NP-Completeness of Perfect Matching Index of Cubic Graphs

Martin Škoviera

Department of Computer Science, Comenius University, Bratislava, Slovakia

Peter Varša

Department of Computer Science, Comenius University, Bratislava, Slovakia

Abstract

The perfect matching index of a cubic graph G, denoted by $\pi(G)$, is the smallest number of perfect matchings needed to cover all the edges of G; it is correctly defined for every bridgeless cubic graph. The value of $\pi(G)$ is always at least 3, and if G has no 3-edge-colouring, then $\pi(G) \geq 4$. On the other hand, a long-standing conjecture of Berge suggests that $\pi(G)$ never exceeds 5. It was proved by Esperet and Mazzuoccolo [J. Graph Theory 77 (2014), 144–157] that it is NP-complete to decide for a 2-connected cubic graph whether $\pi(G) \leq 4$. A disadvantage of the proof (noted by the authors) is that the constructed graphs have 2-cuts. We show that small cuts can be avoided and that the problem remains NP-complete even for nontrivial snarks – cyclically 4-edge-connected cubic graphs of girth at least 5 with no 3-edge-colouring. Our proof significantly differs from the one due to Esperet and Mazzuoccolo in that it combines nowhere-zero flow methods with elements of projective geometry, without referring to perfect matchings explicitly.

2012 ACM Subject Classification Mathematics of computing → Matchings and factors; Theory of computation → Problems, reductions and completeness; Mathematics of computing → Graph coloring

Keywords and phrases cubic graph, edge colouring, snark, perfect matching, covering, NP-completeness

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.56

Funding This research was supported by the grants APVV-19-0308 and VEGA 1/0743/21.

Acknowledgements The authors wish to thank Edita Máčajová for helpful discusions.

1 Introduction

It is well known [7] that every bridgeless cubic graph has a perfect matching that contains an arbitrarily preassigned edge. As a consequence, each such graph can be expressed as a union of a collection of its perfect matchings. The smallest number of perfect matchings needed for this purpose is its *perfect matching index*, denoted by $\pi(G)$. Although no constant bound on $\pi(G)$ is known, a fascinating conjecture of Berge (see [8]) suggests that five perfect matchings should do for every bridgeless cubic graph G.

Clearly, $\pi(G) = 3$ if and only if G is 3-edge-colourable, so if G has chromatic index 4, the value of $\pi(G)$ is at least 4. Understanding the cubic graphs that require more than four perfect matchings to cover their edges is fundamental for any approach that might lead to proving or disproving Berge's conjecture. However, nontrivial examples of cubic graphs with perfect matching index at least 5 appear to be very rare and are difficult to find. In the list comprising all 64 326 024 *nontrivial snarks* – cyclically 4-edge-connected cubic graphs of girth at least 5 with no 3-edge-colouring – on up to 36 vertices, generated by Brinkmann et al. [2], there are only two graphs that cannot be covered with four perfect matchings: the Petersen graph and the *windmill snark* W_{34} on 34 vertices displayed in Figure 1. The latter snark provides the starting point for several infinite families of snarks with $\pi \geq 5$, see [1, 2, 3, 5].

39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022).
Editors: Petra Berenbrink and Benjamin Monmege; Article No. 56; pp. 56:1–56:12
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

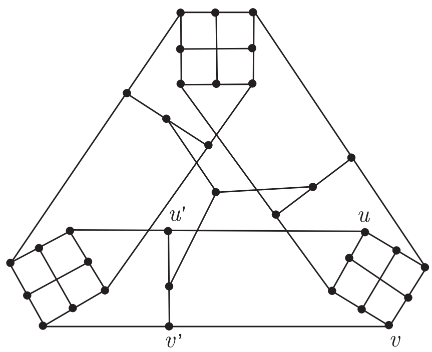

Figure 1 Windmill snark W_{34} on 34 vertices.

It transpires that the structure of graphs with perfect matching index at least 5 is far from being simple. In fact, deciding whether $\pi(G) \leq 4$ is an NP-complete problem, which was proved by Esperet and Mazzuoccolo [3] in 2014. However, as the authors write in [3], "*The gadgets used in the proof of NP-completeness have many* 2*-edge-cuts, so our* [first] *result does not say much about* 3*-edge-connected cubic graphs.*" In particular, they leave the NP-completeness problem open for nontrivial snarks. In this context it may be useful to realise that it is the class of nontrivial snarks which is particularly important for the problem. Indeed, several profound conjectures in graph theory, including the celebrated cycle double cover conjecture and the shortest cycle cover conjecture (also known as the 7/5-conjecture), can be reduced to nontrivial snarks with perfect matching index at least 5, see Steffen [9, Theorem 3.1].

The purpose of this contribution is to prove that deciding whether $\pi(G) \leq 4$ remains NP-complete even in the family of nontrivial snarks. Like the proof of NP-completeness due to Esperet and Mazzuoccolo [3], our proof employs reduction to 3-edge-colourability, which is known to be NP-complete by a result of Holyer [4]. On the other hand, its characteristic feature consists in avoiding direct use of perfect matchings, replacing them with nowhere-zero flows possessing an additional geometric structure within the 3-dimensional projective space $\mathbb{P}_3(\mathbb{F}_2)$ over the 2-element field. Although our methods heavily depend on the theory of tetrahedral flows developed in [5], we include all the necessary definitions and results from [5] to make the present paper self-contained.

We finish this section with a brief list of basic definitions used throughout the paper.

Our graphs will be mostly cubic, simple, although parallel edges and loops are not automatically excluded. A *circuit* is a connected 2-regular graph. A graph G is said to be *cyclically k-edge-connected* if the removal of fewer than k edges from G cannot create a graph with at least two components containing circuits. An edge cut S in G that separates two circuits from each other is *cycle-separating*. A (proper) *edge-colouring* of a graph G is a mapping from the edge set of G to a set of colours such that adjacent edges receive distinct colours. A *k-edge-colouring* is an edge colouring using k colours. A 2-connected cubic graph that does not admit a 3-edge-colouring is called a *snark*. A snark is *nontrivial* if it is cyclically 4-edge-connected and has no circuits of length smaller than 5.

For for more details and general context we refer the reader to [5].

2 Main results

Our point of departure is the result of Esperet and Mazzuoccolo [3, Theorem 2] which establishes NP-completeness of deciding whether $\pi(G) \leq 4$ in the class of bridgeless cubic graphs. We briefly summarise their proof.

Let G be an arbitrary bridgeless cubic graph. Inflate every vertex of G to a triangle, thereby producing a graph G'. Next, construct a new cubic graph G'' as follows. Take the *Tietze graph* T, which arises from the Petersen graph by inflating one of the vertices to a triangle, and remove an edge e lying on the triangle of T. For each edge x lying on a triangle of G' take a copy T_x of $T - e$, remove x from G', and connect the two 2-valent vertices of $G' - x$ to the two 2-valent vertices of T_x in such a way that 3-regularity is restored. The graph G'' is now obtained by repeating the just described procedure with each edge of G' lying on a triangle, see Figure 2. The substantial part of the proof of NP-completeness presented in [3] consists in checking that $\pi(G'') = 4$ if and only if G is 3-edge-colourable.

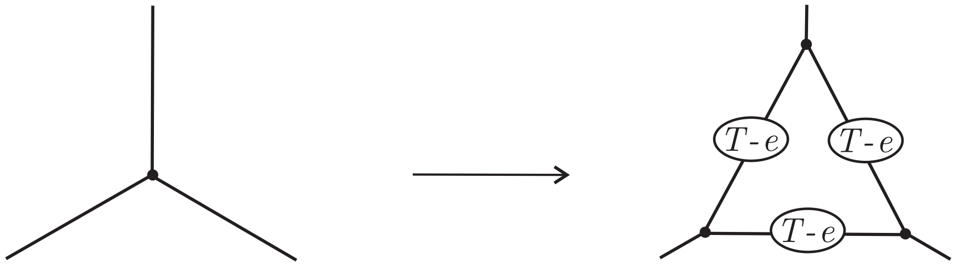

Figure 2 The construction of Esperet and Mazzuoccolo; $T - e$ denotes the Tietze graph with one edge removed.

.

Clearly, each copy T_x of $T - e$ is separated from the rest of G'' by a 2-edge-cut, so G'' has a plenty of 2-cuts. Considering the importance of nontrivial snarks for Berge's conjecture and other related conjectures it is a legitimate question to ask whether small cuts in the proof of NP-completeness can be avoided. Our main result answers this question in the positive.

▶ **Theorem 1.** *Deciding whether a nontrivial snark G satisfies $\pi(G) \leq 4$ is an* NP*-complete problem.*

The previous theorem is a direct consequence of the following more detailed statement.

▶ **Theorem 2.** *For every* 2*-connected cubic graph G of order n one can construct a derived graph $G^\sharp$ on* $102n$ *vertices which is a nontrivial snark. Moreover, $\pi(G^\sharp) = 4$ if and only if G is* 3*-edge-colourable.*

The derived graph $G^\sharp$ will be constructed by substituting the vertices of G with "fat vertices" (vertex gadgets, which we call *tripoles*) and the edges of G with "fat edges" (edge gadgets, which we call *dipoles*). Dipoles and tripoles, and more generally *multipoles*, are structures similar to graphs: like graphs, they consist of vertices and edges, each edge having two *half-edges*. In addition to proper edges, multipoles may contain *dangling edges*, with only one half-edge incident with a vertex, and even *isolated edges*, which are not incident with any vertex at all. Thus a dangling edge has one free half-edge while an isolated edge has two free half-edges. A *dipole* is a multipole whose free half-edges are partitioned into two subsets, the *input connector* and the *output connector*, while a *tripole* has its free half-edges distributed into three *connectors*. All multipoles in this paper are *cubic*, that is to say, every vertex is incident with exactly three half-edges. Moreover, all connectors will be of size 2.

We now describe the conctruction of $G^\sharp$ in detail. Let G be an arbitrary 2-connected cubic graph.

For vertex gadgets we use copies of the tripole W_0 consisting of three isolated edges a, b, and c (and no vertices) in which each connector consists of half-edges that belong to two distinct edges; the tripole W_0 is shown in Figure 3.

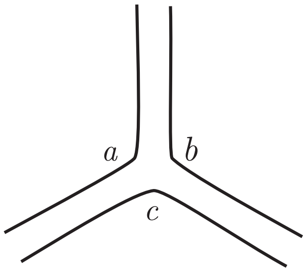

Figure 3 The tripole W_0.

For edge gadgets we take copies of a dipole H_{68} on 68 vertices; it is somewhat more complicated to describe. First, take the windmill snark W_{34} of order 34 shown in Figure 1 and severe the edges uu' and vv' indicated in the figure thereby producing four dangling edges distributed into two pairs. The resulting dipole, which we denote by D_{34}, is displayed in Figure 4. The input connector $\{d, k\}$ is formed from the half-edges belonging to the dangling edges incident with u and v, respectively, while the output connector $\{q, m\}$ is formed from those incident with u' and v'. To finish the construction of H_{68}, take two copies of D_{34} and weld the half-edges of their input connectors identically, that is to say, join d to d and k to k. The result is the dipole on 68 vertices whose both connectors are copies of the output connector of D_{34} – this is the required H_{68}.

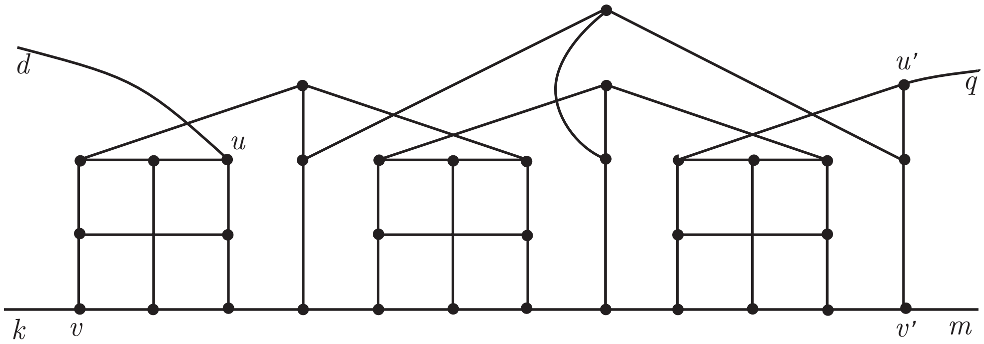

Figure 4 The dipole D_{34}.

Finally, we assemble $G^\sharp$ from the building blocks.

1. For each vertex v of G we take a cyclic permutation R_v of the edges incident with v. The collection $R = (R_v)_{v \in V(G)}$ is the *rotation system* for G, as it is generally known in topological graph theory [6]. The choice of R is irrelevant, yet R is important for keeping the track of the construction. (We will not pursue the topological connection any further. No knowledge of topological graph theory is therefore required.)
2. Next, we create the corresponding vertex gadget X_v as a copy of the tripole W_0. We associate each connector of X_v with an edge of G incident with v: for each such edge z we let $\{z', z''\}$ be the corresponding connector and we further require that the half-edges z' and $(R_v(z))''$ constitute one edge, see Figure 5.

3. For any edge $e = uv$ of G incident with v we create the corresponding edge gadget Y_e as a copy of H_{68} and associate its connectors with the endvertices of e.
4. Finally, we glue Y_e to X_v. If $\{q', m'\}$ is the connector of Y_e associated with v, we take the connector $\{e', e''\}$ of X_v associated with e and attach q' to e' and m' to e''. As soon as the gluing procedure is performed with each edge e and with both its endvertices, the construction of the derived graph $G^\sharp$ is completed. It is easy to see that if G has n vertices, then $G^\sharp$ has $102n$ vertices. It follows that $G^\sharp$ can be constructed from G in a polynomial number of steps with respect to the number of vertices of G.

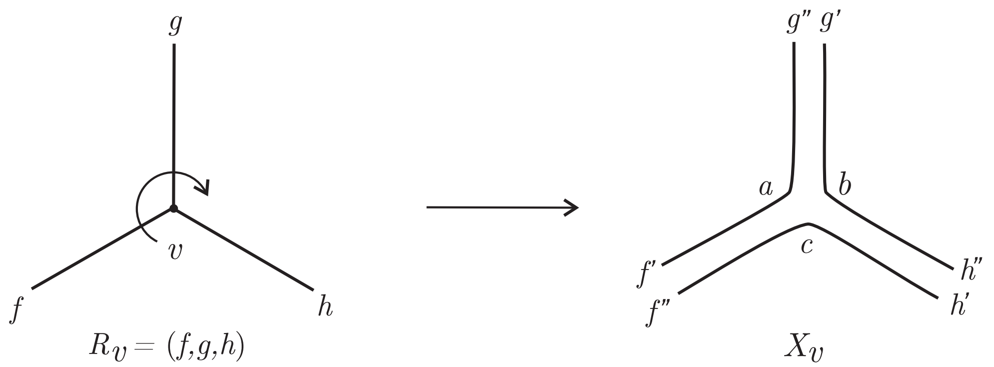

Figure 5 Substituting a vertex of G with a vertex gadget.

Recall that deciding whether a cubic graph is 3-edge-colourable is a well known NP-complete problem [4]. Therefore, from the construction of $G^\sharp$ it is clear that to prove Theorem 1 it suffices to show that $\pi(G^\sharp) = 4$ if and only if G is 3-edge-colourable.

3 Geometric background

In this section we prepare the machinery required for the proof of Theorem 1. As already indicated, the main idea of our proof consists in representing covers of a cubic graph with four perfect matchings by flows possessing a certain geometric structure. We now explain this idea in detail.

First of all, recall that a *flow* on a graph G is a function $\phi\colon E(G) \to A$, with values in an abelian group A, together with an orientation of G, such that the following property is fulfilled: at each vertex of G the sum of all incoming values equals the sum of all outgoing ones (*Kirchhoff's law*). More specifically, ϕ is an *A-flow*. A flow ϕ is *nowhere-zero* if $\phi(e) \neq 0$ for each edge e of G. The choice of an orientation for a flow is immaterial because the orientation of any edge can be reversed and its value can be replaced with the inverse without violating the Kirchhoff law. Furthermore, if $x = -x$ for every $x \in A$, one can ignore orientation altogether. This is possible precisely when A is isomorphic to an elementary abelian 2-group $\mathbb{Z}_2^n$.

Let G be a cubic graph that admits a covering $\mathcal{C} = \{P_1, P_2, P_3, P_4\}$ of its edges with four perfect matchings; note that the matchings need not be pairwise distinct. Clearly, $\mathcal{C}$ can be unambiguously represented by the mapping

$$\xi_{\mathcal{C}}\colon E(G) \to \mathbb{Z}_2^4$$

where the i-th coordinate of $\xi_{\mathcal{C}}(e)$ equals $1 \in \mathbb{Z}_2$ whenever the edge e does not belong to the perfect matching P_i. It is not difficult to see that $\xi_{\mathcal{C}}$ is a nowhere-zero $\mathbb{Z}_2^4$-flow on G.

In order to reveal important properties of this flow it is convenient to identify the set $\mathbb{Z}_2^4 - \{0\}$ with the point set of the 3-dimensional projective space $PG(3,2)$ over the 2-element field. Recall that the *n-dimensional projective space* $PG(n,2) = \mathbb{P}_n(\mathbb{F}_2)$ over the 2-element field $\mathbb{F}_2$ is an incidence geometry whose *points* can be identified with the nonzero vectors of the $(n+1)$-dimensional vector space $\mathbb{F}_2^{n+1}$ and whose *lines* are formed by the triples $\{x, y, z\}$ of points such that $x + y + z = 0$. The *3-dimensional projective space* $PG(3,2)$ consists of 15 points and 35 lines. The crucial observation concerning the flow $\xi_\mathcal{C}$ is that for every vertex v of G the values assigned by $\xi_\mathcal{C}$ to the edges incident with v form a line of $PG(3,2)$. The theory which we now outline generalises this observation.

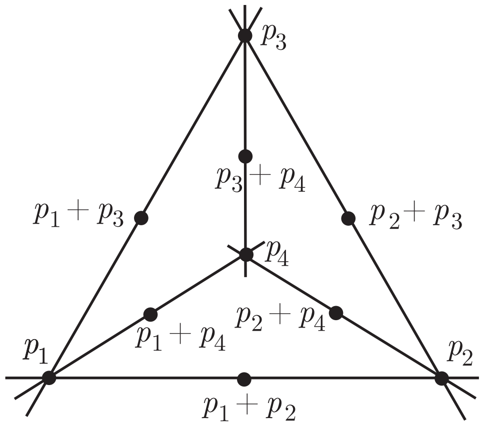

Figure 6 The tetrahedron in $PG(3,2)$ spanned by points p_1, p_2, p_3, and p_4.

We start with the necessary geometric concepts. A *tetrahedron* $T = T(p_1, p_2, p_3, p_4)$ in $PG(3,2)$ is a configuration consisting of ten points and six lines spanned by a set $\{p_1, p_2, p_3, p_4\}$ of four points of $PG(3,2)$ in general position; the latter means that the set $\{p_1, p_2, p_3, p_4\}$ constitutes a basis of the vector space $\mathbb{F}_2^4$. These four points are the *corner points* of T. Any two distinct corner points $c_1, c_2 \in \{p_1, p_2, p_3, p_4\}$ belong to a unique line $\ell = \{c_1, c_2, c_1 + c_2\}$ of T whose third point $c_1 + c_2$ is the *midpoint* of ℓ. Each line of T is uniquely determined by its midpoint. The tetrahedron $T(p_1, p_2, p_3, p_4)$ is depicted in Figure 6.

Given a tetrahedron T, a *T-flow* on a cubic graph G is a mapping $\phi\colon E(G) \to P(T)$ from the edge set of G to the *point set* $P(T)$ of T such that for each vertex v of G the three edges e_1, e_2, and e_3 incident with v receive values that form a line of T. The latter means that $\phi(e_1) + \phi(e_2) + \phi(e_3) = 0$, which amounts to the Kirchhoff law for ϕ. Thus a T-flow is indeed a flow. A *tetrahedral flow* on G is a T-flow for some tetrahedron T in $PG(3,2)$. Note that any tetrahedal flow is also a proper edge colouring, which is why we occasionally refer to a T-flow as a colouring.

The next theorem provides a characterisation of cubic graphs with perfect matching index at most 4 in terms of tetrahedral flows. Due to the result of Esperet and Mazzuocollo [3], this characterisation is not efficient in the strict algorithmic sense, nevertheless, it is very useful.

▶ **Theorem 3.** *A cubic graph G can have its edges covered with four perfect matchings if and only if it admits a tetrahedral flow. Moreover, there exists a one-to-one correspondence between coverings of G with four perfect matchings and T-flows, where T is an arbitrary fixed tetrahedron in $PG(3,2)$.*

The way in which we apply Theorem 3 to the investigation of cubic graphs with perfect matching index at least 5 is based on the following idea: Suppose that a graph G in question has a cycle-separating 4-edge-cut S. Severing the edges of S produces two multipoles X_1 and X_2 with four dangling edges each, which we arrange into dipoles correspondingly. Choosing an input connector in each of them permits us to analyse how pairs of points of a tetrahedron in $PG(3,2)$ are transformed via a tetrahedral flow from the input connector to the output, and to check whether the tetrahedral flows through X_1 are always in conflict with the tetrahedral flows through X_2. This is why we need to examine which pairs of points can occur on the connectors of a dipole equipped with a tetrahedral flow.

For the rest of this section fix an arbitrary tetrahedron $T(p_1, p_2, p_3, p_4) = T$ in $PG(3,2)$. Consider a dipole $X = X(I, O)$ with input connector $I = \{g_1, g_2\}$ and output connector $O = \{h_1, h_2\}$. Given subsets $\{x, y\}$ and $\{x', y'\}$ of $P(T)$, we say that X *has a transition*

$$\{x, y\} \to \{x', y'\}$$

or that $\{x, y\} \to \{x', y'\}$ is a *transition through* X, if there exists a T-flow ϕ on X such that $\{\phi(g_1), \phi(g_2)\} = \{x, y\}$ and $\{\phi(h_1), \phi(h_2)\} = \{x', y'\}$. By the Kirchhoff law, for each transition $\{x, y\} \to \{x', y'\}$ through X we have $x + y = x' + y'$. This common value is called the *trace* of the transition; it can be any element of $\mathbb{Z}_2^4$.

In order to get better insight into possible transitions through a dipole it is useful to classify pairs of points of T according to their geometric shape. We say that two sets A and B of points of a tetrahedron T have the *same shape* if there exists a collineation (in other words, an automorphism) of $PG(3,2)$ that preserves T and takes A to B. A *geometric shape*, or simply a *shape*, is an equivalence class of all point sets having the same shape. The *shape* of a set of points of T is a geometric shape it belongs to. It is proved in [5] that each pair $\{x, y\}$ of points of T, where possibly $x = y$, falls into one of the following seven shapes: *line segment* `ls`, *half-line* `hl`, *angle* `ang`, *altitude* `alt`, *axis* `ax`, *double corner point* `dc`, and *double midpoint* `dm`. Their typical representatives are, respectively, the following pairs: $\{p_1, p_2\}$, $\{p_1, p_1 + p_2\}$, $\{p_1 + p_2, p_1 + p_3\}$, $\{p_1, p_2 + p_3\}$, $\{p_1 + p_2, p_3 + p_4\}$, $\{p_1, p_1\}$, and $\{p_1 + p_2, p_1 + p_2\}$. The set

$$\mathbf{\Sigma} = \{\mathtt{ls}, \mathtt{hl}, \mathtt{ang}, \mathtt{alt}, \mathtt{ax}, \mathtt{dc}, \mathtt{dm}\}$$

comprises all shapes of point pairs of T.

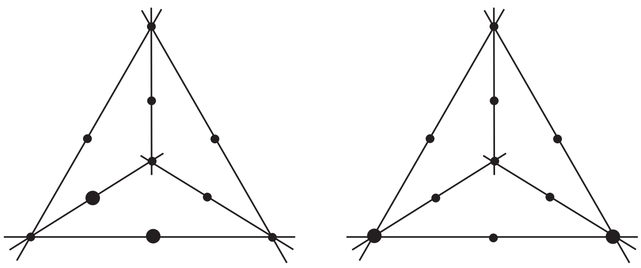

Figure 7 An angle (left) and a line segment (right).

Each transition $\{x, y\} \to \{x', y'\}$ through a dipole X between point pairs induces a transition between their shapes. To be more precise, for elements `s` and `t` of $\mathbf{\Sigma}$ we say that X *has a transition*

$$\mathtt{s} \to \mathtt{t}$$

if X has a transition $\{x, y\} \to \{x', y'\}$ such that `s` is the shape of $\{x, y\}$ and `t` is the shape of $\{x', y'\}$. It can be shown (see [5, Theorem 5.1]) that all transitions through any dipole have the form `s` $\to$ `s` except possibly the transitions `ls` $\to$ `ang` or `ang` $\to$ `ls`, and the transitions `dc` $\to$ `dm` and `dm` $\to$ `dc` between the degenerate point pairs. The two exceptional non-degenerate shapes `ang` and `ls` are illustrated in Figure 7.

For a detailed account of the theory of tetrahedral flows we refer the reader to [5].

4 The proof

In this section we prove Theorems 1 and 2. As previously mentioned, it suffices to prove the latter.

Our first step is to stablish several useful properties of the edge gadget, the dipole H_{68}. Recall that H_{68} consists of two copies of the dipole D_{34} whose input connectors have been identically glued together. The dipole D_{34} is the smallest example of what in [5] is called an *extended Halin dipole*. Every extended Halin dipole arises from a Halin snark (defined in [5]) by severing two edges in a manner similar to the construction of D_{34} from W_{34} (the latter being the smallest Halin snark). The crucial property of an extended Halin dipole, proved in [5, Theorem 8.8], is that every transition through it has the form

`ang` $\to$ `ls`.

One such transition through D_{34} is displayed in Figure 8. The flow is encoded as follows: the label i stands for the corner point p_i of the tetrahedron $T(p_1, p_2, p_3, p_4)$ while the label ij stands for the midpoint $p_i + p_j$. We will use this encoding in the rest of this paper.

The fact that `ang` $\to$ `ls` is the only possible transition of shapes through D_{34} can be checked directly, but without deeper involvement of the theory outlined in the previous section the proof would be quite tedious.

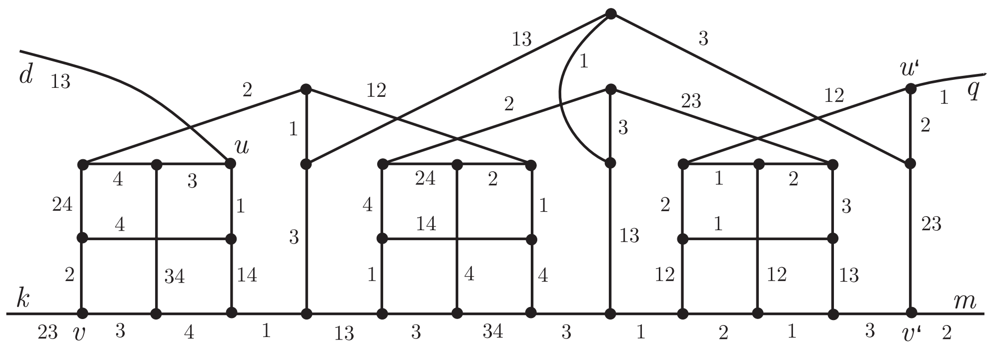

Figure 8 Transition `ang` $\to$ `ls` through D_{34}; encoding of colours: $i \mapsto p_i$, $ij \mapsto p_i + p_j$.

We need the following lemma.

▶ **Lemma 4.** *Every transition $\{x, y\} \to \{x', y'\}$ through the dipole H_{68} has the form*

`ls` $\to$ `ls`.

Such a transition exists for any line segment $\{x, y\}$ in T. Moreover, $\{x, y\} = \{x', y'\}$, and the two points may occur on the connectors of H_{68} in any order.

Proof. Recall that H_{68} is created from two copies of the dipole D_{34} whose input connectors are joined identically. Since $\texttt{ang} \to \texttt{ls}$ is the only possible transition through D_{34}, it follows that the only transition through H_{68} is of the form $\texttt{ls} \to \texttt{ls}$. If $\{x, y\} \to \{x', y'\}$ is any such transition, then $x + y = x' + y'$ by the Kirchhoff law. Each line of a tetrahedron is uniquely determined by its midpoint, so $x + y$ and $x' + y'$ must be midpoints of the same line, and in turn the line segments $\{x, y\}$ and $\{x', y'\}$ must be identical.

Now, let ψ^+ denote the flow displayed in Figure 8. In terms of ordered pairs of points, ψ^+ induces the transition $(p_1 + p_3, p_2 + p_3) \to (p_1, p_2)$, where the input pair stands for $(\psi^+(d), \psi^+(k))$ and the output pair stands for $(\psi^+(q), \psi^+(m))$. The dipole D_{34} has another tetrahedral flow ψ^-, namely one that represents the transition $(p_1 + p_3, p_2 + p_3) \to (p_2, p_1)$ with the output values swapped. This flow can easily be obtained from ψ^+ by interchanging p_1 and p_2 on the unique path in D_{34} which starts with the half-edge q of the output connector, leads through two internal edges, one incident with u' and the other incident with v', and terminates in the output connector with the half-edge m. The flows ψ^+ and ψ^- can be combined into four distinct tetrahedral flows on H_{68} which transform the line segment $\{p_1, p_2\}$ into itself in such a way that both the input pair and the output pair occur in any preassigned ordering. By symmetry, the same is true for any other line segment of T. The lemma follows. ◀

Now we are ready to prove Theorem 2.

Proof of Theorem 2. We first prove that $G^\sharp$ is not 3-edge-colourable irrespectively of the choice of G. Observe that in the language of tetrahedral flows a cubic graph is 3-edge-colourable if and only if it admits a tetrahedral flow using a single line of the tetrahedron. As we already know, every tetrahedral flow on D_{34} induces a transition of the form $\texttt{ang} \to \texttt{ls}$. Since the points of an angle do not lie on the same line of T, every tetrahedral flow on D_{34} must use points of at least two lines of the tetrahedron. Thus D_{34} is not 3-edge-colourable, and consequently neither is $G^\sharp$.

Next we prove that $G^\sharp$ is a nontrivial snark. Obviously, the girth of $G^\sharp$ is 5. To see that G is cyclically 4-edge-connected it is sufficient to realise that the underlying graph G is 2-connected and that each edge gadget arises from a cyclically 4-edge-connected cubic graphs by severing two independent edges. A straightforward case analysis, which we leave to the reader, shows that $G^\sharp$ has no k-edge-cut with $k < 4$ that separates a subgraph containing a cycle from the rest of $G^\sharp$. Summing up, $G^\sharp$ is a nontrivial snark.

We proceed to proving that $\pi(G^\sharp) = 4$ if and only if G is 3-edge-colourable. We do it in two steps.

▷ Claim. If $\pi(G^\sharp) = 4$, then G is 3-edge-colourable.

Proof. Assume that $\pi(G^\sharp) = 4$. By Theorem 3, $G^\sharp$ admits a T-flow ϕ where $T = T(p_1, p_2, p_3, p_4)$. For each edge e of G let $\phi'(e)$ denote the trace of the transition through the edge gadget Y_e of $G^\sharp$ induced by ϕ. Since all edge gadgets of $G^\sharp$ are copies of H_{68}, for each edge e of G the value $\phi'(e)$ is a midpoint of T.

Consider an arbitrary vertex v of G, and let e_1, e_2, and e_3 be the edges incident with v. By Kirchhoff's law, the outflow from the vertex gadget X_v must be 0, which in turn implies that $\phi'(e_1) + \phi'(e_2) + \phi'(e_3) = 0$. The values $\phi'(e_1)$, $\phi'(e_2)$, and $\phi'(e_3)$ are nonzero and therefore pairwise distinct. It follows that $\phi' \colon E(G) \to \mathbb{Z}_2^4$ is a nowhere-zero flow and the same time a proper edge colouring. As a colouring, ϕ uses (at most) six colours, the midpoints of T.

Next we prove that $\phi'(e_1)$, $\phi'(e_2)$, and $\phi'(e_3)$ are midpoints of the same triangle of T. By a *triangle* we mean a configuration of three lines of T spanned by three distinct corner points of T. Clearly, three distinct lines of T form a triangle if and only if any two of them intersect, but the intersection of all three is empty. Set $\phi'(e_i) = m_i$ for each $i \in \{1,2,3\}$; as already mentioned, m_1, m_2, and m_3 are pairwise distinct. Let ℓ_i denote the unique line of T containing m_i. We first prove that any two of the lines ℓ_1, ℓ_2, and ℓ_3 intersect. Suppose not, and assume that, say, ℓ_1 and ℓ_2 are disjoint. The position of any pair of disjoint lines of T implies that exists a permutation σ of $\{1,2,3,4\}$ such that ℓ_1 is the line through $p_{\sigma(1)}$ and $p_{\sigma(2)}$, and ℓ_2 is the line through $p_{\sigma(3)}$ and $p_{\sigma(4)}$. Since $m_1 + m_2 + m_3 = 0$, we conclude that

$$m_3 = p_{\sigma(1)} + p_{\sigma(2)} + p_{\sigma(3)} + p_{\sigma(4)} = p_1 + p_2 + p_3 + p_4,$$

which is not a midpoint of any line of T. Therefore ℓ_1, ℓ_2, and ℓ_3 are distinct pairwise intersecting lines. Next we prove that $\ell_1 \cap \ell_2 \cap \ell_3 = \emptyset$. Indeed, if there is a point $p \in \ell_1 \cap \ell_2 \cap \ell_3$, then p is a corner point of T and ℓ_1, ℓ_2, and ℓ_3 are the three lines of T containing p. But then $m_1 + m_2 + m_3 = p_1 + p_2 + p_3 + p_4 \neq 0$, which is impossible. The only remaining possibility is that ℓ_1, ℓ_2, and ℓ_3 form a triangle. The latter means that there exist three distinct corner points c_1, c_2, and c_3 of T such that $m_1 = c_2 + c_3$, $m_2 = c_1 + c_3$, and $m_3 = c_1 + c_2$.

Now we are ready to produce a proper 3-edge-colouring of G. Let us define the mapping ψ from the set of all midpoints of T to the set $\{1,2,3\}$ as follows:

$$\begin{aligned} p_1 + p_2,\ p_3 + p_4 &\quad\mapsto\quad 1,\\ p_1 + p_3,\ p_2 + p_4 &\quad\mapsto\quad 2,\\ p_1 + p_4,\ p_2 + p_3 &\quad\mapsto\quad 3. \end{aligned}$$

Since every triangle is determined by three distinct corner points c_1, c_2 and c_3 of T, its midpoints $c_1 + c_2$, $c_2 + c_3$, and $c_1 + c_3$ receive from ψ three distinct values. In other words, $\psi\phi'$ is a proper 3-edge-colouring of G, which establishes the claim. ◁

▷ Claim. If G is 3-edge-colourable, then $\pi(G^\sharp) = 4$.

Proof. Assume that G is 3-edge-colourable. By Theorem 3, it is sufficient to find a tetrahedral flow on $G^\sharp$. Our aim is to construct a tetrahedral flow $\gamma^\sharp$ of $G^\sharp$ by departing from a 3-edge-colouring $\gamma\colon E(G) \to \{1,2,3\}$.

First of all, we colour the vertex gadgets. With respect to the chosen rotation system R for G the vertices of G fall into two types depending on whether the cyclic order of colours around the vertex is $(1,2,3)$ (*Type 1*) or $(1,3,2)$ (*Type 2*). Consider a vertex v of G, which is incident with edges f, g, and h, and let X_v be the corresponding vertex gadget. Recall that the connectors of X_v are $\{f', f''\}$, $\{g', g''\}$, and $\{h', h''\}$, and that for each edge x of G incident with v, the free half-edges x' and $(R(x))''$ constitute one edge of X_v, see Figure 5. Let a, b, and c be the edges of X_v that have the half-edges f', g', and h', respectively.

Without loss of generality we may assume that $\gamma(f) = 1$. We intend to colour the edges of X_v with three distinct corner points of the tetrahedron T, say p_1, p_2, and p_3, in such a way that the connector $\{f', f''\}$ receives colours from the line segment $\{p_2, p_3\}$. This choice implies that under the colouring $\gamma^\sharp$ the edge b must receive colour p_1. The colours of the remaining two edges will depend on the type of v. If v is Type 1, then $\gamma(g) = 2$, and $\gamma(h) = 3$, and we set $\gamma^\sharp(a) = p_3$, $\gamma^\sharp(b) = p_1$, and $\gamma^\sharp(c) = p_2$, see Figure 9. If v is Type 2, then $\gamma(g) = 3$, and $\gamma(h) = 2$, and we set $\gamma^\sharp(a) = p_2$, $\gamma^\sharp(b) = p_1$, and $\gamma^\sharp(c) = p_3$, see Figure 10. We have thus coloured every vertex gadget of X_v in such a way that the connector $\{x', x''\}$ of X_v corresponding to the edge x of G incident with v receives colours from the line segment $\{p_1, p_2, p_3\} - \{p_i\}$ if and only if $\gamma(x) = i \in \{1,2,3\}$.

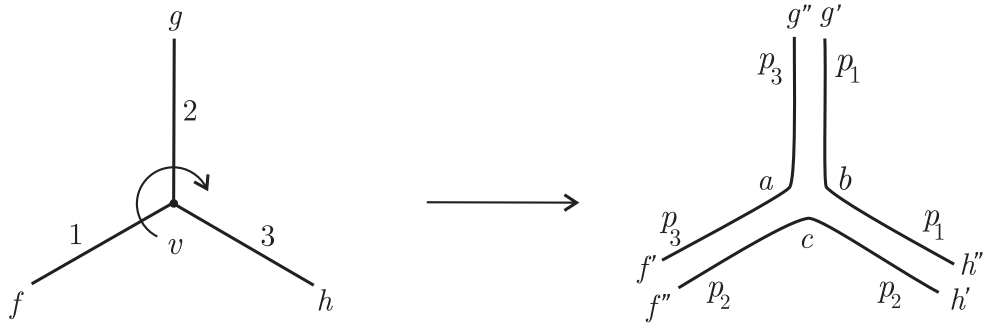

Figure 9 Colouring a vertex gadget that corresponds to a vertex of Type 1.

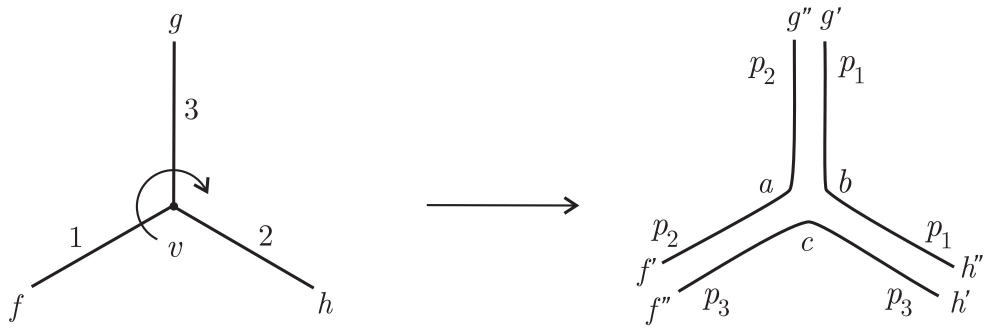

Figure 10 Colouring a vertex gadget that corresponds to a vertex of Type 2.

Next we colour the edge gadgets of $G^\sharp$. Consider an arbitrary edge $x = uv$ of G of colour, say, $\gamma(x) = 1$. We know that the half-edges in the connectors of both X_u and X_v corresponding to x receive colours from the line segment $\{p_2, p_3\}$. We now choose the tetrahedral colouring for the edge gadget Y_x which transforms the line segment $\{p_2, p_3\}$ into itself in such a way that the ordering of colours in both the input and the output of Y_x fits the ordering in the corresponding connectors of X_u and X_v. As argued in Lemma 4, such a colouring always exists. For edges of colours 2 and 3 we proceed analogously. In this way we transform a 3-edge-colouring γ of G into a tetrahedral colouring (that is, a tetrahedral flow) of $G^\sharp$. By Theorem 3, $\pi(G^\sharp) = 4$. This completes the proof of the claim as well as that of Theorem 2. ◁

◀

5 Final remark

The statement of Theorem 2 implies that the derived graph $G^\sharp$ is cyclically 4-edge-connected. If we relax cyclic 4-connectivity to 2-connectivity or 3-connectivity, the corresponding statement becomes significantly easier to prove. Indeed, take an arbitrary 2-edge-connected cubic graph G and substitute each vertex v with a copy X_v of the 3-pole Q obtained from the Petersen graph by removing a vertex. Identify the dangling edges of X_v with the edges of G incident with v, thereby producing a 2-connected cubic graph G^+; if G is 3-connected, so is G^+. Since Q is uncolourable, G^+ is a snark, though a trivial one. By employing our geometric theory it can be proved that $\pi(G^+) = 4$ if and only if G is 3-edge-colourable.

References

1 M. Abreu, T. Kaiser, D. Labbate, and G. Mazzuoccolo. Treelike snarks. *Electron. J. Combin.*, 23:#P3.54, 2016.

2 G. Brinkmann, J. Goedgebeur, J. Hägglund, and K. Markström. Generation and properties of snarks. *J. Combin. Theory Ser. B*, 103:468–488, 2013.

3 L. Esperet and G. Mazzuoccolo. On cubic bridgeless graphs whose edge-set cannot be covered by four perfect matchings. *J. Graph Theory*, 77:144–157, 2014.

4 I. Holyer. The NP-completeness of edge-coloring. *SIAM J. Comput.*, 10:718–720, 1981.

5 E. Máčajová and M. Škoviera. Cubic graphs that cannot be covered with four perfect matchings. *J. Combin. Theory Ser. B*, 150:144–176, 2021.

6 B. Mohar and C. Thomassen. *Graphs on Surfaces*. Johns Hopkins University Press, 2001.

7 J. Plesník. Connectivity of regular graphs and the existence of 1-factors. *Mat. Čas.*, 22:310–318, 1972.

8 P. D. Seymour. On multi-colourings of cubic graphs, and conjectures of Fulkerson and Tutte. *Proc. London Math. Soc.*, 38:423–460, 1979.

9 E. Steffen. 1-Factor and cycle covers of cubic graphs. *J. Graph Theory*, 78:195–206, 2015.