GRAPH CURVES, COLORINGS, AND MATROIDS

CIRO CILIBERTO AND RICK MIRANDA

Abstract. Coloring properties of planar maps determined by planar graphs are examined in terms of algebro-geometric properties of a degenerate algebraic curve associated to the graph. In particular, q-colorings of the planar map are shown to be equivalent to the existence of certain 1-forms on the curve. A relationship between representations of certain matroids defined on the edges of the graph and the Picard group of the graph curve is developed.

Contents

1. Introduction

In this article we would like to present some relationships between the theory of graphs and the theory of algebraic curves. To every simple trivalent graph G one can associate a (degenerate) algebraic curve C_G, which is a union of lines (one for each vertex of G) joined as the edges of G indicate.

These so-called "graph curves" have been useful in proving various theorems about Riemann surfaces, since they are natural degenerations of smooth Riemann surfaces, and any semi-continuous property of curves which is maximal for any graph curve must be maximal for the general Riemann surface also. This was the basis of the proof in [CHM] that the general curve of genus $g \geq 10$, $g \neq 11$ had a surjective Gaussian map: for these genera, there exists a graph curve for which the map is surjective. See also [M2] for a discussion of graph curves and Gaussian maps in general, [CF1], [CF2] and [M1] for computations involving gaussian maps and graph curves, and [CV] for applications to Prym-canonical curves. The reader may also wish to consult [CM1], [CM2], and [DM] for a discussion of Gaussian maps

1991 *Mathematics Subject Classification.* 05C15,14H99,14G15.
Key words and phrases. graphs, colorings, algebraic curves, forms,matroids.
Research supported in part by the NSF under grant DMS-9104058.
December 1, 2005. Set in type by $\mathcal{A}\mathcal{M}\mathcal{S}$-LaTeXVersion 1.0.

for various families of curves. A preliminary version of parts of this article appeared in [M3]. The basic reference for Gaussian maps is Wahl's paper [Wa].

This article focuses on the colorability problems for planar graphs, and we show that colorings are related to the existence of certain regular 1-forms on the graph curve. We obtain no new results, but we do put into one package a set of coloring criteria which a priori are quite different. The results always reduce the coloring of the planar map to a certain vertex-edge assignment problem in a finite field. We also develop a relationship between certain matroid representations and the Picard Group of the graph curve. In particular the component of the Picard Group containing the canonical bundle is naturally associated to certain representations of matroids defined on the edges of the graph, for which the stars of the vertices are dependent sets.

The authors would like to thank David Eisenbud and Joe Harris for several useful conversations.

2. Graphs and Graph Curves

Let G be a connected graph. Formally, this would be a pair (V, E), where V is a set (of vertices) and E is the set of edges. Each edge e can be viewed as a map from an ordered two-point set to V, representing the vertices at the ends of e. The two vertices at the ends of e will be denoted by e_1 and e_2. The edge e is a *loop* if both ends are at the same vertex. We allow multiple edges, so that there may be many edges joining two vertices. The graph G will be called *simple* if it has no loops or multiple edges; this is an important special case.

The *valency* of a vertex $v \in V$ is the number of edges containing v. A loop at a vertex v would count for 2 in the valency of v. A vertex is *trivalent* if there are exactly 3 edges containing it, and it is *sub-trivalent* if there are at most 3 edges containing it. A graph G is *trivalent*, respectively *sub-trivalent*, if every one of its vertices are.

Denote by $\mathrm{val}(v)$ the valency of a vertex v. Thus a connected subtrivalent graph satisfies $1 \leq \mathrm{val}(v) \leq 3$ for all vertices v.

The *arithmetic genus* (or simply *genus*) of a connected sub-trivalent graph is the integer

$$g = |E| - |V| + 1.$$

Note that this is always a non-negative integer. Also note that this definition differs from the graph-theoretic use of the word genus. If G is actually trivalent, then we have

$$3|V| = 2|E|$$

and so $|V| = 2g-2$ and $|E| = 3g-3$. Note in this case that there are exactly $6g-6$ incident vertes-edge pairs (v, e).

Fix a field $\mathbb{K}$. To every connected sub-trivalent graph G we can associate a connected degenerate algebraic curve C_G defined over $\mathbb{K}$ as follows. To every vertex v of G associate an abstract projective line $C_v \cong \mathbb{P}^{\Bbbk}_{\mathbb{K}}$. Mark $\mathrm{val}(v)$ points on each C_v, corresponding to the $\mathrm{val}(v)$ edges of G which meet at the vertex v. (Since the automorphism group of $\mathbb{P}^{\Bbbk}$ acts triply transitively on the points, there is a unique way to do this, up to isomorphism of the line C_v. Indeed, we may choose a coordinate on C_v so that these $\mathrm{val}(v)$ points are, in some order, $\{0\}$ or $\{0, 1\}$ or $\{0, 1, \infty\}$.) Let e be an edge of G joining vertices v and w. Therefore on the line C_v there is a marked point corresponding to e, as there is on the line C_w. (If e is a loop, then there are two marked points on C_v corresponding to e.) For each such edge, glue the line C_v to the line C_w transversally at these marked points associated to the edge e. This union of lines, glued via the edge data, is the curve C_G.

The curve C_G is an algebraic curve, reducible of course, and is called the *graph curve* associated to the connected sub-trivalent graph G. It is connected, since G is. It is an easy exercise to check that the graph curve C_G associated to a connected sub-trivalent graph G with genus g has arithmetic genus g; this is the source of the notation. Moreover if G is trivalent, it is a stable curve in the sense of Deligne and Mumford.

Note also that the construction of C_G does not depend on the orientation of the edges of G. The graph curve is actually associated to the unoriented graph.

The "smallest" simple trivalent graph curve constructible in this way is the only graph curve of arithmetic genus 3, consisting of 4 vertices and all six possible edges joining any 2 of the 4. This graph is the 1-skeleton of the tetrahedron. Its associated graph curve can be realized in the plane as 4 general lines (each meeting the other three).

Note that if a vertex v has a loop at it, then in the construction of the graph curve C_G, one would glue C_v to itself at the two points $x_1(e)$ and $x_2(e)$ corresponding to the loop, creating a rational nodal curve $\overline{C}_v$. This would then be glued to another line C_w corresponding to the third edge (unless it is the only vertex.)

Secondly, suppose that G has multiple edges joining vertices v and w. Then, in the construction of the graph curve C_G, the lines C_v and C_w would be glued together at more than one point.

In any case we see that each component of a graph curve is either a line or a rational nodal curve. If a component is a line, it meets at most three others; if a component is a rational nodal curve, it meets at most one other.

The smallest trivalent graph curves have arithmetic genus 2, with two vertices v and w. The two possibilities are that there are three edges each joining v to w, or that there are loops at both v and w with one edge joining them. The first has associated graph curve equal to two lines meeting transversally at 3 points; the second is two rational nodal curves meeting transversally at one point.

We will use the following notation in general.

G	the graph
C_G	the graph curve associated to G
$\overline{C}_v$	the component of C_G corresponding to the vertex v
C_v	the normalization of $\overline{C}_v$, isomorphic to $\mathbb{P}^1$
n_e	the node of C_G corresponding to the edge e
e_1, e_2	the first and second vertex lying on e
$x_1(e), x_2(e)$	the points of C_{e_1} and C_{e_2} corresponding to the edge e
$x_{v,e}$	the point of C_v corresponding to the non-loop edge e

3. Line Bundles on Graph Curves

Let L be a line bundle on a graph curve C_G. By restriction and pull-back, we obtain line bundles L_v on the normalization C_v of each component $\overline{C}_v$ of C_G. (I.e., we restrict, and in case there are loops in G, so that the "component" $\overline{C}_v$ is a nodal rational curve, we also pull back L to the normalization C_v.) The bundle L_v is a line bundle on a $\mathbb{P}^1$, and therefore has a degree, which we will denote by $d_v(L)$. The *total degree* of the line bundle L on C_G is the $|V|$-tuple of integers $(d_v(L))_{v \in V}$, or, equivalently, the function $d(L) : V \to \mathbb{Z}$, sending v to $d_v(L)$.

We therefore see that the Picard group of C_G decomposes into components, each component corresponding to a possible total degree. The connected component of the identity, $\mathrm{Pic}^0(C_G)$, forms a connected algebraic group over $\mathbb{K}$, which we describe next.

The total degree does not in general determine the line bundle L. The individual bundles L_v on the components C_v are of course determined by their degrees d_v; however they must be glued together, and this glueing data must be specified.

To specify this glueing data for elements of $\mathrm{Pic}^0(C_G)$, note that any line bundle $L \in \mathrm{Pic}^0(C_G)$ has degree 0 on each component. Therefore we may choose an isomorphism

$$\phi_v : L_v \to C_v \times \mathbb{K}$$

for each vertex v. This in particular induces an isomorphism of each fiber of L_v with $\mathbb{K}$, which we will also denote by ϕ_v.

Now recall that each edge is oriented. This is one place in the theory where an orientation helps! (Its affect will disappear in the end.)

At each node n_e of C_G, corresponding to an edge e of G, the 1-dimensional vector spaces $L_{e_1}|_{x_1(e)}$ and $L_{e_2}|_{x_2(e)}$ are glued together, via a linear isomorphism, to form L. Let us denote this fiber $L_{e_i}|_{x_i(e)}$ of the bundle L_v by simply $L_i(e)$. The glueing data for these two 1-dimensional spaces is then determined by a nonzero number λ_e in the field $\mathbb{K}$; the glueing isomorphism is the map

$$t_e = {\phi_{e_2}}^{-1} \circ \lambda_e \circ \phi_{e_1} : L_1(e) \to L_2(e)$$

where λ_e in the above expression of course means multiplication by the constant λ_e in $\mathbb{K}$.

Therefore we see that glueing data is determined by specifying, for each edge e, a non-zero element λ_e in the field $\mathbb{K}$. This is essentially a 1-cochain for the graph G, with values in the multiplicative group $\mathbb{K}^\times$.

The only choices made in the above assignments was the choice of isomorphism ϕ_v. Suppose the isomorphism ϕ_v changes; it can only change by multiplying by a non-zero scalar μ_v. That is, the isomorphism ϕ_v may be replaced by an isomorphism $\phi'_v = \mu_v \phi_v$. In this case, the glueing isomorphism becomes

$$\begin{aligned} t_e &= {\phi_{e_2}}^{-1} \circ \lambda_e \circ \phi_{e_1} \\ &= \mu_{e_2} {\phi'_{e_2}}^{-1} \circ \lambda_e \circ {\mu_{e_1}}^{-1} \phi'_{e_1} \\ &= {\phi'_{e_2}}^{-1} \circ \lambda'_e \circ \phi'_{e_1} \end{aligned}$$

where $\lambda'_e = \mu_{e_2} {\mu_{e_1}}^{-1} \lambda_e$. Thus under this change, the set of glueing constants $\{\lambda_e\}$ is replaced by the set $\{\lambda'_e\}$. We see that what we have is exactly that the 1-cochain $\{\lambda_e\}$ changes by a coboundary, namely the coboundary of the 0-chain $\{\mu_v\}$. Thus we have the following.

Proposition 3.1. *Let G be a connected sub-trivalent graph with associated graph curve C_G. Then*

$$\mathrm{Pic}^0(C_G) \cong H^1(G, \mathbb{K}^\times).$$

It follows from the general theory that, given another possible total degree $\underline{d}$, the line bundles on C_G with total degree $\underline{d}$ form a subset $\mathrm{Pic}^{\underline{d}}(C_G)$ which is a torsor for $\mathrm{Pic}^0(C_G)$. The torsor structure is easy to see. Suppose that L is any line bundle in $\mathrm{Pic}^{\underline{d}}(C_G)$. Then we can vary L, using a cocycle $\{\lambda_e\}$ in $H^1(G, \mathbb{K}^\times)$, by changing the isomorphisms describing the glueing data for L by the λ's.

There is one line bundle, the canonical bundle, which of course has more importance than others. Note that since C_G is a singular curve, we mean the dualizing sheaf ω_{C_G} when we speak of the canonical bundle.

Lemma 3.2. *Let G be a connected sub-trivalent graph. Then the total degree of the dualizing sheaf ω of the associated graph curve C_G is $d_v(\omega) = \mathrm{val}(v) - 2$ for all vertices v. In particular, if G is trivalent, $d_v(\omega) = 1$ for all v.*

Proof. Sections of the dualizing sheaf on C_G correspond to rational 1-forms on each C_v, with only simple poles allowed at the $\mathrm{val}(v)$ identified points $x_i(e)$; furthermore, the residues at these points must cancel in the two components identified there.

Hence we see that the dualizing sheaf ω, when restricted to each C_v, is the sheaf $\omega_{C_v}(D)$, where D is the divisor of the $\mathrm{val}(v)$ points identified on C_v. Since C_v has genus 0, the degree of ω_{C_v} is -2, so that the degree of the restriction of the dualizing sheaf is $\mathrm{val}(v) - 2$. □

4. 1-forms, Twisted 1-forms, and Residues on Graph Curves

The sections of the dualizing sheaf ω_{C_G} on a graph curve C_G will be called *regular 1-forms* on C_G, by abuse of language. As noted above, a regular 1-form on C_G is a rational 1-form on each component C_v, with at worst simple poles at the nodes; at the nodes, the two residues must cancel. This residue cancellation is the compatibility condition for patching together the "local" 1-form data. For more details, the reader may consult [BE], [CHM], or [M1]. We will denote the space of global sections of ω_{C_G} by Ω_G.

We will concentrate on the case when G is simple and trivalent; the reader can make the necessary modifications in the case where G is sub-trivalent, with possible loops or multiple edges. Recall that in this case C_G is a stable curve of genus g.

A convenient way to represent a 1-form on C_G is by using residues. Let $\omega \in \Omega_G$ be a regular 1-form on C_G. For any component C_v of C_G (corresponding to a vertex v of G), denote by ω_v the restriction of ω to the component C_v. As noted above, ω_v is a rational 1-form on C_v with at worst simple poles at the three points of C_v identified together in C_G. Let e be an edge of G emanating from v, so that n_e is one of the nodes of C_G lying on the component $\overline{C}_v$, and $x_{v,e}$ is one of the three points of C_v identified in C_G. Denote by $\mathrm{Res}_\omega(v, e)$ the residue of ω_v at the point $x_{v,e}$. Thus to every regular 1-form ω we have attached $6g - 6$ residue numbers, namely the elements $\mathrm{Res}_\omega(v, e)$ of the field k as (v, e) ranges over the $6g-6$ incident vertex-edge pairs.

Since the sum of the residues of any rational 1-form is zero, we have the equation

$$\mathrm{Res}_\omega(v, e_1) + \mathrm{Res}_\omega(v, e_2) + \mathrm{Res}_\omega(v, e_3) = 0 \tag{4.1}$$

if $\{e_1, e_2, e_3\}$ are the 3 edges of G emanating from the vertex v.

The compatibility condition passing from one component to another is expressed by the formula

$$\mathrm{Res}_\omega(v_1, e) + \mathrm{Res}_\omega(v_2, e) = 0 \tag{4.2}$$

if $\{v_1, v_2\}$ are the 2 vertices lying on the edge e.

This gives a total of $5g-5$ conditions on the possible residue numbers, which are not quite independent: the sum of all the residue numbers must be zero, and this is seen using either of the above sets of conditions. In fact this is the only dependence of these conditions, and so there are exactly $5g - 6$ independent linear conditions on the $6g - 6$ residue numbers. Thus the space of residue numbers satisfying the above conditions has dimension g. This is also the dimension of the space Ω_G of regular 1-forms on any stable curve of genus g, so the map from the vector space Ω_G of regular 1-forms to the space of residue numbers satisfying the above conditions is an isomorphism, since it is clearly injective. (A 1-form with zero residues must be zero.) Thus we have:

Proposition 4.3. *Let G be a connected simple trivalent graph. Then the space of global sections Ω_G of the dualizing sheaf ω_{C_G} for the associated graph curve C_G has dimension g. Moreover it is isomorphic, via the above assignment of residue numbers, to the subspace of $\mathbb{K}^{6g-6}$ (indexed by incident vertex-edge pairs (v, e)) consisting of those vectors $(R(v, e))$ satisfying the residue conditions (4.1) and (4.2).*

Thus we can "encode" the 1-forms in Ω_G by their residue numbers.

We want to include in this section a discussion of line bundles L on C_G whose total degree is the same as that of the canonical bundle. By Lemma 3.2, if G is trivalent, these line bundles are exactly those which restrict to a bundle of degree one on each component. For this reason we denote the set of isomorphism classes of such bundles by $\mathrm{Pic}^1(C_G)$. This is the connected component of the Picard group containing the canonical bundle. We will call a line bundle $L \in \mathrm{Pic}^1(C_G)$ a *twisted canonical bundle*.

Suppose that $L \in \mathrm{Pic}^1(C_G)$. Therefore L differs from the canonical bundle by a bundle M in $\mathrm{Pic}^0(C_G)$. Assume that the edges of G are oriented, and that with this orientation the bundle M is given by a 1-cochain $\{\lambda_e\}$, as described in Section 3. The difference between the values λ_e and 1 is the "twisting" of the bundle from the canonical bundle. The space of global sections of a twisted canonical bundle can be interpreted as a space of "twisted" homology on the graph itself; see [BE] and [CV].

A section of a twisted canonical bundle will be called a *twisted* 1*-form* on C_G. A twisted 1-form is, locally, a 1-form; the patching data moving from component to component across the nodes is varied, however. In particular, every twisted 1-form is represented by a 1-form on each component, with only simple poles allowed at the nodes. Hence to each twisted 1-form there are also residue numbers, as is the case for regular 1-forms. I.e., if σ is a twisted 1-form, the residues $\mathrm{Res}_\sigma(v, e)$ are defined for all incident vertex-edge pairs (v, e). The sum of the residues on each component must still be zero, and so the condition (4.1) must be satisfied.

It is in the compatibility condition (4.2) that the twisting shows up. This condition must be replaced by the condition

$$\mathrm{Res}_\sigma(v_1, e) + \lambda_e \mathrm{Res}_\sigma(v_2, e) = 0 \qquad (4.4)$$

if e is the edge directed from v_1 to v_2.

Clearly these are the only conditions on the residue numbers, and this gives a concrete description of all the twisted 1-forms for a given twisted canonical bundle L.

Using Riemann-Roch, we have the following elementary remark:

Lemma 4.5. *Let G be a simple sub-trivalent connected graph of genus g. Let L be a twisted canonical bundle on C_G. Then:*

(4.5.1) If L is the canonical bundle, $h^0(L) = g$.
(4.5.2) If L is not the canonical bundle, $h^0(L) = g - 1$.

Thus for a twisted canonical bundle which is not the canonical bundle, the $5g-5$ conditions (4.1) and (4.4) on the $6g-6$ residue numbers are independent.

There are many questions one can ask concerning twisted canonical bundles in general: when are they base-point-free, when are they very ample, etc. For our purposes we want to address the following problem: when is the linear series cut out on each component of C_G by a twisted canonical bundle complete? For this, we require a definition.

Definition 4.6. The *support* of a 1-cochain $\{\lambda_e\}$ defining a twisted canonical bundle L on C_G (by comparison with the canonical bundle) is the set of edges e with $\lambda_e \neq 1$. Let L be a twisted canonical bundle on C_G. The *twisted norm* of L is the minimum number k of edges e such that there exists a 1-cochain $\{\lambda_e\}$ defining L with exactly k λ's not equal to 1.

I.e., the twisted norm is the minimum cardinality of support among all 1-cochains defining a twisted canonical bundle. Thus the only twisted canonical bundle on C_G with twisted norm 0 is the canonical bundle itself.

Proposition 4.7. *Suppose that G is a simple trivalent 2-edge-connected graph. Suppose that $L \in \mathrm{Pic}^1(C_G)$, and that either:*

a) *L is the canonical bundle, i.e., L has twisted norm 0; or*
b) *L has twisted norm 2, but no 1-cochain defining L has support on exactly two incident edges; or*
c) *L has twisted norm at least 3.*

Then the map from $H^0(C_G, L)$ to $H^0(C_v, L_v)$ is surjective for all vertices v.

Proof. Assume first that L has twisted norm at least 3, and that C_G has genus g.

Fix a vertex v, and let e_1, e_2, and e_3 be the edges incident to v. Note first that the graph G is not disconnected by the removal of any 2 of these edges. Indeed, if it were, then it is in fact disconnected by the removal of the third, and this cannot happen by the assumption of 2-edge-connectedness.

Now to show that the restriction map from $H^0(C_G, L)$ to $H^0(C_v, L_v)$ is onto, it suffices to show that two of these edges, say e_1 and e_2, impose independent conditions on $H^0(C_G, L)$. Let G' be the graph obtained by removing e_1 and e_2; G' is a connected sub-trivalent graph, and its associated graph curve C' maps naturally to C_G, identifying the points as indicated by the two missing edges. The genus of C' is $g - 2$.

Now the subspace V of $H^0(C_G, L)$ consisting of those sections vanishing at the nodes corresponding to e_1 and e_2 is isomorphic to the space of global sections of the pull-back L' of L to C'. A 1-cochain defining L' is obtained by taking a 1-cochain for L and deleting the elements λ_{e_1} and λ_{e_2}. Therefore, since L has twisted norm at least 3, L' has twisted norm at least one. In particular, L' is not the canonical bundle of C', and so by Lemma 4.5, $H^0(C', L')$ has dimension $g-3$. This is two less than the dimension of $H^0(C_G, L)$ by Lemma 4.5, so the two edges do impose independent conditions.

This proves case c). the proof of case b) is identical; the hypothesis of this case is exactly to ensure that the pull-back bundle L' has twisted norm at least one. The proof of case a) is also similar, noting that both L and L' are the canonical bundles, so the dimension drops by the genus drop, which is two. In the case of the canonical bundle, this is a result of Bayer and Eisenbud (see [BE, Proposition 2.4]). □

Note that if L has twisted norm equal to 1, Then the sections of L have a base point, namely that node corresponding to the edge e where $\lambda_e \neq 0$, but all other λ's are zero.

The set of line bundles in $\mathrm{Pic}^1(C_G)$ with twisted norm at most k forms a closed set. The general line bundle has twisted norm at least g (since $\mathrm{Pic}^1(C_G)$ is a g-dimensional space).

Hence, if $g \geq 3$, the set of line bundles with twisted norm at least 3 forms a Zariski open set in $\mathrm{Pic}^1(C_G)$. In particular, the general twisted canonical bundle has twisted norm at least 3 and those that do form an irreducible set.

We call a twisted line bundle *component-complete* if the restriction map from $H^0(C_G, L)$ to $H^0(C_v, L_v)$ is surjective for all vertices v. It is easy to see that the above conditions characterize the component-complete twisted canonical bundles on C_G. Moreover, we see that the component-complete twisted canonical bundles form an irreducible subset of $\mathrm{Pic}^1(C_G)$: the union of the open set of bundles with twisted norm at least 3, with the single point representing the canonical bundle, and the union of the closed sets representing bundles with twisted norm 2 having a defining 1-cochain with support on exactly two incident edges. We will denote the component-complete twisted canonical bundles by $\mathrm{Pic}^1_{cc}(C_G)$.

5. Matroid representations associated to line bundles

In this section we want to mention a relationship between line bundles on graph curves and representations of a certain matroid associated to the graph.

Recall that a *matroid* M is a finite set E, together with a non-empty collection $\mathcal{I}$ of subsets of E (called the *independent sets*), such that any subset of $\mathcal{I}$ is also in $\mathcal{I}$, and which satisfies the following:

> If I and J are independent sets and $|J| > |I|$, then there exists an element $x \in J - I$ such that $I \cup \{x\}$ is independent.

The sets not in $\mathcal{I}$ are the *dependent* sets. For an introduction to the theory of matroids, one may consult [A], [We], or [Wi2].

An important example of a matroid is the *cutset* matroid of a graph G. This matroid is defined on the set of edges E; the dependent sets are those sets of edges which, when removed from G, will disconnect G. (Some authors call this the *bond* matroid of G.) An example of such a cutset is the set of edges E_v incident to a single vertex v of G: the vertex v is split off when these edges are removed.

We will call such a set of edges, incident to a single vertex v, the *star* of the vertex. Any matroid defined on the edges E of a graph G which has all the stars as dependent sets will be called a *star matroid.* The cutset matroid of the graph is the first example of a star matroid.

An *n-dimensional projective representation* of a matroid $M = (E, \mathcal{I})$ over a field $\mathbb{K}$ is a function $\phi : E \to \mathbb{P}^{\ltimes}_{\mathbb{K}}$ such that dependent sets in E go to dependent sets in $\mathbb{P}^{\ltimes}_{\mathbb{K}}$. The representation ϕ is *nondegenerate* if its image does not lie in a hyperplane of $\mathbb{P}^{\ltimes}$. The representation ϕ is *faithful* if independent sets in E go to independent sets in $\mathbb{P}^{\ltimes}_{\mathbb{K}}$. (The usual use of the word representation in matroid theory corresponds to a faithful representation; we want to allow non-faithful representations, where independent sets may become dependent, so we will abuse the standard language in this way.) Two representations are said to be *isomorphic* if they differ by a linear automorphism of the projective space.

Let $\phi : E \to \mathbb{P}^{\ltimes}$ be an n-dimensional projective representation of M over $\mathbb{K}$, and suppose that $\pi : \mathbb{P}^{\ltimes} \to \mathbb{P}^{\ltimes-\urcorner}$ is a linear projection, with center not intersecting $\phi(E)$. Then $\pi \circ \phi$ is an $(n-k)$-dimensional projective representation of M, since dependency is preserved under π. Such a representation will be called *projected*, and a representation which is non-degenerate and which is not projected will be called *linearly normal.*

We can now describe how representations of star matroids associated to a graph correspond to certain line bundles on the graph curve C_G. We assume that G is a connected simple trivalent graph of genus g.

Recall that $\mathrm{Pic}^1(C_G)$ is the set of isomorphism classes of line bundles L on C_G which have total degree identically one, i.e., $d_v(L) = 1$ for each vertex $v \in V$. From the results of the last section, this set of line bundles forms a component of the Picard group of C_G, and is in fact the component containing the canonical bundle ω_{C_G}.

Suppose first that $\phi : E \to \mathbb{P}^{\ltimes}$ is a non-degenerate projective representation of a star matroid of the graph G. Suppose further that it is 1-1 on each star, that is, the three edges incident to any vertex go to distinct points in $\mathbb{P}^{\ltimes}$. Note that for each vertex v of G, the three edges incident to v are dependent in this star matroid, and hence must be sent to collinear points in $\mathbb{P}^{\ltimes}$ by ϕ. Hence we can "extend" the map ϕ to a map $\psi : C_G \to \mathbb{P}^{\ltimes}$ by sending each component C_v onto the corresponding line in $\mathbb{P}^{\ltimes}$, such that the three nodes of C_v go to the three images of the corresponding edges. Moreover this extension is unique, since a linear map on $\mathbb{P}^{\nVdash}$ is determined by where it sends 3 points. Let L_ϕ be the hyperplane bundle $\mathcal{O}_{C_G}(1)$ for this extension, i.e., $L_\phi = \psi^*(\mathcal{O}_{\mathbb{P}^{\ltimes}}(1))$. Since each component C_v is mapped by ψ to a line, we have $d_v(L_\phi) = 1$ for all vertices v; moreover, L_ϕ is obviously component-complete. I.e., $L_\phi \in \mathrm{Pic}^1_{cc}(C_G)$.

Call a representation of a star matroid *star-separating* if it is 1-1 on each star.

If ϕ is linearly normal, then the map ψ cannot arise by mapping C_G to a higher-dimensional space (such that every component maps to a line) and then projecting; this would imply that the representation ϕ of this star matroid was projected. Hence if ϕ is linearly normal, we have that the map ψ arises from the full space of global sections of the hyperplane bundle L_ϕ, and that the dimension of the representation $n = h^0(L_\phi) - 1$.

By Lemma 4.5, if $L \in \mathrm{Pic}^1(C_G)$, then either $L = \omega_{C_G}$ (in which case $h^0(L) = g$) or $h^0(L) = g - 1$. Therefore:

Proposition 5.1. *Let G be a connected simple trivalent graph of genus g. Then the only star-separating linearly normal projective representations of star matroids of G are $(g-1)$- or $(g-2)$-dimensional. Moreover, if ϕ is a star-separating $(g-1)$-dimensional linearly normal projective representation of a star matroid, then the hyperplane bundle L_ϕ is the canonical bundle ω_{C_G}.*

Notice that if $\mathcal{M}_\infty$ and $\mathcal{M}_\in$ are two star matroids, such that the dependent sets of $\mathcal{M}_\infty$ are a subset of the dependent sets of $\mathcal{M}_\in$, then any representation of $\mathcal{M}_\in$ is also a representation of $\mathcal{M}_\infty$. We will say that a representation of a star matroid $\mathcal{M}$ is *maximal* if it cannot be extended to a "larger" star matroid.

Clearly two isomorphic representations give the same line bundle. Hence we have a constructed a map

$$\left\{ \begin{array}{c} \text{isomorphism classes of} \\ \text{maximal star-separating linearly normal} \\ \text{projective representations} \\ \text{of star matroids of } G \end{array} \right\} \longrightarrow \mathrm{Pic}^1_{cc}(C_G)$$

sending a representation ϕ to the line bundle L_ϕ.

Proposition 5.2. *Suppose that G is a simple trivalent 2-edge-connected graph. Then the above assignment of isomorphism classes of maximal star-separating linearly normal projective representations of star matroids of G to component-complete twisted canonical bundles on C_G is bijective.*

Proof. The inverse map is easily described. Let L be a component-complete twisted canonical bundle on C_G. Consider the map $\psi : C_G \to \mathbb{P}^n$ given by choosing a basis for the $n+1$-dimensional space of global sections of L. This map exists because L has no base points, and it maps each component C_v of C_G to a line, since L is component-complete. If E is the edge-set of G, then composing ψ with the identification of E with the nodes of C_G gives a map $\phi(L) : E \to \mathbb{P}^n$, which is 1-1 on each star. This induces a star matroid $\mathcal{M}(\mathcal{L})$ on E, by declaring the dependent sets to be those which are dependent in $\mathbb{P}^n$ after applying the map $\phi(L)$. The map $\phi(L)$ clearly becomes a maximal star-separating linearly normal projective representation of $\mathcal{M}(\mathcal{L})$. The isomorphism class of $\phi(L)$ depends only on L, and obviously by the construction we have both $L_{\phi(L)} \cong L$ and $\phi(L_\phi) \cong \phi$. □

The canonical bundle ω_{C_G} is component-complete, and hence corresponds to a representation of a star matroid. That matroid is the cutset matroid, by [BE, Proposition 2.3], which states that the conditions imposed on the canonical series by a set of nodes are independent if and only if that set of nodes does not disconnect the graph. Hence:

Corollary 5.3. *Let G be a simple trivalent 2-edge-connected graph of genus g. Then there is a unique $(g-1)$-dimensional maximal star-separating linearly normal projective representation of the cutset matroid of G, up to isomorphism.*

6. The Generic Star Matroid

This whole story is clarified in our view by introducing the scheme X_G of $(g-1)$-dimensional representations of star matroids of G. Consider the ambient space $(\mathbb{P}^{g-1})^{3g-3}$, representing (ordered) sets of $3g-3 = |E|$ points in a projective space of dimension $g-1$, which is the maximum possible dimension for representations as we have seen in Proposition 5.1. Think of homogeneous coordinates for such a set of points as a rectangular matrix of size $(g) \times (3g-3)$, whose columns index the points, and whose rows index the ordinary homogeneous coordinates of the points.

Consider the subscheme X of $(\mathbb{P}^{g-1})^{3g-3}$ defined by the vanishing of all 3×3 minors coming from 3 columns representing stars of points in G. That is, X is exactly the subscheme of those sets of points for which the stars are collinear. We call X the *scheme of representations* of star matroids of G. Let Y be the subset of X representing those sets of points which are not star-separating and linearly normal; the complement $X - Y$ maps to $\mathrm{Pic}^1_{cc}(C_G)$ via the associated representation. The fibers of this map are all orbits under the projective group, and therefore are irreducible.

Hence we see that, since $\mathrm{Pic}^1_{cc}(C_G)$ is irreducible, there is a "generic" component X' of $X - Y$, (namely that one dominating $\mathrm{Pic}^1_{cc}(C_G)$). For the general point in this component X', there is an induced matroid, which is constant on an open set of the component.

Definition 6.1. The matroid $\mathcal{S}(\mathcal{G})$ on the edges of the graph G of which the general element in X' is a representation is called the *generic star matroid* of G.

The scheme $X - Y$ itself is not irreducible; the cutset matroid and its representations provide us with a component different from X', since it has independent sets which do not appear in the other component (namely sets of g independent points).

Indeed, a naive parameter count gives that both the cutset matroid representations and the component X' dominating $\mathrm{Pic}^1_{cc}(C_G)$ have dimension g^2-1. For the cutset matroid, this is given by considering the unique canonical embedding of the graph curve C_G (if the curve is 3-edge-connected) and the $g^2 - 1$ projective automorphisms of the ambient space. For the generic star matroid component X', one has $g-1$ parameters for the choice of hyperplane in which the curve lives, $(g-1)^2 - 1 = g^2 - 2g$ projective automorphisms in that hyperplane, a final g parameters coming from the choice of twisted canonical bundle; this gives a total of $g^2 - 1$ also.

We do not know whethere there are other components of the scheme of representations X; however any others must lie in the set Y.

In our view the generic star matroid $\mathcal{S}(\mathcal{G})$ is a very natural object associated to the graph G. There does not seem to be any discussion of this matroid in the standard texts. It is not the cutset matroid, as we have seen above; for example, the cutset matroid has sets of g independent points, which the generic star matroid does not.

Let us give several small examples to fill out our understanding. Consider first the graph G_3 of genus 3 given by the 4 vertices and 6 edges of a tetrahedron. The generic star matroid for G_3 is the 2-uniform matroid; that is, the dependent sets are those sets of cardinality 3 or more. Indeed, the representations of the twisted canonical bundles are all in dimension $g-2=1$, so any three points will be collinear, hence dependent. Notice that any cutset of this matroid has at least 3 edges, so that the cutset matroid is strictly "smaller" (in the sense of dependent sets) than the generic star matroid.

Now consider the graph G_4 of genus 4 which is a prism over a triangle. This graph has 6 vertices $\{a,b,c,d,e,f\}$ and 9 edges $\{ab, ac, bc, de, df, ef, ad, be, cf\}$. The cutset matroid has, as a dependent set, the cutset $\{ad, be, cf\}$. The generic star matroid however is realized by the map of the graph curve C to $\mathbb{P}^{\nvDash}$ sending each of the 6 components to a line. The reader can easily check that for a generic choice of 6 lines $\{a,b,c,d,e,f\}$ in $\mathbb{P}^{\nvDash}$, the intersection points $\{a\cap d, b\cap e, c\cap f\}$ are *not* collinear. Hence the cutset matroid is not smaller than the generic star matroid in this case.

In general, there are always sets of cardinality g which span canonical space, and hence are independent for the cutset matroid; any such set however must be dependent in the generic star matroid, since it is faithfully represented in dimension $g-1$. Therefore, the generic star matroid is never smaller than the cutset matroid.

Thus for this example of the graph G_4, we see that the generic star matroid and the cutset matroid are incomparable.

Finally take the other graph $K_{3,3}$ of genus 4, which is the complete bipartite graph with 6 vertices, 3 in each part. The canonical image of the associated graph curve consists of 6 lines on a smooth quadric, 3 in each of the two rulings. Notice that the only cutsets of cardinality 3 are the stars themselves. Since any subset of cardinality 4 or more is dependent in the generic star matroid, we see that every cutset must be dependent in the generic star matroid. Thus in this example, the cutset matroid is strictly smaller than the generic star matroid.

We consider the determination of the generic star matroid for a general connected trivalent graph G to be a fundamental problem in the projective geometry of point sets. Indeed, any nontrivial statement about the generic star matroid of any particular graph would yield theorems similar in spirit to those of Pappus and Desargues.

7. Planar Graphs

Let G be a connected simple trivalent graph of genus g. Suppose that this trivalent graph G is embedded in the 2-sphere S^2 to form a map M, with a set of faces (or countries) F, bordered by the set of edges E of G and the set of vertices V of G.

Lemma 7.1. *There are $g+1$ faces in the map M.*

Proof. Since G is trivalent, we have $|V| = 2g-2$ and $|E| = 3g-3$. Euler's equation for the decomposition of the sphere gives

$$|V| - |E| + |F| = 2,$$

forcing $|F| = g+1$. □

Index the faces of F from 0 to g. Let k be a field. Consider the projective space $\mathbb{P}^g = \mathbb{P}^g_k$ over k with coordinates also indexed from 0 to g. For the i-th face of F, denote the corresponding coordinate point of $\mathbb{P}^g$ by f_i. Thus f_i is the point with homogeneous coordinates all 0, except for a 1 in the i-th position.

Note that each vertex of V lies a priori on three faces of F. If these faces are different, the vertex v defines a plane in $\mathbb{P}^g$, namely the span of the three points f_i corresponding to the three faces on which v lies.

There is a convenient criterion on the graph G which ensures that, for each vertex v, the three faces of the map on which v lies are different. Recall that a graph is *n-edge-connected* or simple *n-connected*, if the removal of less than n edges never disconnects the graph. A graph is 1-connected if and only if it is connected. A graph is 2-connected if and only if there is no disconnecting edge.

Lemma 7.2. *Let G be a connected simple trivalent planar graph, forming the map M. Then G is 2-connected if and only if for every vertex v of G, the three faces of the map M on which v lies are different.*

Proof. Let v be a vertex of G such that two of the faces on which v lies are the same. Then removing the edge emanating from v which separates (locally) those two regions will disconnect G. Conversely, if removing an edge e separates G, then the two faces on each side of e are actually the same, so the two vertices of e fail to lie on 3 distinct faces of the map M. □

So if G is trivalent and 2-connected, each vertex v of G does define a plane P_v. Note that in general, if G is trivalent and 2-connected, then G is 3-connected if and only if any two distinct faces meet along at most one edge.

Now assume that G is trivalent and 2-connected, so that by the above lemma each vertex v of G lies on 3 distinct faces of the map M. Then v determines the plane P_v, the span of the 3 coordinate points associated to the 3 faces. Define the surface S_M to be the union of the planes P_v over all vertices v in V:

$$S_M = \bigcup_{v \in V} P_v.$$

This "face surface" S_M is a union of planes, one plane for each vertex. It is a geometric realization of the combinatorial object M, and we would like it to reflect as much as possible the combinatorics of the planar map. In particular, if v and w are distinct vertices, we would like the planes P_v and P_w to be distinct, and if so, to intersect along a line if and only if v

and w are joined by an edge in G. Indeed, we would like the edges of G to encode the line intersections of the planes of S_M. If these conditions are true, we will say that S_M *faithfully realizes* the map M.

Of course, if v and w are joined by the edge e, then the planes P_v and P_w both contain the line spanned by the coordinate points corresponding to the two faces on either side of e. If two planes P_v and P_w intersect in a line, they must intersect in a coordinate line, and so the vertices v and w must both lie at the boundary of two faces f and g of the map M. If v and w are not joined by an edge, then f and g share two edges in common: one emanating from v and one from w. Therefore, G is not 3-connected. In particular, if the two planes P_v and P_w are the same, they certainly share lines, and so G cannot be 3-connected. Conversely, if G is not 3-connected, then there are two faces f and g which meet along more than one edge. This implies that there are two vertices which are not joined, but which both have the line joining the coordinate points of f and g in their associated planes. This proves the following.

Proposition 7.3. *Let G be a 2-connected simple trivalent planar graph, inducing the map M. Then the surface S_M faithfully realizes the map M if and only if G is 3-connected.*

For the rest of this article we will assume that G is a 3-connected simple trivalent planar graph, so that the surface S_M faithfully realizes the map M.

Let H_∞ be the hyperplane of $\mathbb{P}^g$ defined by setting the sum of the coordinates to zero:

$$H_\infty : \sum_{i=0}^{g} x_i = 0.$$

Note that none of the face points f_i lie on H_∞, so H_∞ cuts each plane P_v in a line C_v. Define the curve C_G to be the union of the lines C_v over all vertices v in V (or, equivalently, to be the intersection of S_M with H_∞):

$$C_G = S_M \cap H_\infty = \bigcup_{v \in V} C_v.$$

The curve C_G is a union of $2g-2$ lines in a projective space of dimension $g-1$ (namely H_∞). Two lines C_v and C_w meet in C_G exactly when the corresponding vertices v and w are connected by an edge in G. Indeed, we see that this embedded curve is exactly the associated graph curve C_G to G, realized in a projective space.

Note that if e is an edge bounding faces i and j, then the node n_e of the graph curve C_G corresponding to e is the point in H_∞ with homogeneous coordinates equal to 1 in position i, -1 in position j, and 0 in all other positions.

The curve C_G is embedded in the projective space $H_\infty \cong \mathbb{P}^{g-1}$ with sections of a bundle which is degree 1 on each $\mathbb{P}^1$ component of C_G: the image of each component is a line. This bundle is actually therefore the canonical bundle of the stable curve C_G, since as in the general theory the canonical bundle is the only bundle with this total degree, having g linearly independent global sections.

Therefore the embedding of C_G into $\mathbb{P}^{g-1}$ is given by the complete linear system of regular 1-forms on C_G, i.e., it is the canonical embedding of C_G. (See [BE] for a discussion in general of the canonical embedding for graph curves of trivalent 3-connected graphs.)

The upshot of these remarks is that any canonical divisor on C_G, that is, a divisor of zeroes of a regular 1-form on C_G, is exactly the intersection of C_G with a hyperplane in H_∞. Moreover, since two 1-forms have the same zero locus if and only if they are multiples of each other, we see that the hyperplanes in H_∞ are in $1-1$ correspondence with the 1-dimensional subspaces $\mathbb{P}\Omega_{\mathbb{G}}$ of the space Ω_G of regular 1-forms on C_G.

If x is a point of a component C_v of C_G identified to a node, and ω is a regular 1-form on C_G, then the restriction of ω to C_v near x can be written as $f(z)dz/z$, where z is a local coordinate on C_v at x and f is a regular function. In other words, at the identified point x, the form dz/z is a local basis for the space of regular 1-forms. Hence we see that the 1-form ω has a "zero" at x if and only if the coefficient function f in the above expression is zero at x, i.e., if $f(0) = 0$. This is equivalent to having the residue of ω_v being zero at the point x. Since the residue of ω_v at x is opposite to the residue of the corresponding 1-form at the point identified to x by (4.2), the regular 1-form ω vanishes at the node of C_G to which these points are identified if and only if these residues are zero.

Rephrased in terms of hyperplanes, this is the following.

Proposition 7.4. *Let e be an edge of G, incident to vertices v and w, corresponding to the node n_e of C_G. Let the hyperplane $H \subset H_\infty$ correspond the the regular 1-form $\omega \in \Omega_G$. Then H passes through n_e if and only if the residue numbers $\mathrm{Res}_\omega(v,e)$ and $\mathrm{Res}_\omega(w,e)$, which are opposite in any case by (4.2), are both zero. The set of hyperplances in H_∞ which do not pass through any node of C_G are in $1-1$ correspondence with the 1-dimensional subspaces of Ω_G which are spanned by regular 1-forms all of whose residue numbers are non-zero.*

Because of this proposition, we will denote the set of hyperplanes in H_∞ by $\mathcal{H}_M$, and those which do not pass through any node of C_G by $\mathcal{H}'_M$. Let us also denote the subset of Ω_G consisting of 1-forms all of whose residue numbers are non-zero, by Ω'_G. The above proposition shows that $\mathcal{H}'_M$ is in $1-1$ correspondence with $\mathbb{P}(\Omega'_{\mathbb{G}})$.

When the graph G is planar, there is another convenient description of the space Ω_G of regular 1-forms on C_G, which we now develop. Fix once and for all an orientation of the 2-sphere S^2, giving therefore an orientation to each face of the map M.

Given a face f, we can then orient (clockwise, say) the edges at its boundary, which puts directions on two of the three edges emanating from each vertex on the face f. (The edge not on the boundary of the face f is not directed by this procedure.) Define a regular 1-form $\omega_f \in \Omega_G$ by assigning its residue numbers as follows.

If v is not on the boundary of f, set $\mathrm{Res}_{\omega_f}(v,e) = 0$ for all 3 edges e emanating from v.

If v is on the boundary of f, and e is the edge emanating from v which is not on f, set $\mathrm{Res}_{\omega_f}(v,e) = 0$.

If v is on the boundary of f, and e is an edge emanating from v on the boundary of f, set $\mathrm{Res}_{\omega_f}(v,e) = 1$ if e is directed away from v, and set $\mathrm{Res}_{\omega_f}(v,e) = -1$ if e is directed towards v.

It is easily checked that these assignments satisfy the residue number conditions (4.1) and (4.2), and define a regular 1-form $\omega_f \in \Omega_G$ on C_G. The following lemma is proved in [M1, (Lemma 2.5)].

Lemma 7.5. *The 1-forms $\{\omega_f : f \text{ is a face of } M\}$ span the space of regular 1-forms Ω_G on C_G. Moreover, the only linear relation among the ω_f's is that their sum is zero.*

8. Colorings of the planar map

Suppose now that the field k over which we have been working is the finite field $\mathbb{F}_{\shortparallel}$ with q elements. Let Q be a q-coloring of the map M, that is, an assignment of q symbols to the faces of M which has the property that if two faces share an edge, then they are assigned different symbols. We take the symbols to come from the field $\mathbb{F}_q$, and therefore we will consider a q-coloring of M to be a function Q from the set F of faces of M to the field $\mathbb{F}_q$, satisfying the coloring condition that if two faces share an edge, then Q has different values on them. Denote by $\mathrm{Col}_q(M)$ the set of q-colorings of M.

To every function $Q : F \to \mathbb{F}_q$ we may assign the regular 1-form

$$\omega(Q) = \sum_{f \in F} Q(f)\omega_f$$

which is a 1-form on the graph curve C_G.

Let us analyze the coloring condition in terms of this regular 1-form.

Proposition 8.1. *Let $Q : F \to \mathbb{F}_q$ be a function. Then Q is a q-coloring of M if and only if the regular 1-form $\omega(Q)$ has all non-zero residue numbers.*

Proof. Suppose f and g are two faces of M sharing the edge e and the two vertices v and w at the ends of e. The residue numbers $\mathrm{Res}_{\omega(Q)}(v, e)$ and $\mathrm{Res}_{\omega(Q)}(w, e)$ are determined only by the coefficients of ω_f and ω_g, since if h is a face not having e on its boundary, then $\mathrm{Res}_{\omega_h}(v, e) = \mathrm{Res}_{\omega_h}(w, e) = 0$. Suppose that we have oriented the sphere such that for the face f, the edge e is directed from v to w; then necessarily for the face g, e is directed from w to v. Then by the definition of ω_f and ω_g, we have

$$\mathrm{Res}_{\omega_f}(v, e) = \mathrm{Res}_{\omega_g}(w, e) = 1 \text{ and } \mathrm{Res}_{\omega_f}(w, e) = \mathrm{Res}_{\omega_g}(v, e) = -1.$$

Therefore the residue numbers for $\omega(Q)$ are

$$\mathrm{Res}_{\omega(Q)}(v, e) = Q(f) - Q(g) \text{ and } \mathrm{Res}_{\omega(Q)}(w, e) = Q(g) - Q(f).$$

The coloring condition on Q is that $Q(f) \neq Q(g)$ whenever f and g share an edge. Therefore the coloring condition is equivalent to both of these residue numbers given above being non-zero. Since this is true for any edge e, the coloring condition is equivalent to the regular 1-form ω_Q having all non-zero the residue numbers. □

Using our notation from above, we see that

$$Q \in \mathrm{Col}_q(M) \text{ if and only if } \omega(Q) \in \Omega'_G.$$

Algebraically, let $\mathbb{F}_q^{\mathbb{F}}$ denote the vector space of functions from the set of faces F of the map to the field $\mathbb{F}_q$. The assignment

$$Q \longmapsto \omega(Q)$$

gives an $\mathbb{F}_q$-linear map from $\mathbb{F}_q^{\mathbb{F}}$ to Ω_G, which is onto, with kernel the 1-dimensional space of constant functions on F. The inverse image of the set Ω'_G of 1-forms all of whose residue numbers are nonzero is exactly the set $\mathrm{Col}_q(M)$ of q-colorings of the map M.

This brings us to our main result, which is a consequence of the preceding discussion.

Theorem 8.2. *A q-coloring of the planar map M exists if and only if there is a regular 1-form on the graph curve C_G over $\mathbb{F}_q$ all of whose residue numbers are non-zero (i.e., if and only if $\Omega'_G \neq \emptyset$).*

Indeed, we see that

$$\text{the number of } q\text{-colorings of } M = |\,\mathrm{Col}_q(M)| = q|\Omega'_G| \tag{8.3}$$

since the fibers of the map from q-colorings of M onto Ω'_G have cardinality q. Expressed geometrically in terms of the hyperplanes in H_∞, Proposition 7.4 gives that

$$|\Omega'_G| = (q-1)|\mathcal{H}'_M| \tag{8.4}$$

since every 1-dimensional subspace has $q-1$ spanning vectors. This gives

$$|\operatorname{Col}_q(M)| = q(q-1)|\mathcal{H}'_M|. \tag{8.5}$$

Let us apply these results in the cases of low q, where the interest is highest. All of these results are well-known and rather simple; however the unified treatment is the point here. For graph-theoretic discussions of the statements any number of books may be consulted, for example [Wi1] or [A].

We begin with $q = 2$; no map can be 2-colored, so $\operatorname{Col}_2(M) = \emptyset$. On the other hand it is easy to see that Ω'_G is also empty if $q = 2$. The only nonzero number in $\mathbb{F}_{\nvDash}$ is 1; therefore any 1-form in Ω'_G must have every residue number equal to 1. But $1 + 1 = 0$, so that if we fix a vertex v, the sum of the 3 residue numbers at v is 1; hence the assignment of all residue numbers equal to 1 does not give a regular 1-form on C_G.

It is more interesting in the case $q = 3$. Write $\mathbb{F}_{\nVDash} = \{\nvdash, \pm\nVdash\}$; we see that any 1-form in Ω'_G has only ± 1 as residue numbers. Suppose $\omega \in \Omega'_G$. The only way for the three residue numbers of ω at any vertex v to sum to 0 is if all three numbers are the same: either all are 1 or all are -1. Therefore the residue numbers are actually associated to simple vertices, not vertex-edge pairs. Moreover, since the two residue numbers on any edge must sum to zero, if two vertices are connected by an edge they must have opposite residue numbers. Thus an element of Ω'_G is determined exactly by a 2-coloring of the vertices V of G, namely a function from V to $\{\pm 1\}$ such that if vertices v and w are joined by an edge of G, the function has different values on them. It is elementary to see that a 2-coloring of the vertices of a planar graph G is possible if and only if every face f of the map M has an even number of edges. If this is true, then the 2-coloring is determined by the color of any one vertex v_0: the other vertices are colored according to the parity of the distance from v_0, which is well-defined with the even-face assumption. Thus we recover the well-known 3-coloring criterion for a planar map:

Corollary 8.6. *A planar map M is 3-colorable if and only if every face of M has an even number of edges. If this is so, there are exactly 6 3-colorings of M.*

Finally let us treat the most interesting case of $q = 4$. Write $\mathbb{F}_{\not\trianglerighteq} = \{\nvdash, \ni, , \}$. Any 1-form in Ω'_G has only the nonzero numbers $\{a, b, c\}$ as residues. Since the field $\mathbb{F}_{\not\trianglerighteq}$ has characteristic 2, and the sum of the two residue numbers at each edge must be zero, we see that the two residue numbers at each edge must be the same. Therefore the residue numbers for any ω are determined by assigning numbers to each edge. Suppose that $\omega \in \Omega'_G$. Then the three (non-zero) residue numbers at each vertex must sum to zero, and it is easy to see that this means that at each vertex, all three non-zero numbers a, b, and c must each be used (once). In other words, one obtains a 1-form in Ω^1_G exactly with a 3-coloring of the edges E of G, namely, a function from E to $\{a, b, c\}$ such that if two edges share a vertex, the function has different values on them. Therefore we obtain the (slightly less well-known but equally easy) 4-coloring criterion for a planar map:

Corollary 8.7. *The faces of a planar map M are 4-colorable if and only if the edges of the map are 3-colorable. The number of 4-colorings of the faces is 4 times the number of 3-colorings of the edges.*

Since we have related colorings of G to regular 1-forms on C_G, and we have seen that regular 1-forms on C_G correspond to hyperplanes in the space H_∞, we can express the coloring results in more geometric terms. The next corollary follows directly from Proposition 7.4 and Theorem 8.2

Corollary 8.8. *The faces of a planar map M bounded by a 3-connected simple trivalent graph G can be colored by q colors if and only if there is a hyperplane H in the space H_∞ (which contains the canonically embedded graph curve C_G) which is defined over $\mathbb{F}_{\shortparallel}$ and which misses every node of the graph curve C_G.*

From this point of view, the 4-color theorem says that a sufficiently general hyperplane exists over $\mathbb{F}_{\ntrianglerighteq}$ for any planar graph. Therefore one may view the 4-color theorem as a "Bertini"-type result!

9. The chromatic function

As noted above, Proposition 7.4 and Theorem 8.2 show us that the existence of q-colorings of the map M is related to the existence of hyperplanes in H_∞ missing the nodes of the graph curve C_G. It is possible to count these hyperplanes (those of $\mathcal{H}'_M$), therefore counting possible colorings, by studying the configuration of the set of nodes.

For any set of nodes N of C_G, denote by $\mathcal{L_N}$ the (projective) linear span of N.

It is convenient to introduce the *dual* graph G' of G, with vertices equal to the faces in F and edges corresponding to the edges in E, now viewed as joining faces instead of joining vertices. A set of nodes N of C_G can then be viewed also as a set of edges E'_N of the dual graph G', which form (with their adjacent vertices) a subgraph G'_N of the dual graph G'.

Form the matrix A_N whose columns are the homogeneous coordinates of the nodes in N. The dimension of the linear span $\mathcal{L_N}$ of the nodes is then one less than the rank of this matrix (since we are speaking of projective dimension).

This matrix A_N has, in the column for a node n corresponding to an edge e, one 1 and one -1, in the rows corresponding to the faces of G adjacent to the given edge e. In terms of the dual graph G', the 1 and -1 occur in the rows corresponding to the two vertices of G' which lie on the edge e. Thus this matrix A_N is exactly the matrix for the boundary operator in homology, from the 1-chains of the subgraph G'_N to the 0-chains of the subgraph G'_N (given a suitable orientation of the dual graph). Therefore the kernel of A_N is identified with $H_1(G'_N)$, and the cokernel of A_N is identified with $H_0(G'_N)$. Hence the rank of A_N is equal to $|N| - rank H_1(G'_N)$, and also to $|V(G'_N)| - rank H_0(G'_N)$, where $V(G'_N)$ is the set of vertices of G'_N.

Thus we have the following formulae for the dimension of the linear span $\mathcal{L_N}$ of a set of nodes N:

$$dim(\mathcal{L}_N) = |N| - rank H_1(G'_N) - 1$$

or

$$dim(\mathcal{L}_N) = |V(G'_N)| - rank H_0(G'_N) - 1,$$

i.e.,

$$dim(\mathcal{L}_N) = \#\text{of vertices of } G'_N - \#\text{of connected components of } G'_N - 1.$$

If N is empty, then we will set $dim(\mathcal{L}_N) = -1$; this implies that the empty graph has trivial H_1.

Thus the dimension of the space of hyperplanes in H_∞ which pass through the nodes of N is

$$g-1-|N|+rank H_1(G'_N).$$

Fix a prime power q. Working over the field $\mathbb{F}_q$ with q elements, a projective space of dimension n has

$$\mathcal{P}_q(n)=\frac{q^{n+1}-1}{q-1}$$

elements.

It is now trivial to find a formula for the number $\mathcal{H}'_M$ of hyperplanes of H_∞ which contain no node of the graph curve C_G. For this we simply use the inclusion-exclusion principle: one then obtains

$$|\mathcal{H}'_M|=\sum_{k=0}^{3g-3}\sum_{|N|=k}(-1)^k\mathcal{P}_q(g-1-|N|+rank H_1(G'_N)).$$

One realizes immediately that this sum is simply over all sets of nodes N, and so we can re-write this as

$$|\mathcal{H}'_M|=\sum_{N\subset E}(-1)^{|N|}\mathcal{P}_q(g-1-|N|+rank H_1(G'_N)).$$

Given the results of the previous section, this is actually a version of the chromatic function for the map M. In graph theory, the chromatic function of a map M is the function $c_M(n)$ whose value at an integer n is the number of ways of coloring the map M with n colors. Using our results concerning the hyperplanes and the residues of 1-forms, we see that

$$c_M(q)=q(q-1)|\mathcal{H}_M|=q(q-1)\sum_{N\subset E}(-1)^{|N|}\mathcal{P}_q(g-1-|N|+rank H_1(G'_N))$$

for any prime power q.

Using the definition of $\mathcal{P}_q$, the factor of $q-1$ can be brought into the sum to obtain

$$c_M(q)=q\sum_{N\subset E}(-1)^{|N|}(q^{g-|N|+rank H_1(G'_N)}-1).$$

Since

$$\sum_{N\subset E}(-1)^{|N|}=0$$

by the binomial theorem, this can be re-written as

$$c_M(q)=q\sum_{N\subset E}(-1)^{|N|}q^{g-|N|+rank H_1(G'_N)}.$$

It is interesting to compare this with the standard formula for the chromatic function of a graph. Coloring the planar map M is equivalent to giving a coloring of the vertices of the dual graph G'. The chromatic function of any graph G' is the function $\pi(G',\lambda)$ which is the number of colorings of G' with λ colors. A formula for the chromatic function of a graph is developed in [O, Section 14.4]; it is

$$\pi(G',\lambda)=\sum_{s,c}(-1)^s p(s,c)\lambda^c,$$

where $p(s,c)$ is the number of subgraphs H of G' with s edges and c connected components. This reduces to the above formula in our case of $\lambda=q$, because of the planarity (giving the Euler formula) and the obvious relationship between the arithmetic genus g and the ranks of the homology groups.

References

[A] M Aigner: "Combinatorial Theory". Springer-Verlag, New York, 1979.

[BE] D. Bayer and D. Eisenbud: "Graph Curves", preprint.

[CF1] C. Ciliberto and A. Franchetta: "Curve poligonali, grafi e applicazione di Gauss", preprint.

[CF2] C. Ciliberto and A. Franchetta: "L'applicazione di Gauss per grafi trivalenti", preprint.

[CHM] C. Ciliberto, J. Harris, and R. Miranda: "On the surjectivity of the Wahl map", Duke Math. Journal, vol. 57, (1988), 829-858.

[CM1] C.Ciliberto and R. Miranda: "On the Gaussian map for canonical curves of low genus". Duke Math. J., vol. 61, (1990), 417-443.

[CM2] C.Ciliberto and R. Miranda: "Gaussian maps for certain families of canonical curves". In: *Complex Projective Geometry*, edited by G. Ellingsrud, C Peskine, G. Sacchiero, and S.A. Stromme, London Mathematical Society Lecture Note Series No. 179 (1992), 106-127.

[CV] C. Ciliberto and A. Verra: "On the surjectivity of the Gaussian map for Prym-canonical line bundles on a general curve", preprint.

[DM] J. Duflot and R. Miranda: "The Gaussian map for rational ruled surfaces", to appear in the Transactions of the AMS.

[M1] R. Miranda: "The Gaussian map for certain planar graph curves", in: Algebraic Geometry: Sundance 1988. Contemporary Math., vol. 116 (1991), 115-124.

[M2] R. Miranda: "Six lectures on graph curves and curves on K3 surfaces." In: Lectures on Riemann surfaces, Proceedings of the College on Riemann surfaces, International Centre for Theoretical Physics, Trieste, Italy, 9 Nov. - 18 Dec. 1987; pp. 119 - 176. World Scientific (1989).

[M3] R. Miranda: "Colorings of Planar Maps and Residues of 1-forms". In: *Methods in Module Theory*, G. Abrams, J. Haefner, and K.M. Rangaswamy, eds. Lecture Notes in Pure and Applies Mathematics, Marcel Dekker, Vol. 140, (1993), 237-248.

[O] O. Ore: "Theory of Graphs". AMS Colloquium Publications, vol. XXXVIII (1962).

[Wa] J. Wahl: "The Jacobian Algebra of a graded Gorenstein singularity". Duke Math. Journal, vol. 55 (1987), 843-871.

[We] D.J.A. Welsh: "Matroid Theory". London Mathematical Society Monograph No. 8, Academic Press, London, New York, San Franscisco, (1976).

[Wi1] R. Wilson: "Introduction to Graph Theory". Academic Press (1972).

[Wi2] R. Wilson: "An Introduction to Matroid Theory". American Math. Monthly, Vol. 80, No. 5 (May 1973), 500-525.

Dept. of Mathematics, University of Rome II, Tor Vergata, Via O. Raimondo, 00173 Rome, ITALY

E-mail address: ciliberto@mat.utovrm.it

Dept. of Mathematics, Colorado State University, Ft. Collins, CO 80523, USA

E-mail address: miranda@riemann.math.colostate.edu