) take two TTSP graphs (G, s_G, t_G) and (H, s_H, t_H), identify s_G with s_H, forming the source of the new graph $G \oplus_P H$, and identify t_G with t_H, forming the sink of $G \oplus_P H$;

2. (series connection) take two TTSP graphs (G, s_G, t_G) and (H, s_H, t_H), identify t_G with s_H and let s_G and t_H be the source and sink, respectively, of the new graph $G \oplus_S H$.

Both the series and parallel connection operations are associative.

Series-parallel graphs can be defined in various equivalent ways, for example as graphs with no K_4-minor, or as graphs whose blocks are either single edges, or can be reduced to a loop by a sequence of operations each of which is the suppression of a vertex of degree two or the elimination of an edge in parallel to another edge. However for us, the key property of series-parallel graphs is their relationship to two–terminal series–parallel graphs as expressed in the following lemma, which follows from results of Duffin [3]:

Lemma 3.1. *A graph G is a 2-connected series–parallel graph if and only if for every edge $e = st$, the two–terminal graph $(G - e, s, t)$ is a TTSP graph.* $\square$

Consider a rooted binary tree, with each non-leaf node designated as either an *s*-node or a *p*-node (see Figure 2 for an example). We can use this tree as a "blueprint" to construct a TTSP graph by first associating the graph K_2 with each leaf and then, working up the tree, associating with each *s*-node the series connection of its two children, and with each *p*-node the parallel connection of its two children, and finally reading off the graph associated with the root of the tree. This tree is called the *decomposition* tree of the TTSP graph, and any TTSP graph can be described by such a decomposition tree.

Any series-parallel graph is planar and its planar dual is again a series-parallel graph. However, we need a concept of duality—that we denote *sp*-duality—that respects the terminals and allows us to remain within the class of two-terminal series–parallel graphs. Thus we say that two TTSP graphs G and H are "*sp*-dual" if there are decomposition trees for G and H of identical shape, but with all the *s*-nodes changed to *p*-nodes, and vice versa. There is, however, a close relationship between *sp*-duality and planar duality: in particular, if $H = (H, s_H, t_H)$ is the *sp*-dual of $G = (G, s_G, t_G)$ then there are plane embeddings of G and H such that $G + s_G t_G$ is the planar dual of $H + s_H t_H$. We denote the *sp*-dual of G by $G^{*_{sp}}$.
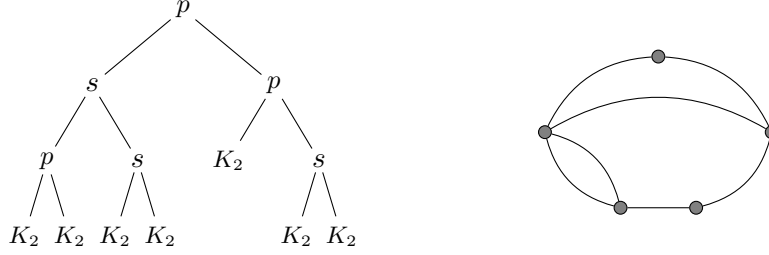
Figure 2: A decomposition tree and associated TTSP graph

Lemma 3.2. *If G, H are TTSP graphs, then*

$$(G \oplus_S H)^{*sp} = G^{*sp} \oplus_P H^{*sp},$$
$$(G \oplus_P H)^{*sp} = G^{*sp} \oplus_S H^{*sp}.$$

Proof. This follows immediately by considering the effect on the sp-tree of changing all s-nodes to p-nodes, and p-nodes to s-nodes. $\square$

A *2-forest* of a TTSP graph G is a spanning subgraph of G with two components, each of which is a tree, with one containing the sink of G and the other containing the source of G. We denote the number of 2-forests of G by $\tau_2(G)$. A *very acyclic orientation* of a two-terminal series-parallel graph is an acyclic orientation in which there is no directed path between the two terminals. We denote the number of very acyclic orientations of G by $\alpha_2(G)$. Finally, an *almost totally cyclic orientation* of a two-terminal series-parallel graph is an orientation in which each edge either lies in a directed cycle or lies on a directed path between the two terminals. We denote the number of almost totally cyclic orientations of G by $\alpha_2^*(G)$. The rationale behind introducing these additional parameters, thereby apparently complicating the problem, is that for a series-parallel graph, it is possible to keep track of these parameters through the operations of series and parallel connection.

Lemma 3.3. *Suppose that $G = G_1 \oplus_S G_2$. Then*

$$\tau(G) = \tau(G_1)\tau(G_2),$$
$$\tau_2(G) = \tau(G_1)\tau_2(G_2) + \tau_2(G_1)\tau(G_2),$$
$$\alpha(G) = \alpha(G_1)\alpha(G_2),$$
$$\alpha_2(G) = \alpha(G_1)\alpha(G_2) - \frac{(\alpha(G_1) - \alpha_2(G_1))(\alpha(G_2) - \alpha_2(G_2))}{2},$$
$$\alpha_2^*(G) = \alpha_2^*(G_1)\alpha_2^*(G_2) - \frac{(\alpha_2^*(G_1) - \alpha^*(G_1))(\alpha_2^*(G_2) - \alpha^*(G_2))}{2},$$
$$\alpha^*(G) = \alpha^*(G_1)\alpha^*(G_2).$$

Proof. The arguments are straightforward counting arguments, and so we just give one example, namely the count for $\alpha_2(G)$. We first count the number of acyclic orientations of $G_1 \oplus_S G_2$ that *do* admit a directed path between the terminals, noting that any such path must either run source-to-sink or sink-to-source but *not both*. The restriction of such an acyclic orientation to G_1 must admit a directed path between its terminals, so there are $\alpha(G_1) - \alpha_2(G_1)$ such acyclic orientations and similarly for G_2. Exactly half of the resulting $(\alpha(G_1) - \alpha_2(G_1))(\alpha(G_2) - \alpha_2(G_2))$ combinations have the paths aligned consistently thereby creating a directed path between the terminals of $G_1 \oplus_S G_2$. Subtracting this number from the total number of acyclic orientations of $G_1 \oplus_S G_2$ gives the stated result. The other arguments are mild variants of this. $\qquad\square$

Lemma 3.4. *Suppose that $G = G_1 \oplus_P G_2$. Then*

$$\tau(G) = \tau(G_1)\tau_2(G_2) + \tau_2(G_1)\tau(G_2),$$
$$\tau_2(G) = \tau_2(G_1)\tau_2(G_2),$$
$$\alpha(G) = \alpha(G_1)\alpha(G_2) - \frac{(\alpha(G_1) - \alpha_2(G_1))(\alpha(G_2) - \alpha_2(G_2))}{2},$$
$$\alpha_2(G) = \alpha_2(G_1)\alpha_2(G_2),$$
$$\alpha_2^*(G) = \alpha_2^*(G_1)\alpha_2^*(G_2),$$
$$\alpha^*(G) = \alpha_2^*(G_1)\alpha_2^*(G_2) - \frac{(\alpha_2^*(G_1) - \alpha^*(G_1))(\alpha_2^*(G_2) - \alpha^*(G_2))}{2}.$$

Proof. Straightforward counting arguments. $\qquad\square$

Using the two previous lemmas, and some induction, we can immediately determine the relationship between the parameters of a TTSP graph and its *sp*-dual.

Lemma 3.5. *If H is the sp-dual of G, then*

$$\tau(H) = \tau_2(G), \qquad \alpha(H) = \alpha_2^*(G), \qquad \alpha^*(H) = \alpha_2(G).$$

$\qquad\square$

4   Replaceability and Reducibility

A TTSP graph G is *replaceable* by a TTSP graph H if for any TTSP graph K, $G \oplus_P K$ satisfies Conjecture 1.3 whenever $H \oplus_P K$ satisfies Conjecture 1.3. We say that a TTSP graph is *reducible* if it is replaceable by one with fewer edges. TTSP graphs which are not reducible are called *irreducible*.

Lemma 4.1. *If the graph G is replaceable by H, then for any G', the TTSP graph $G \oplus_P G'$ is replaceable by $H \oplus_P G'$ and the TTSP graph $G \oplus_S G'$ is replaceable by $H \oplus_S G'$.*

Proof. First, we suppose that $(H \oplus_P G') \oplus_P K$ satisfies Conjecture 1.3. Then as $(H \oplus_P G') \oplus_P K = H \oplus_P (G' \oplus_P K)$ and G is replaceable by H, it follows that $G \oplus_P (G' \oplus_P K) = (G \oplus_P G') \oplus_P K$ also satisfies Conjecture 1.3. Therefore $G \oplus_P G'$ is replaceable by $H \oplus_P G'$.

If $G = (G, s, t)$ is a 2-terminal graph, then temporarily let $r(G)$ denote the *reverse* 2-terminal graph (G, t, s) where the roles of the source and sink have been reversed.

Next we suppose that $(H \oplus_S G') \oplus_P K$ satisfies Conjecture 1.3. Now $(H \oplus_S G') \oplus_P K$ is isomorphic (as a graph, but not as a TTSP) to $H \oplus_P (K \oplus_S r(G'))$ and so $G \oplus_P (K \oplus_S r(G'))$ satisfies Conjecture 1.3. But the latter is isomorphic to $(G \oplus_S G') \oplus_P K$ and so we have shown that $G \oplus_S G'$ is replaceable by $H \oplus_S G'$. $\square$

We say that a TTSP graph (G, s, t) is *extendable* if either G or $G + st$ has the property that every edge lies in a parallel class of size two or three, or lies in a series class of size two or three. If a TTSP graph G is *not* extendable, then nor is $G \oplus_S H$ or $G \oplus_P H$ for any graph H.

Lemma 4.2. *If G is an irreducible extendable TTSP graph with at least two edges, then it is either the series connection or the parallel connection of two smaller irreducible extendable TTSP graphs.*

Proof. As G has at least two edges, then there are TTSP graphs G_1, G_2 such that either $G = G_1 \oplus_S G_2$ or $G = G_1 \oplus_P G_2$. It is clear that G_1 and G_2 are extendable, and by Lemma 4.1 they are irreducible. $\square$

The proof that series-parallel graphs satisfy the Merino-Welsh conjectures hinges on showing that certain specific graphs are reducible, which requires demonstrating that they are replaceable. The next lemma gives a sufficient condition for this

Lemma 4.3. *Let G be a TTSP graph. If there is a TTSP graph H with*

$$(\max\{\tau(G)/\tau(H), \tau_2(G)/\tau_2(H)\})^2$$
$$\leq \min\{\alpha(G)/\alpha(H), \alpha_2(G)/\alpha_2(H)\} \cdot \min\{\alpha^*(G)/\alpha^*(H), \alpha_2^*(G)/\alpha_2^*(H)\}, \quad (5)$$

then G is replaceable by H. (Here it is assumed that if $\alpha_2(H) = 0$ then the term $\alpha_2(G)/\alpha_2(H)$ is ignored, and similarly if $\alpha^*(H) = 0$.)

Proof. Recall that

$$\tau(H \oplus_P K) = \tau(H)\tau_2(K) + \tau_2(H)\tau(K),$$
$$\alpha(H \oplus_P K) = \alpha(H)(\alpha(K) + \alpha_2(K))/2 + \alpha_2(H)(\alpha(K) - \alpha_2(K))/2,$$
$$\alpha^*(H \oplus_P K) = \alpha^*(H)(\alpha_2^*(K) - \alpha^*(K))/2 + \alpha_2^*(H)(\alpha^*(K) + \alpha_2^*(K))/2.$$

In other words, $\tau(H \oplus_P K)$ is a linear combination of $\tau(H)$ and $\tau_2(H)$ whose coefficients depend on K and $\tau(G \oplus_P K)$ is the *same linear combination* with $\tau(G)$ and $\tau_2(G)$ replaced by $\tau(H)$ and $\tau_2(H)$; the same is true for the expressions for α and α^*.

So suppose that G and H satisfy (5) and that $H \oplus_P K$ satisfies the multiplicative Merino-Welsh conjecture. Then

$$(\tau(H)t_1 + \tau_2(H)t_2)^2 \leq (\alpha(H)a_1 + \alpha_2(H)a_2)(\alpha^*(H)c_1 + \alpha_2^*(H)c_2) \qquad (6)$$

where t_1, t_2, a_1, a_2, c_1, c_2 are the coefficients depending on K. Replacing H by G, the three terms of this expression are changed as follows

$$(\tau(H)t_1 + \tau_2(H)t_2)^2 \to (\tau(G)t_1 + \tau_2(G)t_2)^2$$
$$(\alpha(H)a_1 + \alpha_2(H)a_2) \to (\alpha(G)a_1 + \alpha_2(G)a_2)$$
$$(\alpha^*(H)c_1 + \alpha_2^*(H)c_2) \to (\alpha^*(G)c_1 + \alpha_2^*(G)c_2)$$

Therefore, when H is replaced by G, the left-hand side of (6) is multiplied by *at most*

$$(\max\{\tau(G)/\tau(H), \tau_2(G)/\tau_2(H)\})^2 \qquad (7)$$

while the right-hand side is multiplied by *at least*

$$\min\{\alpha(G)/\alpha(H), \alpha_2(G)/\alpha_2(H)\} \cdot \min\{\alpha^*(G)/\alpha^*(H), \alpha_2^*(G)/\alpha_2^*(H)\}, \qquad (8)$$

where any terms involving zero denominators are omitted. By the hypotheses of the lemma, the expression (7) is at most equal to the expression (8) and therefore $G \oplus_P K$ satisfies the multiplicative Merino-Welsh conjecture. $\square$

It is straightforward to systematically construct all series-parallel graphs ordered by increasing number of edges—start with K_2, and at each stage form the series connection and parallel connection of one of the not-yet-processed pairs of graphs with the smallest total number of edges, adding the newly-constructed graphs to the growing list. If a reducible graph G is constructed during this process, then by Lemma 4.1, any further graphs constructed using G are also reducible. As previously noted, if a non-extendable graph is produced during this process, then any further graphs constructed using this are also non-extendable. Therefore a modified procedure that immediately discards any graphs that are either non-extendable or that can be shown by Lemma 4.3 to be reducible will produce a list of graphs that still contains all the extendable irreducible graphs (and perhaps some others). The surprise is that this modified procedure terminates, and indeed terminates quite quickly:

No.	Edges	sp-dual	Built	τ	τ_2	α	α_2	α_2^*	α^*
0	1	0	—	1	1	2	0	2	0
1	2	2	$0 \oplus_P 0$	2	1	2	0	4	2
2	2	1	$0 \oplus_S 0$	1	2	4	2	2	0
3	3	4	$0 \oplus_P 2$	3	2	6	0	4	2
4	3	3	$0 \oplus_S 1$	2	3	4	2	6	0
5	3	6	$0 \oplus_P 1$	3	1	2	0	8	6
6	3	5	$0 \oplus_S 2$	1	3	8	6	2	0
7	4	8	$1 \oplus_S 1$	4	4	4	2	14	4
8	4	7	$2 \oplus_P 2$	4	4	14	4	4	2
9	4	10	$1 \oplus_P 2$	5	2	6	0	8	6
10	4	9	$1 \oplus_S 2$	2	5	8	6	6	0
11	5	12	$1 \oplus_S 5$	6	5	4	2	30	12
12	5	11	$2 \oplus_P 6$	5	6	30	12	4	2
13	6	14	$2 \oplus_P 8$	12	8	46	8	8	6
14	6	13	$1 \oplus_S 7$	8	12	8	6	46	8
15	7	16	$5 \oplus_S 8$	12	16	28	18	30	12
16	7	15	$6 \oplus_P 7$	16	12	30	12	28	18
17	7	18	$2 \oplus_P 12$	16	12	102	24	8	6
18	7	17	$1 \oplus_S 11$	12	16	8	6	102	24

Table 1: List containing all extendable irreducible graphs

Proposition 4.4. *Every extendable TTSP graph is either listed in Table 1 or is reducible.*

Proof. Suppose for a contradiction that there is an extendable irreducible TTSP graph not listed in Table 1, and let G be such a graph with fewest edges. As K_2 is in the table, G has more than one edge and so $G = G_1 \oplus_S G_2$ or $G = G_1 \oplus_P G_2$, where both G_1 and G_2 are extendable and irreducible. By the minimality of G, both G_1 and G_2 occur in Table 1. However it is easy to check for each pair of graphs in Table 1 that both $G_1 \oplus_S G_2$ and $G_1 \oplus_P G_2$ are either included in Table 1, are not extendable, or can be shown to be reducible by Lemma 4.3. This information is summarised in Table 2 in the following manner: the graphs in Table 1 are numbered $0, 1, \ldots, 18$ and the rows and columns of Table 2 are indexed by these graphs. The above-diagonal entries of Table 2 give information about the series connection of the corresponding graphs, while the below-diagonal entries refer to the parallel connection. Each entry is either N, indicating that the graph constructed is not extendable, or is a plain integer x indicating that the graph constructed is reducible because it can be replaced by the smaller graph x (using Lemma 4.3), or an integer preceded by an equals sign $=x$ indicating that the graph is isomorphic to graph x, and hence in Table 1. Table 3 gives the same information for the series connection and parallel connection when the two components are isomorphic. $\qquad\square$

Theorem 4.5. *If G is a series–parallel graph without loops or bridges, then*

$$\alpha(G)\alpha^*(G) \geq \tau(G)^2. \qquad (9)$$

Proof. Suppose for a contradiction that G is a series–parallel graph not satisfying Equation (9), and that among all such graphs G has the fewest edges. Then G is 2-connected, because each of α, α^* and τ is multiplicative over blocks. Let e be an edge of G with endvertices u and v. Then, by Lemma 3.1, $(G - e, u, v)$ is a two-terminal series-parallel graph. By Lemmas 2.3 and 2.4, each edge of G lies in a series class of size two or three, or in a parallel class of size two or three. Consequently every edge of $G - e$ lies in a series class of size two or three, except possibly an edge joining u and v or an edge which is contained in every path from u to v in $G - e$, and therefore $G - e$ is extendable. Since G is a minimal counterexample to the theorem, it follows from 4.1 that $G - e$ is irreducible. Therefore $G - e$ is one of the graphs listed in Table 1. However, for each of these graphs, it is easy to check that if an edge is added between the terminals, then the resulting graph satisfies Equation (9), thereby supplying the required contradiction. $\qquad\square$

As the multiplicative Merino-Welsh conjecture is easily seen to imply the original version, we have now proved the titular result of this paper:

Corollary 4.6. *The Merino-Welsh conjecture holds for series-parallel graphs without loops or bridges.* $\qquad\square$

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0	-	=4	=6	N	=10	4	2	2	2	4	6	2	2	2	2	2	2	2	2
1	=5	-	=10	N	2	=11	6	=14	4	7	2	=18	2	7	2	2	7	7	2
2	=3	=9	-	N	6	2	2	2	2	2	2	2	2	2	2	2	2	2	2
3	=9	5	1	-	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
4	N	N	N	N	-	2	2	2	2	2	2	2	2	2	2	2	2	2	2
5	1	1	5	1	N	-	2	=18	=15	3	2	0	2	0	2	2	0	0	2
6	3	1	=12	1	N	1	-	2	2	2	2	2	2	2	2	2	2	2	2
7	1	1	3	1	N	1	=16	-	2	4	2	2	2	2	2	2	2	2	2
8	1	1	=13	1	N	1	=17	1	-	4	2	2	2	2	2	2	2	2	2
9	5	1	1	1	N	1	1	1	1	-	2	7	2	0	2	2	0	0	2
10	3	1	8	1	N	1	4	3	3	1	-	2	2	2	2	2	2	2	2
11	1	1	1	1	N	1	1	1	1	1	1	-	2	4	2	2	2	2	2
12	1	1	=17	1	N	1	0	1	1	1	8	1	-	2	2	2	2	2	2
13	1	1	1	1	N	1	1	1	1	1	1	1	1	-	2	2	2	4	2
14	1	1	8	1	N	1	0	1	1	1	0	1	3	1	-	2	2	2	2
15	1	1	8	1	N	1	0	1	1	1	0	1	1	1	1	-	2	2	2
16	1	1	1	1	N	1	1	1	1	1	1	1	1	1	1	1	-	0	2
17	1	1	1	1	N	1	1	1	1	1	1	1	1	1	1	1	1	-	2
18	1	1	8	1	N	0	0	1	1	1	0	1	1	1	3	0	1	1	-

Table 2: Each $G_1 \oplus_S G_2$ and $G_1 \oplus_P G_2$ is either in Table 1, is non-extendable, or is reducible

G	0	1	2	3	4	5	6	7	8
$G \oplus_S G$	=2	=7	2	N	2	3	2	2	2
$G \oplus_P G$	=1	1	=8	1	N	1	4	1	1

G	9	10	11	12	13	14	15	16	17	18
$G \oplus_S G$	7	2	2	2	4	2	2	0	0	2
$G \oplus_P G$	1	7	1	1	1	3	0	1	1	0

Table 3: Each $G \oplus_S G$ and $G \oplus_P G$ is either in Table 1, is non-extendable, or is reducible

5 Conclusion

All three of the parameters $\tau(G)$, $\alpha(G)$ and $\alpha^*(G)$ are evaluations of the Tutte polynomial $T_G(x, y)$; we have

$$\tau(G) = T_G(1,1), \qquad \alpha(G) = T_G(2,0), \qquad \alpha^*(G) = T_G(0,2),$$

As the Tutte polynomial is naturally defined for all *matroids*, the conjecture can directly be extended to any matroid M where it becomes

$$T_M(1,1) \leq \max\{T_M(2,0), T_M(0,2)\}$$

even though there are no obvious combinatorial interpretations of $T_M(2,0)$ or $T_M(0,2)$ for general matroids.

A matroid is called a *paving matroid* if its smallest circuit has size at least equal to its rank, and it is generally presumed (though not proved) that asymptotically almost all matroids are paving matroids (this is conjectured in [6]). Chávez-Lomelí, Merino, Noble and Ramírez-Ibáñez [1] proved (among other results) that for a coloopless paving matroid M, the function $T_M(1-x, 1+x)$ is *convex* in the region $-1 \leq x \leq 1$, thus proving that paving matroids satisfy the Merino-Welsh conjecture.

So the conjecture is proved for the vast class of paving matroids, and now for the tiny class of series-parallel graphs, but for general graphs and matroids, every variant of the conjecture remains open. While it may be possible to develop bounds for dense or sparse matroids similar to those found by Thomassen for graphs, the heart of the problem lies in the case where the rank is about half the number of elements.

References

[1] L. E. Chávez-Lomelí, C. Merino, S. D. Noble, M. Ramírez-Ibáñez, Some inequalities for the Tutte polynomial, *European J. Combin.* **32** (2011) 422–433. doi:10.1016/j.ejc.2010.11.005.

[2] R. Conde, C. Merino, Comparing the number of acyclic and totally cyclic orientations with that of spanning trees of a graph, *Int. J. Math. Comb.* **2** (2009) 79–89.

[3] R. J. Duffin, Topology of series-parallel networks, *J. Math. Anal. Appl.* **10** (1965) 303–318.

[4] B. Jackson, An inequality for Tutte polynomials, *Combinatorica* **30** (2010) 68–81.

[5] C. Merino, D. J. A. Welsh, Forests, colorings and acyclic orientations of the square lattice, *Ann. Comb.* **3** (1999) 417–429, doi:10.1007/BF01608795

[6] D. Mayhew, M. Newman, D. Welsh, G. Whittle, On the asymptotic proportion of connected matroids, *European J. Combin.* **32** (6) (2011) 882–890. doi:10.1016/j.ejc.2011.01.016.

[7] J. Sedláček, On the spanning trees of finite graphs, *Časopis Pěst. Mat.* **91** (1966) 221–227.

[8] C. Thomassen, Spanning trees and orientations of graphs, *J. Comb.* **1** (2010) 101–111. doi:10.4310/JOC.2010.v1.n2.a1.

A Code

This appendix contains the Mathematica code used to check the assertions made in Table 2 and Table 3 regarding the replaceability of the series and parallel connections of each of the graphs in Table 1. In this code, the parameters for each graph are represented as a vector

$$(\tau(G), \tau_2(G), \alpha(G), \alpha_2(G), \alpha_2^*(G), \alpha^*(G))$$

and the functions `ser` and `par` produce the parameters of the series connection and parallel connection respectively.

```
ser::usage="Returns the parameters for
                the series connection of g and h"

ser[g_,h_] := {
  g[[1]]h[[1]],
  g[[1]]h[[2]]+g[[2]]h[[1]],
  g[[3]]h[[3]],
  g[[3]]h[[3]]-(g[[3]]-g[[4]])(h[[3]]-h[[4]])/2,
  g[[5]]h[[5]]-(g[[5]]-g[[6]])(h[[5]]-h[[6]])/2,
  g[[6]]h[[6]]};

par::usage="Returns the parameters for the
                    parallel connection of g and h"
```

```
par[g_,h_] := {
  g[[1]]h[[2]]+g[[2]]h[[1]],
  g[[2]]h[[2]],
  g[[3]]h[[3]]-(g[[3]]-g[[4]])(h[[3]]-h[[4]])/2,
  g[[4]]h[[4]],
  g[[5]]h[[5]],
  g[[5]]h[[5]]-(g[[5]]-g[[6]])(h[[5]]-h[[6]])/2};


spdual::usage="Returns the parameters for the sp-dual of g"

spdual[g_] := {
 g[[2]],g[[1]],g[[5]],g[[6]],g[[3]],g[[4]]};
```

The function `replaces` is a boolean function testing whether the first argument can be replaced by the second.

```
replaces::usage="Returns true if g can
                      be replaced by h, else false"

replaces[g_,h_] := Module[{t1,t2,t3},
  t1 = Max[ g[[1]]/h[[1]], g[[2]]/h[[2]] ];
  t2 = If[h[[4]]==0,g[[3]]/h[[3]],Min[g[[3]]/h[[3]],g[[4]]/h[[4]]]];
  t3 = If[h[[6]]==0,g[[5]]/h[[5]],Min[g[[5]]/h[[5]],g[[6]]/h[[6]]]];
  t1^2-t2 t3 <= 0]
```

Finally, `gs` is the collection of all the graphs (other than K_2) in Table 1 ordered such that `gs[[x]]` is the graph `x` in the table.

```
k2 = {1,1,2,0,2,0};
g0 = k2;
g1 = par[g0,g0]; g2 = ser[g0,g0];
g3 = par[g0,g2]; g4 = ser[g0,g1];
g5 = par[g0,g1]; g6 = ser[g0,g2];
g7 = ser[g1,g1]; g8 = par[g2,g2];
g9 = par[g1,g2]; g10 = ser[g1,g2];
g11 = ser[g1,g5]; g12 = par[g2,g6];
g13 = par[g2,g8]; g14 = ser[g1,g7];
g15 = ser[g5,g8]; g16 = par[g6,g7];
g17 = par[g2,g12]; g18 = ser[g1,g11];


gs = {g1,g2,g3,g4,g5,g6,g7,g8,g9,g10,g11,g12,g13,g14,g15,g16,g17,g18}
```

It is now straightforward to confirm any of the assertions contained in the tables; for example, one entry in the table claims that $G_1 \oplus_S G_6$ is replaceable by G_6 itself. This can easily be confirmed by ensuring that the statement

```
replaces[ ser[ gs[[1]], gs[[6]] ], gs[[6]] ]
```

returns `True`, and similarly for all the other entries.