Colorings of Planar Maps and Residues of 1-forms

Rick Miranda [1]
Dept. of Mathematics
Colorado State University
Ft. Collins, CO 80523

Abstract

Coloring properties of planar maps determined by planar graphs are examined in terms of algebro-geometric properties of an associated degenerate algebraic curve 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. Although no new coloring theorems are obtained, it is hoped that exposing this apparently unexplored relationship between graph theory and algebraic geometry will be useful in other contexts.

1 Introduction

In this article I would like to present some relationships between the theory of graphs and the theory of algebraic curves. For the remainder of the article a graph will mean a simple graph, i.e., a graph without loops or multiple edges. Formally, this would be a pair (V, E), where V is a set (of vertices) and E (the set of edges) is a subset of $V \times V - \Delta$, where Δ is the diagonal in the product $V \times V$. The *valency* of a vertex $v \in V$ is the number of edges containing v. A vertex is *trivalent* if there are 3 edges containing it, and a graph is *trivalent* if every one of its vertices are. To every 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

[1]Research supported in part by the NSF under grant DMS-9104058.
Keywords: graphs, colorings, algebraic curves, forms.
Subject Classifications: 05C15,14H99,14G15

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, and [DM] for a discussion of Gaussian maps for curves lying on ruled surfaces. The basic reference for Gaussian maps is Wahl's paper [Wa].

In another article [M1] I have shown that the Gaussian map for a planar graph curve (one associated to a planar graph G) satisfying some additional mild conditions has corank 1; the graph-theoretic import of this is that, under suitable hypotheses, a planar graph can be embedded in the sphere in only one way. This is a theorem of Whitney, and I find it interesting that such graph-theoretic results can be obtained via techniques of algebraic geometry.

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.

I would like to thank Ciro Ciliberto for several useful conversations on this subject.

2 Graphs, planar maps, and graph curves

Let G be a trivalent graph. 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. We assume that there are $g+1$ faces in the map M.

Since G is trivalent, we have

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

and Euler's equation for the decomposition of the sphere gives

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

so that since $|F| = g+1$, we have

$$|V| = 2g-2 \text{ and } |E| = 3g-3.$$

We will call the integer g the *arithmetic genus* of the graph G (or of the map M); note that this differs from the standard definition of "genus" in graph theory.

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* 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 2.1 *Let G be a planar trivalent 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. Q.E.D.

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:

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

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

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

An abstract variety S'_M can be formed from the map M by taking a union of planes, joined as the map M indicates. The faithful realization of M by S_M is exactly equivalent to having the map of S'_M to projective space given by the line bundle which has degree 1 on each plane, be a closed embedding.

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 L_v. Define the curve C_G to be the union of the lines L_v over all vertices v in V (or, equivalently, to be the intersection of S_M with H_∞):

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

The curve C_G is a union of $2g-2$ lines in a projective space of dimension $g-1$ (namely H_∞). Two lines L_v and L_w meet in C_G exactly when the corresponding vertices v and w are connected by an edge in G. Thus the edge set E corresponds to the set of nodes of C_G. If e is an edge of G, we denote by n_e the corresponding node of C_G. Note that if e is an edge bounding faces i and j, then n_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.

Each line of C_G meets exactly 3 other lines, and so C_G, as an abstract variety, is a stable curve in the sense of Mumford. As an abstract variety, it is completely determined by G: one component isomorphic to $\mathbb{P}^1$ for each vertex of G, with the components meeting as the edges of G indicate. Since the automorphism group of $\mathbb{P}^1$ is triply transitive over any field, there are no essential choices in the way that the components are put together. This curve C_G is called a *graph curve*. In general, a graph curve is a union of smooth rational curves isomorphic to $\mathbb{P}^1$, each meeting exactly 3 others in ordinary nodes.

3 One-forms on graph curves

The curve C_G is embedded in the projective space $H_\infty \cong \mathbb{P}^{g-1}$ with sections of the 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. To be precise, a regular 1-form on C_G is a 1-form on each component, with at worst simple poles at the nodes; at the nodes, the two residues must cancel for the compatibility condition to hold. For more details, the reader may consult [CHM] or [M1]. In particular, since the bundle of 1-forms on each component has degree -2, but we are allowing 3 simple poles, we obtain the bundle of degree 1 on each component. In fact the embedding of C_G into $\mathbb{P}^{g-1}$ is given by the complete linear system of these "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_∞.

Therefore to understand these intersections, we need to understand in more detail the space Ω_G of regular 1-forms on C_G and in particular the zeroes of a regular 1-form.

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 L_v of C_G (corresponding to a vertex v of G), denote by ω_v the restriction of ω to the component L_v. As noted above, ω_v is a rational 1-form on L_v with at worst simple poles at the three nodes of C_G lying on L_v. Let e be an edge of G emanating from v, so that n_e is one of the three nodes of C_G lying on the component L_v. Denote by $R_\omega(v, e)$ the residue of ω_v at the point n_e. Thus to every regular 1-form ω we have attached $6g-6$ residue numbers, namely the elements $R_\omega(v, e)$ of the field k as (v, e) ranges over incident vertex-edge pairs.

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

$$R_\omega(v, e_1) + R_\omega(v, e_1) + R_\omega(v, e_1) = 0 \tag{3.3}$$

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

$$R_\omega(v_1, e) + R_\omega(v_2, e) = 0 \tag{3.4}$$

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

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 $R_{\omega_f}(v, e) = 0$ 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 $R_{\omega_f}(v, e) = 0$.
>
> If v is on the boundary of f, and e is an edge emanating from v on the boundary of f, set $R_{\omega_f}(v, e) = 1$ if e is directed away from v, and set $R_{\omega_f}(v, e) = -1$ if e is directed towards v.

It is easily checked that these assignments satisfy the residue number conditions, 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 3.5 *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.*

If one chooses a coordinate z on one of the components L_v of C_G such that the 3 nodes on L_v are at $z = 0$, 1, and ∞, then any regular 1-form ω on C_G has, on L_v, the form

$$\omega_v = \frac{az + b}{z(z-1)} dz.$$

The residue of ω_v at $z = 0$ is $-b$, and the residue of ω_v at $z = 1$ is $a + b$. At $z = \infty$, we use the coordinate $w = 1/z$, and the formula for ω_v is $\frac{a+bw}{w(w-1)} dw$; the residue at $w = 0$ is then $-a$. We see from these formulae that ω has a "zero" on the component L_v at a node n_e if and only if the residue number $R_\omega(v, e) = 0$.

Let us now return to the embedded description of C_G as lying in the projective space H_∞. Since C_G is canonically embedded there, the intersection of C_G with a hyperplane in H_∞ is exactly the locus of zeroes of a regular 1-form $\omega \in \Omega_G$. 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 subpaces $\mathbb{P}\Omega_G$ of the space Ω_G of regular 1-forms on C_G.

In particular, we have the following.

Proposition 3.6 *The set of hyperplanes 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 $\mathcal{H}^1_M$. Let us also denote the subset of Ω_G consisting of 1-forms all of whose residue numbers are non-zero, by Ω^1_G. The above proposition shows that $\mathcal{H}^1_M$ is in $1-1$ correspondence with $\mathbb{P}\Omega^1_G$.

4 Colorings of the planar map

Suppose now that the field k over which we have been working is the finite field $\mathbb{F}_q$ 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 may as well take the symbols to come from the field $\mathbb{F}_q$; 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.

Hence to every q-coloring Q of M 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.

We want to analyze the coloring condition in terms of this regular 1-form. 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 $R_{\omega(Q)}(v,e)$ and $R_{\omega(Q)}(w,e)$ are determined only by the coefficients of ω_f and ω_g; if h is a face not having e on its boundary, then $R_{\omega_h}(v,e) = R_{\omega_h}(w,e) = 0$. 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

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

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

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

Now the coloring condition says that $Q(f) \neq Q(g)$; hence both of these residue numbers are non-zero. This is true for any edge e, so we see that a q-coloring of the map M gives a 1-form on C_G, all of whose residue numbers are non-zero. I.e., the 1-form $\omega(Q)$ lies in the set Ω_G^1.

Note that formulas above for the residue numbers of ω_Q involve the differences of values of Q at the faces. Therefore if α is any element from the field $\mathbb{F}_q$, and if we define a new coloring $Q'(f) = Q(f) + \alpha$ for any face f, we will obtain the same regular 1-form on C_G. We call this operation *translating* the coloring Q by α; this gives an action of the additive group of the field $\mathbb{F}_q$ on the set of q-colorings, and we will call an orbit of this action a *translation-set* of q-colorings of the map M.

It is clear that two q-colorings in the same translation-set give the same regular 1-form on C_G. Conversely, since the only relation among the 1-forms ω_f is that their sum is zero, it is obvious that if $\omega(Q') = \omega(Q)$, then Q' and Q are in the the same translation-set.

Now suppose that ω is any regular 1-form all of whose residue numbers are non-zero. I claim that $\omega = \omega(Q)$ for some q-coloring $Q \in \mathrm{Col}_q(M)$ of the map M. Indeed, write $\omega = \sum_f a_f \omega_f$ for some coefficients $a_f \in \mathbb{F}_q$; this is possible since the ω_f's span the space Ω_G of regular 1-forms on C_G. Define a function Q on the faces of M to $\mathbb{F}_q$ by setting $Q(f) = a_f$. Then by the analysis above, $\omega(Q) = \omega$ and Q is a proper q-coloring since all of the residue numbers of $\omega(Q)$ are non-zero.

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

$$Q \longmapsto \sum_f Q(f)\omega_f$$

gives an $\mathbb{F}_q$-linear map from $\mathbb{F}_q^F$ 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^1 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 4.7 *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 $\Omega_G^1 \neq \emptyset$).*

Indeed, we see that

$$\text{the number of } q\text{-colorings of } M \ = |\mathrm{Col}\,_q(M)| = q|\Omega^1_G| \tag{4.8}$$

since the fibers of the map from q-colorings of M onto Ω^1_G have cardinality q. Expressed geometrically in terms of the hyperplanes in H_∞, Proposition 3.6 gives that

$$|\Omega^1_G| = (q-1)|\mathcal{H}^1_M| \tag{4.9}$$

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

$$|\mathrm{Col}\,_q(M)| = q(q-1)|\mathcal{H}^1_M|. \tag{4.10}$$

Let us apply these results in the cases of low q, where the interest in 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 [Wi].

We begin with $q = 2$; no map can be 2-colored, so $\mathrm{Col}\,_2(M) = \emptyset$. On the other hand it is easy to see that Ω^1_G is also empty if $q = 2$. The only nonzero number in $\mathbb{F}_2$ is 1; therefore any 1-form in Ω^1_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 $\mathbb{F}_3 = \{0, \pm 1\}$; we see that any 1-form in Ω^1_G has only ± 1 as residue numbers. Suppose $\omega \in \Omega^1_G$. 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^1_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 4.11 *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 $\mathbb{F}_4 = \{0, a, b, c\}$. Any 1-form in Ω^1_G has only the nonzero numbers $\{a, b, c\}$ as residues. Since the field $\mathbb{F}_4$ 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 $\omega \in \Omega^1_G$. 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 4.12 *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.*

5 The chromatic function

As noted above, Proposition 3.6 and Theorem 4.7 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. In order to count these hyperplanes (those of $\mathcal{H}^1_M$) we must study the configuration of the set of nodes.

Let N be a set of nodes of C_G. We will first need to compute the dimension of the linear span $\mathcal{L}_N$ 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 $\mathcal{H}^1_M$ 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

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

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

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

This is actually, given the results of the previous section, 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

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

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

[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).

[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.

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