ELSEVIER

Discrete Mathematics 149 (1996) 141-157

DISCRETE MATHEMATICS

Planar graphs on the projective plane

Bojan Mohar[a,1,*], Neil Robertson[b,2], Richard P. Vitray[b,3]

[a] Department of Mathematics, University of Ljubljana, Jadranska 19, 61111 Ljubljana, Slovenia

[b] Department of Mathematics, The Ohio State University, 231 West Eighteenth Avenue, Columbus, OH 43210, USA

Received 15 July 1993; revised 7 September 1994

Abstract

It is shown that embeddings of planar graphs in the projective plane have very specific structure. By exhibiting this structure we indirectly characterize graphs on the projective plane whose dual graphs are planar. Whitney's Theorem about 2-switching equivalence of planar embeddings is generalized: Any two embeddings of a planar graph in the projective plane can be obtained from each other by means of simple local reembeddings, very similar to Whitney's switchings.

1. Introduction

Graphs considered in this paper are finite and undirected. Loops and parallel edges are permitted. For graph theory terminology we refer to [2]. Let Σ be a closed surface. An *embedding* ψ of G into Σ is a 1–1 continuous mapping $\psi : G \to \Sigma$, where G is assumed to have the natural topology as a 1-dimensional CW-complex. The graph can be identified with its image $\psi(G)$ on Σ, and we will usually not hesitate to speak about vertices or edges of the graph on the surface (under the given embedding). The connected components of $\Sigma \backslash G$ are called *faces* of ψ. In the sequel, G will be a graph, Σ a closed surface, and ψ an embedding of G into Σ. Also, ψ is always assumed to be given and fixed if not otherwise stated.

A *curve* on Σ is a continuous mapping $\gamma : [0, 1] \to \Sigma$. If $\gamma(0) = \gamma(1)$, then γ is said to be *closed*. It is *simple* if it is 1–1 with the only possible exception that it might be closed. We will identify any curve γ with its image $\gamma([0, 1])$ on Σ. Any walk in

* Corresponding author.

[1] The work was done in 1988 when the author was a Visiting Fulbright Scholar at The Ohio State University. Also supported in part by the Ministry of Science and Technology of Slovenia.

[2] This research was partially supported by NSF Grant DMS 8504054.

[3] Current address: Rollins College, Department of Mathematical Sciences, 1000 Holt Avenue – 2743, Winter Park, FL 32789-4499, USA.

SSDI 0012-365X(94)00349-1

a graph G embedded in Σ can be viewed as a curve on Σ. Cycles of G are special cases of simple closed curves on Σ.

Simple closed curves on a surface are also called *circuits*. Non-orientable surfaces contain *one-sided* circuits. Such a circuit has a neighborhood homeomorphic to the Möbius band, and it is said to *go across a cross-cap*. On the projective plane there is only one homotopy class of non-contractible closed curves. It contains one-sided circuits.

Let Σ be a closed surface different from the 2-sphere, and $\psi : G \to \Sigma$ an embedding of G into Σ. For a closed curve γ on Σ, let $\mathrm{cr}(\psi, \gamma)$ denote the number of times that γ intersects the graph G on Σ. More precisely,

$$cr(\psi, \gamma) = \Big|\{x \in [0,1); \gamma(x) \in \psi(G)\}\Big|.$$

The minimal value of $\mathrm{cr}(\psi, \gamma)$, taken over all non-contractible closed curves γ on Σ, is denoted by $\rho(\psi)$ and it is called the *representativity* of the embedding ψ. This quantity, originally introduced by Robertson and Seymour [5] in their work on graph minors, is also called the *face width* of the embedding [1,8]. By elementary topology, $\rho(\psi)$ is also the minimum of $\mathrm{cr}(\psi, \gamma)$ where γ is any non-contractible circuit that passes through vertices and faces only, and that uses no vertex or face more than once. The reader is referred to [6] for more details about the representativity of embeddings.

One of the first results about the representativity, due to the second and the third author [6] (see also [8]), considers the representativity of non-planar embeddings of planar graphs.

Theorem 1.1. *Let G be a planar graph, Σ a closed compact surface different from the 2-sphere, and ψ an embedding of G into Σ. Then $\rho(\psi) \leq 2$.*

This theorem has been further generalized to non-planar graphs [6], stating that any non-genus embedding of G has representativity bounded by a linear function of the genus of G.

Our main concern is to strengthen Theorem 1.1 to obtain, essentially, a simple description of the structure of planar graphs embedded in non-planar surfaces. As a first step we give such a description for the case when Σ is the projective plane. It turns out that planar graphs embed on the projective plane in a very simple way. Every such embedding can be modified into a planar embedding by using simple elementary changes. This gives rise to a generalization of Whitney's theorem [12] stating that any two embeddings of a 2-connected planar graph can be obtained from each other using a sequence of generalized switchings (Corollary 4.3). Our results also yield an answer to the question which 3-connected planar graphs admit a closed 2-cell embedding in the projective plane (Corollary 4.2). It should be mentioned that the main result of this paper, Theorem 3.2, appeared as a part of the Ph. D. thesis of the third author [11].

In our further works it will be shown that embeddings of planar graphs in arbitrary surfaces other than the 2-sphere have a special structure. Indirectly, this will give a characterization of graphs on these surfaces whose dual graphs are planar (cf. Corollary 4.4 for the case of the projective plane). It turns out that these embeddings can be described in terms of non-contractible circuits in the surface, meeting the graph in at most two points (which may be taken to be vertices of the graph). The close connection between the fundamental group of the surface and the planar graph embeddings is perhaps the most interesting aspect of this study. Several consequences follow from these results [13].

2. Basic definitions

Let γ and γ' be closed curves on some surface having only finitely many segments in common. Assume that their base points $\gamma(0)$ and $\gamma'(0)$ do not lie on γ' and on γ, respectively. The curve γ is said to *cross* γ' if there are intervals $I, I' \subset [0,1]$ such that the sets $\gamma(I)$ and $\gamma'(I')$ have a common point or a common segment, and they look on the surface as is shown in Fig. 1. If γ and γ' overlap but do not cross, they are said to *touch* each other. One can similarly define when a curve crosses, or touches itself.

Let ψ be an embedding of G into Σ. Suppose that there is a contractible curve γ on Σ which bounds an open disk D on Σ. Suppose that γ does not cross itself (but it may touch itself). Assume also that γ intersects the graph G only at vertices of G. Denote by H the part of G which lies in (the interior of) D. Consider the following possibilities:

(a) $\mathrm{cr}(\psi,\gamma) = 0$: The deletion of H is called a *0-reduction* of ψ.

(b) $\mathrm{cr}(\psi,\gamma) = 1$: A similar operation as above, the deletion of H, is a *1-reduction* of the embedding ψ. By this operation, the vertex of $\gamma \cap G$ is not removed from G.

(c) $\mathrm{cr}(\psi,\gamma) = 2$ and there is a path P in H connecting the two points x, y of the intersection of γ with G: We can replace H by an edge e connecting x and y and embed this edge onto the segment P. Such an operation is a *2-reduction* of ψ. Of course, if γ intersects G twice at the same point of G (i.e., $x = y$), the path P and e are assumed to go 'across' the disk D.

If ψ admits no non-trivial 0-, 1-, or 2-reductions, it is said to be *reduced*. Every embedded graph can be made reduced by a sequence of reductions. It can be shown that the obtained reduced embedding is (essentially) unique. Its importance lies in the fact that it preserves the essential properties of the embedding. In particular, the representativity does not change after a reduction.

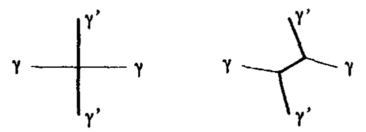

Fig. 1.

Lemma 2.1. *Let ϕ be an embedding in Σ which is obtained from the embedding ψ by a k-reduction ($k = 0$, 1, or 2). Then $\rho(\psi) = \rho(\phi)$.*

The proof is left to the reader.

An embedding ψ is *k-reduced* if no non-trivial j-reductions ($j \leq k$) of ψ are possible. The embedding ψ is *cellular* if the faces of ψ are open disks. It is a *closed-cell* embedding if, moreover, every face is bounded by a circuit in G. It is easy to see that a non-spherical embedding ψ is cellular if and only if it is 0-reduced and $\rho(\psi) \geq 1$. We will need the following simple result:

Proposition 2.2. *Let ψ be an embedding of G in some surface that is not simply connected. Then the following assertions are equivalent:*

(a) ψ is a closed-cell embedding.

(b) ψ is 1-reduced and $\rho(\psi) \geq 2$.

(c) G is 2-connected and $\rho(\psi) \geq 2$.

Proof. Equivalence of (a) and (c) is proved in [6]. We will prove that (a) and (b) are equivalent. In a closed-cell embedding, the closure of each face is a closed disk. By the remark given above, ψ is 0-reduced and has representativity at least 1. It is also straightforward that ψ must be 1-reduced. Moreover, $\rho(\psi) \geq 2$, since no face boundary touches itself. The converse is similar. □

Let γ be a closed curve on Σ without self-crossings but possibly touching itself. Assume, moreover, that γ bounds an open disk, say D, on Σ. Given $k \geq 0$, assume that $\mathrm{cr}(\psi, \gamma) \leq k$ and that γ intersects G only at vertices. Then the part of the graph lying in D together with the vertices on its boundary is called a *k-patch* (or just a *patch*). In any graphical representation, a k-patch will be represented as a shaded area with k vertices on its boundary explicitly shown. Its meaning is that any plane embedded graph attached to the vertices on the boundary can appear in the shaded area. This includes disconnected or even empty graphs and possible vanishing of the vertices on the boundary. Also, any surface contraction of the shaded area or its parts is allowed, with the only restriction that the contraction should not change the homeomorphism type of the surface. In particular, any two vertices shown on the boundary of a patch can be identical. In case when two boundary vertices of a k-patch are identified, we speak about a *degenerate* k-patch. For example, the patch in Fig. 2 represents either of the drawings in Fig. 3. Examples (c) and (d) are degenerate.

Given a non-trivial 2-patch in the disk D, one can obtain another embedding of the same graph by reversing the orientation of the patch by choosing an orientation-reversing homeomorphism of D onto itself which fixes the ≤ 2 vertices of G on the boundary of D. Such an operation is known as the (Whitney) *2-switching*.

We will use another type of local changes of embeddings. Let γ be a closed curve without self-crossings. Suppose that γ bounds a disk D and that γ meets G (at most) four times in vertices a, x, b, x, respectively. See Fig. 4. Note that the vertex x is

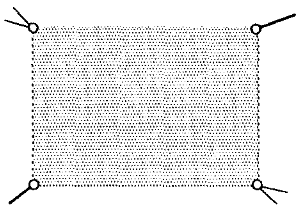

Fig. 2.

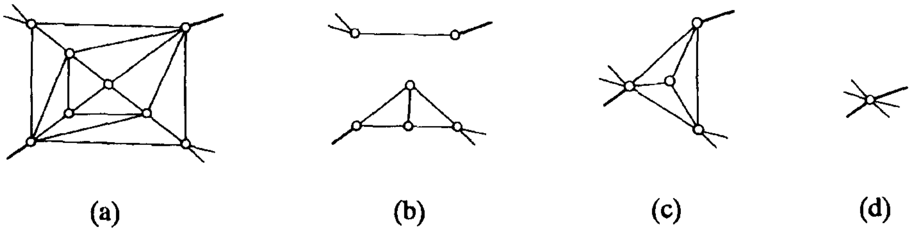

Fig. 3.

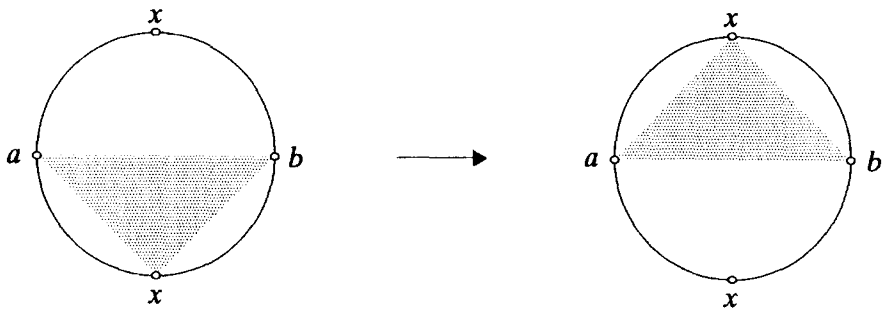

Fig. 4.

repeated. Thus, γ is self-touching. If there is a 3-patch in one of the two triangles axb in D, then the embedding ψ can be changed into another embedding ψ' of the same graph by a 'flip' as shown in Fig. 4. This operation is called a *3-switching*. As a 3-switching we also admit the degenerate case when D is 'squeezed in the middle' so that a and b become identified and D is the union of two disks, $D_1 \cup D_2$. In this case, 3-switching reembeds the 2-patch from D_1 into D_2.

Let H be a subgraph of G. An *H-component* is a subgraph of G which is either an edge not in H but with both ends in H (including loops), or a connected component of $G \backslash H$ together with all edges (and corresponding vertices in H) which have one end in this component and the other end in H. The vertices of an H-component B which lie in H are *vertices of attachment* of B. The set of vertices of attachment of B is denoted by att(B). Each edge of B incident with a vertex of attachment is a *foot* of B.

If C is a cycle in G, the C-components are also called *bridges* of C in G, cf. [2,9]. Two bridges B_1 and B_2 *overlap* if either there are four distinct vertices $a, b, c, d \in V(C)$ in the respective order such that $a, c \in$ att(B_1) and $b, d \in$ att(B_2), or B_1 and B_2 have three or more vertices of attachment in common. Tutte [10] characterized planar graphs as those for which the bridges of any cycle C can be split into two classes, so that no two bridges in the same class overlap. Any planar embedding of G determines such a partition. In one class are those bridges which are outside C, in the other class are the bridges embedded inside.

Lemma 2.3. *Let ψ be a 1-reduced embedding of a graph G in a surface Σ. Let C be a cycle of G bounding a disk D in Σ, and let $P \subset C$ be a segment of C. Then there is a cycle C' in G bounding a disk $D' \subseteq D$ in Σ, which coincides with C outside P, and no bridge of C, having all its vertices of attachment to C in P, intersects the inside of D'.*

Proof. If P is just a point, the statement is true for $C' = C$ since ψ is 1-reduced. Therefore P is a path on C joining distinct vertices a and b on C, say. Let B be a bridge of C, lying inside D and having all its vertices of attachment on P. Let x be the vertex of attachment of B which is as close as possible to a (measured on P). Similarly, let y be the vertex closest to b. Since ψ is 1-reduced, $x \neq y$. Going "clockwise" around x (i.e., starting with the edge on C joining x in the direction toward y, and then continuing with the first foot of B at x), we reach, after the last foot of B at x, a face F, say. Since ψ is 1-reduced and $F \subseteq D$, the boundary of F must be a cycle in G. It is also easily seen that it contains y, and that no other vertex of C appears on an x-y segment, say S, of this cycle. Therefore, the replacement of the x-y segment on P by S gives a new cycle $\bar{C}$ bounding a disk $\bar{D} \subseteq D$. It is easy to see that any bridge of $\bar{C}$ lying in $\bar{D}$ is also a bridge of C. If we repeat the above procedure for the remaining bridges of C, we obtain the required cycle C'. □

3. The projective plane

In this section we show that a planar graph in the projective plane must be embedded in a certain very simple way. To present embeddings in the projective plane we use its standard representation as a closed disk with any two opposite points on the boundary being identified. All our pictures will use this convention. Our notation for the projective plane is $\widetilde{\Sigma}_1$.

Lemma 3.1. *Let ψ be an embedding of G into the projective plane. If $\rho(\psi) \leq 1$, then G is a planar graph.*

Proof. Let γ be a non-contractible circuit on $\widetilde{\Sigma}_1$ with $\mathrm{cr}(\psi, \gamma) = \rho(\psi) \leq 1$. Let Φ be the quotient map which takes $\widetilde{\Sigma}_1$ to $\widetilde{\Sigma}_1/\gamma$, i.e. contracts γ to a point. By elementary topology, $\Phi(\widetilde{\Sigma}_1)$ is homeomorphic to the 2-sphere. Clearly, since $\mathrm{cr}(\psi, \gamma) \leq 1$, the composition $\Phi \circ \psi$ is an embedding of G into the 2-sphere. □

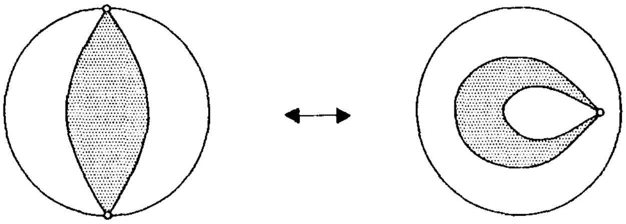

Fig. 5.

Geometrically, the case $\rho(\psi) = 0$ is obviously 'planar' — replace the non-simply connected face by a disk. If $\rho(\psi) = 1$, we just change the (local) rotation of the vertex v of G which lies on a 1-representative curve γ by switching the order at one side of γ, and then replace the non-simply connected face by a disk. Globally, the change is as shown on Fig. 5. We call it a *cross-cap switching*.

Our next theorem is the main result of this paper. It describes the structure of embeddings of planar graphs in the projective plane.

Theorem 3.2. *Let G be a graph embedded in the projective plane. Then G is planar if and only if the embedding has the structure as shown in Fig. 6, or in Fig. 7, where the shaded triangles are 3-patches.*

Remark. The structure of Fig. 7 needs more precise definition. Denote by s the leftmost (= the rightmost) point in Fig. 7 Then there are points $a_0, a_1, \ldots, a_{2k}$, $k \geq 0$, in $\widetilde{\Sigma}_1$ (corresponding to the points on the vertical rim where the 3-patches meet), and the 3-patches in Fig. 7 have points a_i, a_{i+1}, s ($i = 0, 1, \ldots, 2k$, indices modulo $2k + 1$). The patches are positioned interchangeably to the left and to the right of the vertical rim. If the representativity of the embedding is 2, then the points s and $a_0, a_1, \ldots, a_{2k}$ are vertices of the graph. Because of degeneracies of the patches, it may happen that some of the consecutive points a_i, a_{i+1} coincide. Any degeneracy of 3-patches in Fig. 6 or Fig. 7 is allowed. However, if any of the four patches in Fig. 6 is degenerate, the

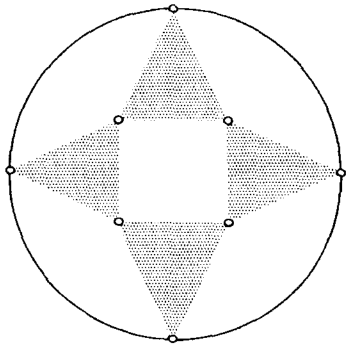

Fig. 6.

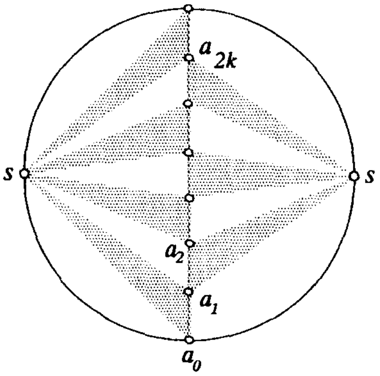

Fig. 7.

embedding also fits into Fig. 7, as it can easily be verified. The vertical rim in Fig. 7 needs not to be a cycle since degeneracies of 3-patches may result in identification of the vertex s with several vertices a_i. Also, the vertices a_i and a_{i+1} may have the only connection in the corresponding 3-patch via the vertex s. However, it turns out that the vertical rim is a cycle of the graph that is disjoint from s if the representativity is equal to 2 and ψ is 1-reduced.

Proof. First of all, we must verify that any graph with a projective embedding having the structure of Fig. 6 or Fig. 7 is planar. Any triangular patch in Fig. 7 can be reembedded (by using 3-switching operation) onto the other side of the vertical rim. After a sequence of 3-switchings, we get an embedding with representativity 0 or 1. By Lemma 3.1, G is planar. Clearly, also any graph with the structure of Fig. 6 can be reembedded in such a way as to obtain an embedding with representativity at most 1. In the plane we get an octahedron, four of whose faces are replaced by triangular patches.

Now, let G be a planar graph and ψ its embedding into the projective plane. We may assume that ψ is reduced and $\rho(\psi) \geq 2$. The case $\rho(\psi) < 2$ is done by Lemma 3.1, and it clearly fits into Fig. 7 (with just one patch).

Let C be a cycle in G which bounds a disk D in $\widetilde{\Sigma}_1$ and is maximal in the sense that no other cycle C' in G bounds a disk D' such that $D \subset D'$. Since $\rho(\psi) \geq 2$, such a cycle must exist by Proposition 2.2. We claim that no vertex of G lies out of D. Consider an arbitrary bridge B of C that is embedded outside D. If B is attached to three or more vertices on C, it contains a vertex v with three disjoint paths connecting v with C. At least one pair of these paths bounds a disk with C. Replacing the part of C lying in this disk by the two paths, gives rise to a cycle bounding a disk $D' \supset D$. This contradicts the maximality of C. By Proposition 2.2, G is 2-connected. Therefore B is attached to exactly two vertices, say u, v, of C. Suppose that B is not just an edge. Since ψ is 2-reduced, B contains an essential cycle C'. By Menger's theorem there are vertex disjoint paths P_1, P_2 joining $\{u, v\}$ with C' (possibly trivial if u or v belongs to C'). Denote by Q_1, Q_2 the segments of C' between the ends of P_1 and P_2 on C'. Then either $P_1 \cup Q_1 \cup P_2$ or $P_1 \cup Q_2 \cup P_2$ is homotopic to the segment from u to v on C. Clearly, this contradicts the maximality of D.

So, outside D there are only edges of G. Denote them by $e_1, e_2, \dots, e_r$ and call them *outside edges*. Each of these edges connects two vertices on C and is going "across the cross-cap", i.e., e_i together with a segment on C joining its endvertices is non-contractible. These edges do not contain a 3-matching (a set of three pairwise non-adjacent edges) since otherwise, C together with these edges would form a subgraph of G homeomorphic to $K_{3,3}$.

For a vertex x on C, let $d(x)$ be the number of the outside edges incident with x. Assume that there is no vertex y on C with $d(y) = 1$. If there is a vertex x with $d(x) \geq 3$, then one of its neighbors will have $d(y) = 1$. Thus, we may assume that all vertices x on C with $d(x) > 0$ have $d(x) = 2$. It is easy to see that in case of an even number of outside edges, they would appear in parallel pairs. Since a pair of

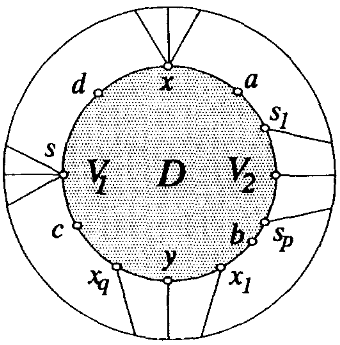

Fig. 8.

parallel edges out of D can be 2-reduced to a single edge, this would contradict the reducibility. The vertices with $d(x) = 2$ and edges e_i generate a 2-regular graph. Since its order is odd, it follows easily that this graph is connected, i.e., it is a cycle. If it contains more than five vertices, then the first, the third and the fifth edge of this cycle form a 3-matching. Consequently, there are at most five of the edges. Five of them are impossible since this gives a subdivision of K_5 in G. So there are three of them, and this is easily seen to fit Fig. 7 (where D is the only nontrivial patch — see Fig. 4.1).

Now we assume that there is a vertex y on C with $d(y) = 1$. Let x be the other end of the outside edge at y. Denote by V_1 the set of vertices on C lying on one $x - y$ segment of C (say on the 'left'; see Fig. 8), and let V_2 be the vertices on the 'right'. We do not include x or y in V_1 or V_2. The graph H generated by $V_1 \cup V_2$ and the edges $e_1, \ldots, e_r$ (not counting those which are incident with x) is clearly bipartite. H does not have a 2-matching since such a 2-matching together with the edge xy gives rise to a 3-matching in $\{e_1, \ldots, e_r\}$. Any bipartite graph H without a 2-matching has a vertex, denote it by s, which covers all the edges in H. So s and x together cover all the outside edges. Since G is planar, the bridges of the cycle C in G, including the outside edges, can be split into two parts, $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$, such that no two bridges of $\boldsymbol{B}_1$ and no two bridges of $\boldsymbol{B}_2$ overlap. Observe that, since $\rho(\psi) \geq 2$, there is at least one edge in H. Any edge of H overlaps with the edge xy. Therefore, if xy is in $\boldsymbol{B}_1$, then all edges covered by s are in $\boldsymbol{B}_2$, except that if $e_j = sx$ for some j, this edge can be either in $\boldsymbol{B}_1$ or $\boldsymbol{B}_2$.

Denote by $s_1, s_2, \ldots, s_p$ the neighbors of s in H, and suppose that they appear on C in the direction from x towards y in the given order. Let $xx_1, xx_2, \ldots, xx_q$ be the outside edges covered by x, excluding the outside edge sx (if present at all). We assume that they are enumerated according to the order of vertices $x_1, x_2, \ldots, x_q$ on C. We assume henceforth that C is oriented so that the order of vertices on C is $x \to s_1 \to s_p \to x_1 \to y \to x_q \to s \to x$. If $a, b \in V(C)$ we will denote by $[a - b]$ the closed segment of C from a to b in the given direction. Observe that $a, b \in [a - b]$ and that $[a - b] \cap [b - a] = \{a, b\}$. Denote by $(a - b) := [a - b] \backslash \{a, b\}$ the open segment on C from a to b.

Define vertices x' and x'' as follows. Let $e = xx_q$ and let $S = [s - s_1]$. If there is a C-bridge in D which overlaps with e and has all its vertices of attachment to C in the segment S, let x' be the vertex of attachment of such a bridge that is as close as possible to s (i.e., the segment $[s - x']$ is as short as possible). Let x'' be the vertex of attachment of such a bridge as close as possible to s_1. If no bridge with the required properties exists, let $x' = x'' = x$. Clearly, $x' \in [s - x]$ and $x'' \in [x - s_1]$. Define similarly the vertices s'_p, s''_p by considering the C-bridges in D overlapping with $e = s_p s$ and with all their vertices of attachment in $S = [x - x_1]$. Finally, define s', s'' in the same way but considering $e = ss_1$ and $S = [x_q - x]$. Note that the order of these vertices on C is as follows: $x', x, x'', s'_p, s_p, s''_p, x_1, x_q, s', s, s'', x'$. Define vertices a, b, c, d on C as $a := x''$, $b := s''_p$, $c := s'$, and either $d := s''$, or $d := x'$. For d we will choose one or the other alternative later. Then $a \in [x - s_1]$, $b \in [s_p - x_1]$, $c \in [x_q - s]$, and $d \in [s - x]$. See Fig. 8.

We claim that no C-bridge is attached to C strictly between s'' and x'. Suppose that B is a bridge with an attachment in $(s'' - x')$. First, B does not have all its attachments in $[s'' - x']$ since ψ is reduced. By the definition of x', B has an attachment to C which does not belong to the segment $[s - s_1]$. Therefore B overlaps with ss_1. Similarly, by the definition of s'', B has an attachment out of $[x_q - x]$. Therefore B overlaps with xx_q. Since ss_1 also overlaps with xx_q, this contradicts the planarity of the graph.

Suppose first that each bridge of C inside D has all vertices of attachment in just one of the segments, either $[a - b]$, $[b - c]$, $[c - d]$, or $[d - a]$. If sx is not among the outside edges, then ψ has the structure of Fig. 6 with x being the top (and the bottom) vertex and s being the vertex on the right and the left of the figure. The four 'inner' vertices of the 3-patches correspond to a, b, c, and d, respectively. If sx is an outside edge, the same is true when $c = s$ (sx goes in the patch determined by b, c, x), or when $a = x$ (sx goes in the patch determined by a, b, s). In the remaining case, when $c \neq s$ and $a \neq x$, we easily see that G is not planar.

It is possible that $d(s_1) > 1$. This case behaves differently than the case when $d(s_1) = 1$ since we cannot tell whether the outside edge xx_1 is in $\boldsymbol{B}_1$ or in $\boldsymbol{B}_2$. We will consider this case at the end of the proof. Up to that point we assume that $d(s_1) = 1$.

It is easy to see that every bridge of C has all its vertices of attachment either in the union of segments $[a - b] \cup [c - s'']$ on C, or in the segments $[b - c] \cup [x' - a]$. By our choices of a, b, c, if a bridge of C does not have all its attachments in a single segment among $[a - b]$, $[b - c]$, $[c - s'']$, or $[x' - a]$, the possible sets of attachments are $[a - b] \cup \{s', s''\}$, or $[b - c] \cup \{x', x''\}$.

We are left with the case when there is a C-bridge B whose vertices of attachment are not contained in a single segment among $[a - b]$, $[b - c]$, $[c - d]$, and $[d - a]$. Then the roles of x and s can be exchanged, if necessary. Because of this symmetry, we may assume that the bridge B is joining the segments $[a - b]$ and $[c - d]$. Then we choose $d = s''$. The rest of the proof splits in four cases, (a)–(d), as follows.

(a) $s' = s'' = s$: Then $c = d = s$. Since B is not attached just to one of the four segments, we either have $\mathrm{att}(B) = \{a, b, c\}$, or B has an attachment in the open

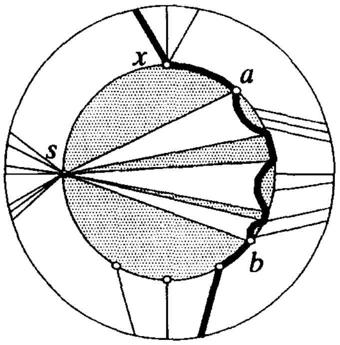

Fig. 9.

segment $(a-b)$. In the first case, B overlaps with xx_q. (If not, then $a=x$ and $b=x_q$. But then, by the definition of s' and s'', we would have $s'=x_q \neq s$.) Therefore B does not overlap with ss_i $(1 \leq i \leq p)$. Consequently, $\{s_1,\ldots,s_p\} \subseteq \{a,b\}$. If we take the segment $[x-x_1]$ together with the outside edge xx_1 as the central rim, we get the structure of Fig. 7. (The two patches on the right are just the outside edges (if present) as and bs. The edge sx, if present, lies in the lower left patch which contains all other edges covered by x.)

If B has an attachment in the open segment $(a-b)$, then it is easy to see that the structure of ψ is as shown in Fig. 9. Let $P_1,\ldots,P_j$ be the maximal subsegments of $[a-b]$ such that no interior point of P_i $(1 \leq i \leq j)$ is a vertex of attachment of a C-bridge that overlaps with the outer edge xx_q. For $i=1,\ldots,j$, let C_i' be the path in G obtained by Lemma 2.3 applied on the cycle C and its segment P_i. Denote by $P_i' = C_i' \backslash (C \backslash P)$ the segment of C_i' that replaces P_i, Let R be the cycle obtained from $xx_1 \cup [x-x_1]$ by replacing $P_1 \cup \cdots \cup P_j$ with $P_1' \cup \cdots \cup P_j'$. In Fig. 9, R is represented as the thick cycle. It is now easy to see that the embedding ψ has the structure of Fig. 7 where the cycle R represents the central rim, and where the vertex s corresponds to the leftmost (= the rightmost) vertex of Fig. 7. The outside edge sx, if present, lies in the lowest left patch.

(b) $s' \neq s \neq s''$ and $\mathrm{att}(B) \supseteq \{s',s''\}$: Then B overlaps with ss_1. Since G is planar, B cannot at the same time overlap with any of the xx_i $(1 \leq i \leq q)$. We have to consider two cases:

(i) $s'' \neq x$. Since B does not overlap with xx_q, it follows that B is attached only to the segment $[x_q-x]$. It also overlaps with ss_1. By the definition of $s'=c$, we see that B is attached to the single segment $[c-d]$. A contradiction to the choice of B.

(ii) $s''=x$. Then $x'=x''=x=a=d$ and we finish the proof as in Case (i).

(c) $s' \neq s \neq s''$, $s' \in \mathrm{att}(B)$, and $s'' \notin \mathrm{att}(B)$: Since $\mathrm{att}(B) \not\subseteq [b-c]$, B has an attachment in $[a-b] \backslash \{b\}$. Therefore B overlaps with ss_p. B cannot at the same time overlap with xx_i (whenever $x_i \neq s_p$). It follows that either $s'=x_q$, or $a=x$ is the only attachment of B in $[a-b]$. In the latter case, B would force that $s''=x$, and B would be confined to the single segment $[c-d]$. So we get $s'=x_q$. Moreover, we either have $q=1$, or we have $q=2$ and $x_1=s_p=b$. In the first case, B forces that $s_p''=x_q$, and B

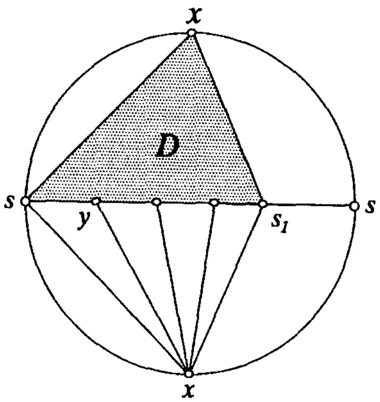

Fig. 10.

would be attached only to the segment $[a-b]$. Therefore we have the other possibility, $q = 2$. Clearly, $xx_q \in \boldsymbol{B}_1$. Then $ss_1 \in \boldsymbol{B}_2$ and $B \in \boldsymbol{B}_1$. Since xx_1 overlaps with ss_1, it cannot overlap with B. But this is possible only when $\mathrm{att}(B) = \{x_q, x\}$ and $x = a$. By the definition of s'' we thus get $s'' = x$. So, B is confined to $[c-d]$. A contradiction.

(d) $s' \neq s \neq s''$, $s' \notin \mathrm{att}(B)$, and $s'' \in \mathrm{att}(B)$: Similarly as in (c) we see that B has an attachment in $[a-b]\backslash\{a\}$. The case $s'' = x$ is easily ruled out. Therefore B overlaps with xx_q. But then it cannot overlap with ss_1. Consequently, $\mathrm{att}(B) \subseteq [s'' - s_1]$. By the definition of a, we thus have $\mathrm{att}(B) \subseteq [s'' - a] = [d - a]$. A contradiction.

It remains to consider the case when $d(s_1) > 1$. It is clear that this implies $p = 1$ and that s_1x, s_1s are the outside edges at s_1. If $d(s) = 1$, then we could have taken s_1 instead of s, and we would be in one of the above cases. Thus, we may assume that $d(s) = 2$. Then the outside edges at s are ss_1 and sx. See Fig. 10. If s and s_1 are on the boundary of a common face $F \subseteq D$, then we have the structure of Fig. 7 with s being the left (= the right) vertex, and $[x - s_1]s_1x$ being the central rim. (The only patch on the right of the rim is the edge s_1s.) Otherwise, there is a C-bridge B' in D overlapping with ss_1. If $x' = x = x''$, then it is easy to see that we get the form of Fig. 7 as in the above Case (a) (with the roles of x and s exchanged). If not, then let B'' be a bridge defining x' and x''. Then $xy \in \boldsymbol{B}_1$, $B'' \in \boldsymbol{B}_2$, $ss_1 \in \boldsymbol{B}_2$, and $B' \in \boldsymbol{B}_1$. Also, either sx, or s_1x overlaps with B'. Assume that sx does. Then $sx \in \boldsymbol{B}_2$. In this case, B' has an attachment on the segment $[s - x']\backslash\{s\}$. Therefore $x' \neq s$ and, consequently, B'' overlaps with sx. This implies that $sx \in \boldsymbol{B}_1$, a contradiction. Similar arguments work if s_1x overlaps with B'.

We exhibited all cases. The proof is complete. □

4. Some applications

As the first corollary to Theorem 3.2 we will prove that embeddings of 4-connected planar graphs into the projective plane are very restrictive. Since the embeddings with representativity ≤ 1 are very close to planar embeddings (cf. Lemma 3.1), we will only describe the case of representativity 2.

Corollary 4.1. *Let G be a 4-connected planar graph embedded in the projective plane. If the representativity of the embedding is equal to 2, then G is either the graph of the octahedron embedded as shown in Fig. 6, or G has a triangle xyz, and the embedding into the projective plane is as shown on Fig. 11, where the patch gives a planar embedding of G except that the edges between the vertices x, y and z use the way out of the patch.*

Proof. Suppose first that the embedding of G has the structure of Fig. 7 If any of the patches contains a vertex of G in its interior, then by 4-connectedness of G all vertices lie in this patch, which is just Fig. 11. Note that x, y and z must be distinct and the edges between them must be out of the patch since the representativity is 2. Otherwise, each patch contains only its boundary vertices (at most three) and some edges between them. All the vertices except the vertex s on the 'left' lie on the vertical rim. If there are more than three vertices on the rim, then s together with the first and the third vertex on the rim separates the second one from the rest which contradicts the 4-connectivity. The small cases with at most three vertices on the rim are also impossible because a 4-connected graph has at least 5 vertices.

The remaining case is when the embedding has the structure of Fig. 6. Any degeneracy of patches leads to the structure of Fig. 7 that we have already covered. Therefore, all six vertices from Fig. 7 are distinct. Since G is 4-connected, no 3-patch contains a vertex apart from the given three. Since G has no vertices of degree three or less, all three vertices in each patch must be pairwise adjacent. Therefore G is the graph of the octahedron. $\square$

Every plane graph has either a vertex of degree at most 3 or a face of size at most 3. Therefore, every 4-connected planar graph contains a triangle xyz (which is necessarily facial). Therefore, such a graph admits a closed-cell embedding in the projective plane as shown in Fig. 11. Richter et al. proved in [4] that every 3-connected

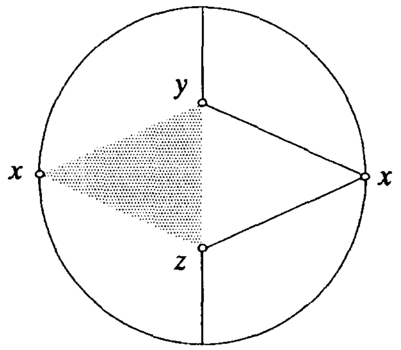

Fig. 11.

planar graph admits a closed-cell embedding in a non-orientable surface of genus at most 3. By Theorem 3.2, we can characterize those planar graphs which have a closed-cell embedding in the projective plane.

Vertex set $\{x,y,z\} \subseteq V(G)$ is a *3-separation* if $G - \{x,y,z\}$ is disconnected.

Corollary 4.2. *A 3-connected planar graph G has a closed-cell embedding in the projective plane if and only if it satisfies one of the following conditions:*
(a) *G contains a triangle,*
(b) *G contains a 3-separation $\{x,y,z\}$ where x and y are adjacent,*
(c) *G contains distinct, pairwise non-adjacent vertices x,y,z,w such that $\{x,y,z\}$, $\{x,w,z\}$, and $\{y,w,z\}$ are 3-separations in G,*
(d) *G contains distinct, pairwise non-adjacent vertices p,q,r,s,t,u such that $\{p,q,t\}$, $\{q,r,u\}$, $\{r,s,t\}$, and $\{s,p,u\}$ are 3-separations in G.*

Proof. If G satisfies any of the above properties, one easily finds a closed-cell embedding of G in the projective plane. In cases (a)–(c) we get the structure of Fig. 7 (usually with three patches). In case (d), we get Fig. 6.

Suppose now that G has a closed-cell embedding in the projective plane and that it satisfies neither (a) nor (b). If the embedding of G has the structure of Fig. 6 (without degeneracies), then we have (d). Otherwise, we have the structure of Fig. 7. By a sequence of 3-switchings we get a closed-cell embedding with exactly three patches. Excluding (a) and (b), every patch is non-degenerate and vertices on its boundary are pairwise non-adjacent. Thus we get (c). □

Note that in Case (c) of Corollary 4.2, the plane embedding of G consists of three 3-patches attaching one another at vertices x,y,z,w. Similarly, the condition of Case (d) is equivalent to have four 3-patches forming a structure like an octahedron.

Another consequence of Theorem 3.2 is a generalization of Whitney's 2-switching Theorem. Embeddings ψ and φ of the same graph are *switching equivalent* if φ can be obtained from ψ by a sequence of switchings. On the projective plane, the following operations represent the admissible *switchings*:
- the Whitney 2-switching,
- the 3-switching (Fig. 4),
- the cross-cap switching (Fig. 5) and its inverse, and
- the operation shown in Fig. 12.

Corollary 4.3. *Any two embeddings of a 2-connected planar graph in the projective plane are switching equivalent.*

Proof. By Theorem 3.2, it is evident that the switching operations described above can be used to transform any embedding into a planar one (representativity 0). We are done by the reversibility of all switching operations and by the Whitney's 2-switching theorem [12], which implies that any two planar embeddings of a 2-connected graph

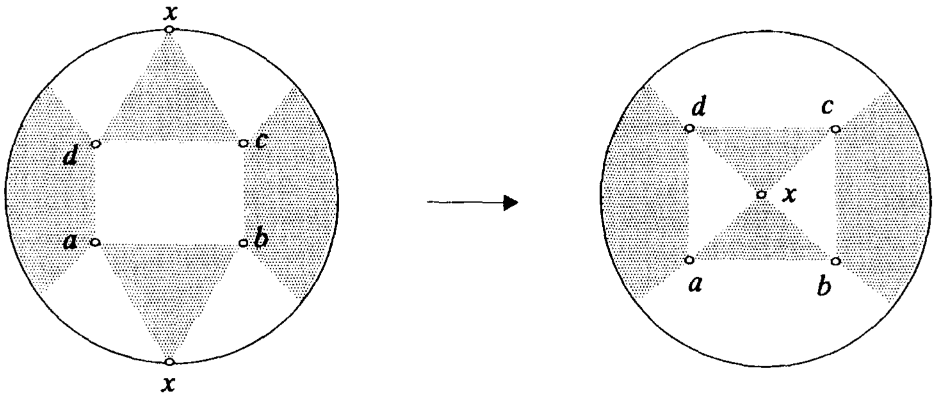

Fig. 12.

can be obtained from each other by a sequence of 2-switchings. Notice that to obtain an embedding with representativity 0, the cross-cap switching and the switching of Fig. 12 must each be used at most once. □

Corollary 4.3 nicely complements results of Negami [3] who considered embedding flexibility of non-planar graphs in the projective plane.

At the end we present some results about graphs with planar duals. Let ψ be an embedding of a graph G into a surface Σ. The *geometric dual* of G with respect to ψ is a graph G^* together with an embedding $\psi^* : G^* \to \Sigma$ which is obtained as follows. Vertices of G^* correspond to the components of face boundaries of ψ. (If $\rho(\psi) \geq 1$ and G is connected, then each face has only one boundary component.) The edge set of G^* is just $E(G^*) = E(G)$, and two vertices A, B are joined in G^* by an edge e if A and B are the face boundaries of ψ containing the edge e. We get a dual embedding ψ^* by taking a regular neighborhood of each of the face boundary components, and embedding the corresponding vertex and the edges in it, so that $\psi^*(e)$ traverses $\psi(e)$ for each edge e in $E(G)$. It should be mentioned that $\rho(\psi) = \rho(\psi^*)$, including the possibility when the representativity is 0.

Theorem 3.2 indirectly characterizes graphs in the projective plane having the dual graph which is planar. Let us describe this characterization more closely.

Corollary 4.4. *A graph G can be embedded in the projective plane in such a way that its geometric dual is planar if and only if one of the following cases holds:*

(a) *G is planar.*

(b) *There are plane graphs G_1, G_2, G_3, G_4, each G_i ($i = 1, \ldots, 4$) has a face (in its plane embedding) with distinct vertices a_i, b_i, c_i on its boundary, and G is obtained from the disjoint union of G_i by identifying vertices $a_1, \ldots, a_4$ into a single vertex, and similarly identifying $b_1, \ldots, b_4$ and $c_1, \ldots, c_4$, respectively, into two other vertices.*

(c) *There are plane graphs $G_1, G_2, \ldots, G_{2n+1}$ ($n \geq 1$), each G_i ($i = 1, \ldots, 2n+1$) has a face (in its plane embedding) with distinct vertices a_i, b_i, c_i on its boundary, and G is obtained from the disjoint union of G_i by identifying, for each i, the vertex b_i with a_{i-1} and c_{i+1} (indices modulo n).*

Proof. Let G be a graph with an embedding $\psi : G \to \widetilde{\Sigma}_1$ such that the dual graph G^* is planar. If $\rho(\psi^*) \leq 1$, then also $\rho(\psi) \leq 1$. By Lemma 3.1, G is a planar graph.

Assume now that $\rho(\psi) = \rho(\psi^*) \geq 2$. By Theorem 3.2, ψ^* has the structure of Fig. 6, or Fig. 7. If ψ^* does not have the structure of Fig. 7, then it must have the structure of Fig. 6 without any degeneracy. The three faces between the patches must also be distinct since $\rho(\psi^*) \geq 2$. Let G_i, $i = 1, \ldots, 4$, be the plane graphs obtained from the four patches as their duals, together with the vertices of G corresponding to the faces of ψ^* surrounding each of the patches. Denoting these three vertices by a_i, b_i, c_i, respectively, we see that the graphs G_i determine G in the way as described in (b).

The last case to consider is when the representativity is (at least) two and ψ^* is as shown in Fig. 7. The duals of the patches together with the vertices corresponding to the surrounding ψ-faces as the graphs G_i determine G as explained in (c). Since $\rho(\psi^*) = \rho(\psi) = 2$, the three ψ^*-faces surrounding a patch are distinct, and so the corresponding vertices a_i, b_i, c_i of G_i are distinct. Note that there are at least three patches of ψ^*, and that their total number is odd.

The converse is obvious in Case (a). It is also straightforward in the other two cases where one easily finds an embedding of the graph in the projective plane that is dual to Figs. 6 and 7, respectively. $\square$

At the end, let us mention why maps with planar duals are important. If E is a set, let $\boldsymbol{P}(E)$ denote the linear space over GF(2) generated by E. If $E = E(G)$, the cycle space $\boldsymbol{C}(G)$ and the cocycle space $\boldsymbol{C}^*(G)$ are subspaces of $\boldsymbol{P}(E)$. An arbitrary subspace $\mathbf{C}$ of $\boldsymbol{P}(E)$ is *graphic* (resp. *cographic*) if there is a graph with the edge set E such that $\boldsymbol{C} = \boldsymbol{C}(G)$ (resp. $\boldsymbol{C} = \boldsymbol{C}^*(G)$).

Let G be a connected graph with edge set E, which is embedded in the projective plane. The subspace $\boldsymbol{B}(G)$ of $\boldsymbol{C}(G)$ generated by all face boundaries is of codimension 1 in $\boldsymbol{C}(G)$, providing the representativity of the embedding is not 0 (which we assume henceforth). In $\boldsymbol{B}(G)$, there are exactly all contractible cycles of our embedded graph. Let G^* be the geometric dual of G. Any cycle $C \in \boldsymbol{B}(G)$ separates the projective plane into two (not necessarily connected) parts. If C is the sum of the boundaries of faces $F_1, \ldots, F_k$, then one of these parts is just the union of faces F_i, $i = 1, \ldots, k$. In other words, C as a subset of $E(G^*)$ is a cutset (cocycle) separating vertices $F_1, \ldots, F_k$ of G^* from the rest. Conversely, each cutset in G^* determines in the same way a cycle in $\boldsymbol{B}(G)$. Therefore $\boldsymbol{B}(G)$ is equal to the cocycle space $\boldsymbol{C}^*(G^*)$ of G^*. Now, if G^* is planar, then we obtain a subspace of codimension 1 in $\boldsymbol{C}(G)$ which is at the same time graphic and cographic.

W.T. Tutte raised the question when a subspace $\boldsymbol{C}$ of $\boldsymbol{P}(E)$ contains a cographic subspace of codimension 1. Shih [7] obtained a partial result in this direction by solving the Tutte's problem when C is cographic, too. It is remarkable that his solution shows exactly the same structure as our Corollary 4.4. In Shih's result, the structure of connected graphs, for which the cocycle space has a cographic subspace of codimension 1, is given by Cases (b) and (c) of our Corollary 4.4 without the assumption on the planarity of the building blocks G_i. There is another rule which should be added, and this rule is just the dual rule of the cross-cap switching: an identification of two distinct vertices into a new vertex. This one corresponds to Case (a) of Corollary 4.4.

References

[1] D. Archdeacon, Densely embedded graphs, J. Combin. Theory Ser. B 54 (1992) 13–36.
[2] J.A. Bondy and U.S.R. Murty, Graph Theory with Applications, (MacMillan New York, 1976).
[3] S. Negami, Re-embedding of projective-planar graphs, J. Combin. Theory Ser. B 44 (1988) 276–299.
[4] R.B. Richter, P.D. Seymour and J. Širáň, Circular embeddings of planar graphs in non-spherical surfaces, Discrete Math., in press.
[5] N. Robertson and P.D. Seymour, Graph minors VII. Disjoint paths on a surface, J. Combin. Theory. Ser. B 45 (1988) 212–254.
[6] N. Robertson and R.P. Vitray, Representativity of surface embeddings, in: B. Korte, L. Lovász, H.J. Prömel and A. Schrijver, eds. Paths, Flows and VLSI-Layout, (Springer, Berlin, 1990) 293–328.
[7] C.-H. Shih, On graphic subspaces of graphic spaces, Ph.D. Thesis, The Ohio State University, Columbus, OH, 1982.
[8] C. Thomassen, Embeddings of graphs with no short non-contractible cycles, J. Combin. Theory. Ser. B 48 (1990) 155–177.
[9] W.T. Tutte, Connectivity in Graphs, (Univ. of Toronto Press, Toronto, 1966).
[10] W.T. Tutte, Matroids and graphs, Trans. Amer. Math. Soc. 88 (1958) 144–174.
[11] R.P. Vitray, Representativity and flexibility of drawings of graphs on the projective plane, Ph.D. Thesis, The Ohio State University, Columbus, OH, 1987.
[12] H. Whitney, 2-isomorphic graphs, Amer. J. Math. 55 (1933) 245 254.
[13] B. Mohar and N. Robertson, Planar graphs on nonplanar surfaces, submitted.