An Algebraic Reformulation of the Four Color Theorem

Howard Levi

March 13, 2005

Abstract

An algebraic equivalent of the four-color theorem is presented. The equivalent is the assertion of non-membership of a family of polynomials in a family of polynomial ideals over a particular finite field. A more precise statement must wait until specialized machinery has been introduced.

Geometry and algebra give to airy nothing a local habitation and a name.

adapted from Shakespeare.

1 About this paper

Most papers don't begin with a history of their preparation but, as you will see, it is necessary for this one. For the last several years of his life, Howard Levi attempted to devise a proof of the four-color theorem of a traditional kind—not involving the extensive use of computers. To this end he formulated an algebraic equivalent which he tried to prove directly. He discussed his attempts with Don Coppersmith, who repeatedly discovered gaps in the proof attempts, which Howard repeatedly attempted to fill, as far as we know without success. However, the consultation with Don helped Howard clarify his ideas, and the equivalence of the four-color theorem and the algebraic version was solidly established. When Howard died in late 2002, his work was left in an incomplete, and somewhat chaotic, condition. We, his colleagues and associates, Don Coppersmith, Melvin Fitting, and Paul Meyer, have tried to create a coherent presentation of the successful portion of Howard's work, from his notes and papers, in the hope that someone else will be able to successfully complete the project. All ideas and results here are due to Howard Levi; any infelicities of style are our responsibility.

We conclude this somewhat unusual preface with the acknowledgements that Howard originally intended to be at the conclusion of his paper.

"Although the author has never referred to the details of the computer based proof of the four color theorem [1] the existence of that proof has often served as a prefabricated light at the end of the tunnel in the search for the proofs in this paper. The author has used computers to find plausible conjectures and then often to shoot them down, employing the languages APL, SCRATCHPAD and MATHEMATICA

He wishes to thank his friends and colleagues for their kind assistance, including Melvin Fitting, Paul Meyer and Alfonso Vasquez , of CUNY, and William Burge, Alan Hoffman, Dan Prener and Barry Trager of IBM. He wishes especially to thank Don Coppersmith of IBM for his relentless detection of weak spots in the original manuscript."

2 Introduction

Our main result is a reformulation of the four color map theorem in a purely algebraic form (Theorem 10.1), with the four colors replaced by the four elements of the Galois field of order 4. After proving this equivalence, we have an algebraic statement that is true, because the four color theorem has been established, but which if given a purely algebraic proof would provide a computer-independent proof of the four color theorem.

According to a result of Hassler Whitney [3], the four color theorem for general planar graphs follows from a more restrictive statement: to prove the four-color theorem for planar graphs it is sufficient to prove it for planar graphs which have a Hamiltonian circuit, that is, a subset of its edges which link its set of vertices to form a simple closed curve. Suppose we have a planar graph G having a Hamiltonian circuit. Think of the circuit as dividing the graph into an "inside" and an "outside," and separate the graph into these two pieces, with the Hamiltonian circuit counted as belonging to each. For each of the two pieces, the original Hamiltonian circuit can be reshaped into a polygon, with the other "inside" or "outside" edges becoming internal edges. If each of these two pieces could be four-colored in ways that agree on the (common) bounding polygon, the separate four-colorings could be recombined to yield a four-coloring of the original graph G. Further, by adding extra internal edges to the "inside" and to the "outside" polygon, we can fully triangulate each of these two pieces. Clearly, if these triangulated polygons could be four colored in ways that agree on the bounding polygon, that would yield a four coloring of the original polygons without the additional edges, and so a four coloring of G.

So the strategy is this: suppose we have two fully triangulated polygons having the same boundary. If it could be shown that there always exist four colorings of such triangulated graphs that agree on the bounding polygon, the four color theorem for planar graphs with Hamiltonian circuits would follow, and hence the four color theorem generally. We will endow triangulated polygons with some algebraic structure, calling the resulting entities *triagons*. We will then give a necessary and sufficient polynomial test for the simultaneous four colorability of two triagons that share a polygon. This converts the four color theorem into a purely algebraic condition which, we hope, will eventually lead to a purely algebraic proof of the four color theorem. Our work is offered as an algebraic continuation of the deep, non-algebraic result of Whitney.

3 Algebraic background

We use as our colors the four elements $\{0, 1, \zeta, \zeta^2 = \zeta + 1\}$ of the Galois Field $F_4 = GF[2, 2]$. We note that if non-zero elements ζ^a, ζ^b, ζ^c of F_4 have sum 0 then $\langle \zeta^a, \zeta^b, \zeta^c \rangle$ must be a permutation of $\langle 1, \zeta, \zeta^2 \rangle$. In F_4 every element is its own additive inverse, so we routinely identify x and $-x$ when working in this field.

A four-coloring of graph nodes induces a three color labeling of graph edges (see Section 5). For three-colorings we use F_3, the field of three elements $\{0, 1, 2 = -1\}$. We generally identify the members of F_3 with the non-zero members of F_4 via the mapping that sends $x \in F_3$ to ζ^x in F_4.

We introduce the ring QR_n as the quotient ring of $F_3[u_1, \ldots, u_n]$ modulo the ideal I_3 generated by $\{u_i^3 - u_i \mid i = 1, 2, \ldots, n\}$. Note that for every element a of QR_n we have $a^3 = a$. This ring has the useful property that it contains a polynomial formula for every function from F_3^n to F_3. To see this let $\langle e_1, \ldots, e_n \rangle$ be an element of F_3^n and multiply together all the expressions $1 - (e_i - u_i)^2$ for $i = 1, \ldots, n$. This product is 1 on the given n-tuple and 0 on all other n-tuples. It is straightforward to use these characteristic functions to construct, for each function from F_3^n to F_3, a polynomial of QR_n which represents that function. In particular we assert that such a function returns the value

0 for every n-tuple $\langle e_1, \ldots, e_n \rangle$ of non-zero elements of F_3 if and only if its polynomial formula is in the ideal I_2 generated by $\{u_i^2 - 1 \mid i = 1, \ldots, n\}$. This follows from the facts that, modulo I_2, every element of QR_n is congruent to a polynomial all of whose monomials are square free, and an induction argument shows that a non trivial linear combination of such monomials cannot vanish on all n-tuples of non-zero elements of F_3. For every element a of QR_n we denote by $\mathsf{red2}(a)$ this square free residue of a modulo I_2.

We define a sequence of polynomials, POL_n, with coefficients in F_3, and with POL_n having variables $u_1, \ldots, u_n$, as follows.

$$\begin{aligned} \mathsf{POL}_0 &= 1 \\ \mathsf{POL}_1 &= u_1 \\ \mathsf{POL}_{n+1} &= u_{n+1} \cdot \mathsf{POL}_n - \mathsf{POL}_{n-1} \end{aligned} \tag{1}$$

We sometimes need to make use of algebraic objects constructed by applying the rules defining QR_n, I_2, I_3, $\mathsf{red2}(a)$, POL_n, but referring to a list of variables different from $\{u_1, u_2, \ldots\}$. If we introduce a new list of variables *var* then we designate the results of such constructions by $QR_n(var)$, $I_2(var)$, $I_3(var)$, $\mathsf{red2}(a, var)$, $\mathsf{POL}_n(var)$. As special, but natural, notation, we will use $\mathsf{POL}_n(s_1, \ldots, s_n)$ as an alternative to $\mathsf{POL}_n(\{s_1, \ldots, s_n\})$. All of our polynomials are understood to be in some $QR_n(var)$ or one of its quotient rings.

4 Geometric background

A_2 is the real affine plane in which a counterclockwise orientation has been designated. We use $\mathcal{P}_{n+2}$ to designate a polygon in A_2 of $n+2$ vertices with convex interior, no three of its vertices colinear, and with one of its vertices selected to be p_0. When we present a list $[p_0, p_1, \ldots, p_{n+1}]$ of vertices of $\mathcal{P}_{n+2}$ we assume that the list defines a counter-clockwise traversal of the polygon. We reserve the right to rename any of the vertices of $\mathcal{P}_{n+2}$ as p_0 provided that we also rename the other vertices to be in accord with this assumption. We call a graph which results from a triangulation of $\mathcal{P}_{n+2}$ a $\mathcal{TP}$ *graph*. More specifically a $\mathcal{TP}$ graph is a graph having the vertices and edges of $\mathcal{P}_{n+2}$ and whose other (internal) edges are a maximal subset of non-crossing segments which connect non-adjacent vertices of $\mathcal{P}_{n+2}$ (vertices of the polygon don't count as crossing points of internal edges). We will systematically use $\mathcal{T}_{n+2}$ to denote a $\mathcal{TP}$ graph whose polygon is $\mathcal{P}_{n+2}$.

A *triangle* of a $\mathcal{TP}$ graph is a triangle whose sides and vertices are edges and vertices of the $\mathcal{TP}$ graph. An *eye* of a $\mathcal{TP}$ graph is a vertex on no internal edge. A $\mathcal{TP}$ graph with $n+2$ vertices has n triangles, $n-1$ internal edges and at least two eyes. If a $\mathcal{TP}$ graph has more than one triangle, no two of its eyes can be adjacent. It is a useful fact that if p is an eye of a $\mathcal{TP}$ graph of n triangles, and pp_a, pp_b are the two edges of the polygon of this graph which contains p, then deleting p, pp_a, and pp_b yields a new $\mathcal{TP}$ graph of $n-1$ triangles.

5 Colorings

To four-color a graph means, in what follows, to assign an element of F_4 to each of the vertices of the graph so that the endpoints of each of its edges have different colors. Given an assignment of members of F_4 to the vertices of a graph, whether or not it is a four-coloring, we will associate with each edge the sum of the members of F_4 assigned to the endpoints of that edge—we call this an *edge label*. Endpoints of an edge have different colors if and only if the edge label assigned to that edge is non-zero, that is, it is ζ^h for some h of F_3. Thus a four-coloring induces a labeling of edges with non-zero members of F_4, and this is isomorphically a labeling of edges with members of F_3.

As an auxilliary notion, we will sometimes need to speak of a three-coloring of a graph, which is an assignment of elements of F_3 to *vertices* so that endpoints of each edge have different colors. For a three-coloring the *difference* of the colors of the endpoints of an edge must be a non-zero member of F_3, that is, 1 or -1.

6 Colorings of polygons

This section concerns itself with the coloring of *polygons*, not full $\mathcal{TP}$-graphs. Denote the vertices of $\mathcal{P}_{n+2}$ by $[p_0, p_1, \ldots, p_{n+1}]$. Let C_4 be a four-coloring of $\mathcal{P}_{n+2}$ for which $C_4(p_0) = 0$, $C_4(p_1) = 1$. This determines an edge labeling by non-zero members of F_4. We define ζ^{s_r} by

$$\zeta^{s_r} = \frac{C_4(p_{r+1}) + C_4(p_r)}{C_4(p_{r-1}) + C_4(p_r)} \qquad r = 1, \ldots, n+1 \tag{2}$$

where p_{n+2} means p_0. That is, given a node p_r, ζ^{s_r} is the ratio of the edge labels of the two edges of $\mathcal{P}_{n+2}$ having p_r as an endpoint, $p_r p_{r+1}$ and $p_{r-1}p_r$. It follows from (2) that

$$C_4(p_{r+1}) = 1 + \zeta^{s_1} + \ldots + \zeta^{s_1+\ldots+s_r} \qquad r = 1, \ldots, n+1 \tag{3}$$

For each ζ^{s_r}, we identify s_r with the corresponding member of F_3 and call it the *local sum* associated with vertex p_r. (The reason for the name will appear later.) Local sums are one of our main tools. We say that the $(n+1)$-tuple of F_3 elements $\langle s_1, \ldots, s_{n+1}\rangle$ is the *index* of the coloring C_4. We shall see that indices provide a way of connecting the colorings of two $\mathcal{TP}$ graphs which share their polygon.

We use our family of polynomials (1) to test whether a given $(n+1)$-tuple of F_3 elements $\langle s_1, \ldots, s_{n+1}\rangle$ is in fact the index of a coloring of a given polygon.

Theorem 6.1 *Let $\langle s_1, \ldots, s_{n+1}\rangle$ be an $(n+1)$-tuple of F_3 elements and let $\mathcal{P}_{n+2}$ be a polygon with vertices $[p_0, p_1, \ldots, p_{n+1}]$. Then there is a four coloring C_4 of the polygon $\mathcal{P}_{n+2}$ such that $C_4(p_0) = 0$, $C_4(p_1) = 1$, for which our $(n+1)$-tuple is the index, iff $\mathsf{POL}_{n+1}(s_1, \ldots, s_{n+1}) = 0$.*

Proof We note that the assignments of equations (3) for all values of $r < n+1$ can be made without conflict. It is only the requirement

$$C_4(p_{n+2}) = 1 + \zeta^{s_1} + \ldots + \zeta^{s_1+\ldots+s_{n+1}} = C_4(p_0) = 0$$

that is in doubt. Proceeding by induction we start with $n = 0$. We are to show that $1 + \zeta^{s_1} = 0$ iff $\mathsf{POL}_1(s_1) = 0$. This is clearly the case. We now examine the transition from n to $n+1$, and begin with the case $s_{n+1} = 0$. Because the values assigned by C_4 are from a field of characteristic 2, the last two terms of our sum cancel, and we conclude that the sum for $C_4(p_{n+2})$ is the same as the sum for $C_4(p_n)$, which from the induction hypothesis is 0 iff $\mathsf{POL}_{n-1}(s_1, \ldots, s_{n-1}) = 0$. From the recursion formula for POL_{n+1}, together with the assumption $s_{n+1} = 0$, we conclude that in this case the sum for $C_4(p_{n+2})$ is zero iff $\mathsf{POL}_{n+1}(s_1, \ldots, s_{n+1}) = 0$. There remains the cases in which s_{n+1} is not zero. It is a fact that, in F_4 for every x and non-zero y, we have $\zeta^x + \zeta^{x+y} = \zeta^{x-y}$ (check this by dividing both sides by ζ^x). Therefore the last two terms in our sum for $C_4(p_{n+2})$ combine to complete the sum for the n-tuple $\langle s_1, \ldots, s_{n-1}, s_n - s_{n+1}\rangle$. Our induction hypothesis implies that this sum is zero iff $\mathsf{POL}_n(s_1, \ldots, s_{n-1}, s_n - s_{n+1}) = 0$. The recursion formula for this case reads

$$\mathsf{POL}_n(s_1, \ldots, s_{n-1}, s_n - s_{n+1}) = (s_n - s_{n+1}) \cdot \mathsf{POL}_{n-1}(s_1, \ldots, s_{n-1}) - \mathsf{POL}_{n-2}(s_1, \ldots, s_{n-2})$$

which using (1) reduces to

$$\mathsf{POL}_n(s_1, \ldots, s_n) - s_{n+1} \cdot \mathsf{POL}_{n-1}(s_1, \ldots, s_{n-1}). \tag{4}$$

Since s_{n+1} is in F_3 and is not 0, $(s_{n+1})^2 = 1$, so the product of (4) with s_{n+1} reduces to $\mathsf{POL}_{n+1}(s_1, \ldots, s_{n+1})$, which is also, therefore, zero iff the sum for $C_4(p_{n+2})$ is 0.

It should also be verified that the assignment given by (3) is a four-coloring—no adjacent nodes share a color. But since each member of F_4 is its own additive inverse the expressions for $C_4(p_{r+1})$ and $C_4(p_{r+2})$ sum to $\zeta^{s_1+\ldots+s_{r+1}}$. This is not the 0 of F_4 and hence $C_4(p_{r+1})$ and $C_4(p_{r+2})$ are different. ■

7 Enumerating the triangles of a $\mathcal{TP}$ graph

We present a procedure for assigning a sequence of integers to each $\mathcal{TP}$ graph, which serves to identify the graph uniquely (up to isomorphism). We call this sequence the *identifier* of the $\mathcal{TP}$ graph. As part of the mechanism of defining this assignment, we produce an enumeration of the triangles of the $\mathcal{TP}$ graph, which plays a very significant role as well.

Let $\mathcal{T}_{n+2}$ be a $\mathcal{TP}$ graph resulting from a triangulation of a polygon $\mathcal{P}_{n+2}$. If $n = 1$ our graph has one triangle, which we name Δ_1. We use the sequence $\langle 1 \rangle$ as the identifier of this $\mathcal{TP}$ graph, and set $\mathsf{Identifier}(\mathcal{T}_3) = \langle 1 \rangle$.

Suppose now that n is bigger than 1 and that $\mathsf{Identifier}(\mathcal{T}_{n+1})$ has been defined. Let p_e be the eye of the graph $\mathcal{T}_{n+2}$ with largest e for which $1 \leq e \leq n$. Then the points p_{e-1}, p_e, p_{e+1} are vertices of a triangle of the graph. We select this triangle to be Δ_n. We delete p_e along with the edges and the triangle which contain it from $\mathcal{T}_{n+2}$. The resulting figure becomes a new $\mathcal{TP}$ graph $\mathcal{T}_{n+1}$ after we rename the points p_i as p_{i-1} for $e+1 \leq i \leq n+1$. This new graph has all the triangles of the original except Δ_n and no eye p_f with $e < f$. If $\mathsf{Identifier}(\mathcal{T}_{n+1}) = \langle e_1, \ldots, e_{n-1} \rangle$, we set $\mathsf{Identifier}(\mathcal{T}_{n+2}) = \langle e_1, \ldots, e_{n-1}, e \rangle$. In brief,

$$\mathsf{Identifier}(\mathcal{T}_3) = \langle 1 \rangle$$
$$\mathsf{Identifier}(\mathcal{T}_{n+2}) = \mathsf{Identifier}(\mathcal{T}_{n+1}) || e$$

where p_e is selected as described above, and $||$ is the operation of appending to the end of a sequence.

Note that $\mathsf{Identifier}(\mathcal{T}_{n+2})$ is an ordered sequence of positive integers $\langle e_1, \ldots, e_n \rangle$ with $e_1 = 1$, with $e_{i-1} \leq e_i$ and with $e_i \leq i$; $i = 2, \ldots, n$. Note also that every such ordered sequence is the identifier of a $\mathcal{TP}$ graph.

8 More about triangles of a $\mathcal{TP}$ graph

In this section, when three-coloring $\mathcal{TP}$ graphs we will assume the color assigned to p_0 is 0, and the color assigned to p_{n+1} is 1. This specification allows us to give the following.

Lemma 8.1 *Every $\mathcal{TP}$ graph has a unique three-coloring.*

Proof Let $\mathcal{T}_{n+2}$ be a $\mathcal{TP}$ graph. If $n = 1$, $\mathcal{T}_{n+2}$ has one triangle, which clearly has a unique three coloring. For larger n let p_e be one of the eyes of $\mathcal{T}_{n+2}$, with $0 < e < n+1$. If we delete the two edges of $\mathcal{T}_{n+2}$ which end at p_e the residual figure is a $\mathcal{TP}$ graph with one fewer vertex. We can extend the unique three-coloring of this $\mathcal{TP}$ graph to a unique coloring of $\mathcal{T}_{n+2}$ by assigning to p_e the unique color different from those of p_{e-1} and p_{e+1}. ■

Let $\mathcal{T}_{n+2}$ be a $\mathcal{TP}$ graph with vertices $[p_0, \ldots, p_{n+1}]$ and enumerated triangles $[\Delta_1, \ldots, \Delta_n]$. We consider the system of equations over F_3 in the variables $\{s_0, s_1, \ldots, s_{n+1}\}$ and $\{t_1, \ldots, t_n\}$ given by:

$$s_i = \sum{}' t_k, \qquad i = 0, 1, \ldots, n+1 \tag{5}$$

where each sum $\sum'$ is extended over those t_k for which the triangle Δ_k has p_i as one of its vertices. We first note that each triangle of $\mathcal{T}_{n+2}$ is represented in exactly three of the equations of (5) so that, because we work in F_3, we have $s_0 + \ldots + s_{n+1} = 0$. Also, let C_3 be the unique three-coloring of $\mathcal{T}_{n+2}$. Because each triangle is represented once in an equation for which $C_3(p_i) = 1$, once in an equation for which $C_3(p_i) = -1$, and once in an equation for which $C_3(p_1) = 0$, we have $C_3(p_0) \cdot s_0 + \ldots + C_3(p_{n+1}) \cdot s_{n+1} = 0$. Clearly these two equations relating the s_i can be solved for s_0 and s_{n+1} in terms of $\{s_1, \ldots, s_n\}$. Guided by these results we form system (6) by dropping the equations for s_0 and s_{n+1} from (5).

$$s_i = \sum{}' t_k, \qquad i = 1, \ldots, n \tag{6}$$

We now show that system (6) can be solved for $\{t_1, \ldots, t_n\}$ in terms of $\{s_1, \ldots, s_n\}$. To describe this solution we note that $p_0, \ldots, p_{n+1}$ are linearly ordered by the counter-clockwise traversal of the polygon of $\mathcal{T}_{n+2}$. For $0 \le a < b \le n+1$ we denote the set of vertices which follow vertex p_a and precede vertex p_b by $[a, b]$ (this excludes the end points).

Theorem 8.2 *For each Δ_k of $\mathcal{T}_{n+2}$ let its list of vertices be $[p_a, p_b, p_c]$ where $a < b < c$. Also let C_3 be the three-coloring from above, determined by setting $C_3(p_0) = 0$ and $C_3(p_{n+1}) = 1$. Then a solution of (6) is given by*

$$t_k = \sum{}' s_i - \sum{}'' s_i - \sum{}''' s_i \qquad k = 1, \ldots, n$$

where

$\sum'$ *is extended over the i for which p_i is in $[a, c]$ and $C_3(p_i) = C_3(p_b)$*
$\sum''$ *is extended over the i for which p_i is in $[a, b]$ and $C_3(p_i) = C_3(p_c)$*
$\sum'''$ *is extended over the i for which p_i is in $[b, c]$ and $C_3(p_i) = C_3(p_a)$*

Proof Let us say a summation $\sum s_i$ *references* t_j if there is a term s_m in the summation such that the formula for s_m from system (6) has t_j on its right side. Note that in the proposed formula for t_k above, only $\sum'$ references t_k, while if $j \ne k$ then $\sum'$ references t_j if and only if exactly one of $\sum''$ and $\sum'''$ references t_j. It follows that when all the components of $\sum' s_i - \sum'' s_i - \sum''' s_i$ are collected the only survivor is t_k. ■

Note 1: The expression for t_k in Theorem 8.2 depends only on the portion of $\mathcal{T}_{n+2}$ which lies in the open half-plane bounded by the line $p_a p_c$ and which contains p_b. In particular we infer that neither s_0 nor s_{n+1} occurs in such an expression because there is no such open half-plane which contains p_0 or p_{n+1}.

Note 2: Observe that our construction of the t_k as expressions in the s_i can be imitated in obtaining expressions for t_k in the s_i of that open half plane of $p_c p_a$ which contains p_b, and also for t_k in the s_i of the open half plane of $p_b p_c$ which contains p_a. Thus there are, in general, three distinct forms in the s_i for each t_k, each of which evaluates to t_k. More explicitly, in the Theorem we took the triangle to be $[p_a, p_b, p_c]$; the construction works if we take it to be $[p_b, p_c, p_a]$ or $[p_c, p_a, p_b]$, and these give different formal expressions for the t_k. It should be mentioned that these alternate representations do involve s_0 and s_{n+1}, in contrast to the original representation, which did not.

Definition 8.3 *Let v be the identifier of the $\mathcal{TP}$ graph $\mathcal{T}_{n+2}$, as defined in Section 7. The graph (or equally well, the identifier) determines the set of equations (5), and hence an $n+2$ by n matrix over F_3 mapping the vector $\langle t_1, \ldots, t_n \rangle$ to the vector $\langle s_0, s_1, \ldots, s_{n+1} \rangle$. We call this matrix* $\mathsf{matriag}[v]$*, so that $\langle s_0, s_1, \ldots, s_{n+1} \rangle = \mathsf{matriag}[v] \cdot \langle t_1, \ldots, t_n \rangle$.*

In a similar way, $\mathsf{sqmat}[v]$ *is the n by n matrix over F_3 mapping $\langle t_1, \ldots, t_n \rangle$ to $\langle s_1, \ldots, s_n \rangle$ according to (6), so $\langle s_1, \ldots, s_n \rangle = \mathsf{sqmat}[v] \cdot \langle t_1, \ldots, t_n \rangle$. By Theorem 8.2 this matrix is invertible, and $\mathsf{sqmat}[v]^{-1} \cdot \langle s_1, \ldots, s_n \rangle = \langle t_1, \ldots, t_n \rangle$.*

We now introduce the final major component of our machinery—*labels for triangles*, or just *labels*. Let C_4 be a four coloring of $\mathcal{T}_{n+2}$. For each Δ_k with vertices $[p_a, p_b, p_c]$, the edge labels assigned to the edges of this triangle are $C_4(p_a) + C_4(p_b)$, $C_4(p_b) + C_4(p_c)$, $C_4(p_c) + C_4(p_a)$, and we know that each of these is non-zero and their sum is zero. This sequence of edge labels must therefore constitute a permutation of $\langle 1, \zeta, \zeta^2 \rangle$. If this permutation is even we say that the *label induced on Δ_k by C_4* is 1, if it is odd we say that this label is -1. We generally denote the label of Δ_k by t_k. Notice that the assignment of a label does not change if we renumber the vertices of the $\mathcal{TP}$ graph starting from a different point than p_0 because it only depends on the orientation that the list $[p_a, p_b, p_c]$ of vertices confers on its triangle, and our agreement about numbering implies that renumbering preserves this orientation.

If we require only that C_4 color the polygon of our $\mathcal{TP}$ graph then a sum attached to an edge of a triangle could be zero, and, for such a triangle, the definition above could not be applied. In this case we assign 0 to Δ_k as its induced label. Thus every four coloring of the polygon $\mathcal{P}_{n+2}$ of a $\mathcal{TP}$ graph $\mathcal{T}_{n+2}$ yields as its labeling an n-tuple $\langle t_1, \ldots, t_n \rangle$ of elements of F_3, all of which are different from zero iff the given coloring actually four-colors the whole $\mathcal{TP}$ graph.

In Figure 1 a portion of a $\mathcal{TP}$-graph is shown, including nodes p_{i-1}, p_i, and p_{i+1}. It is assumed the graph has been four-colored, and edge labels and triangle labels are indicated. Because we have a four-coloring, all edge labels are non-zero. We have shown the edge label of $p_{i-1}p_i$ as ζ^h. If the edge labels associated with the edges of the leftmost triangle, taken with clockwise orientation, are an even permutation of $\langle 1, \zeta, \zeta^2 \rangle$, the label on this triangle will be $t_a = +1$, and the label on the edge out of p_i immediately clockwise to $p_{i-1}p_i$ must be ζ^{h+1}, that is, ζ^{h+t_a}. Similarly if the permutation is an odd one, and $t_a = -1$; either way the next edge label out of p_i must be ζ^{h+t_a}. This reasoning continues across all the triangles, and it is straightforward to check that all the edge labels must be as indicated.

In Section 6 local sums and indexes were defined. Applying that definition to Figure 1, local sum s_i was characterized by the condition that ζ^{s_i} must be the ratio of the edge label on p_ip_{i+1} to the edge label on $p_{i-1}p_i$. This ratio is seen to be $\zeta^{t_a + \ldots + t_q}$, and so $s_i = t_a + \ldots + t_q$ (accounting for the terminology *local sum*). In other words, local sums and labels are connected by the equations of (5). We summarize some of this discussion as a theorem.

Theorem 8.4 *Let $\langle t_1, \ldots, t_n \rangle$ be the labels induced on triangles $\langle \Delta_1, \ldots, \Delta_n \rangle$ of $\mathcal{T}_{n+2}$ by a four coloring C_4 of $\mathcal{T}_{n+2}$, and let $\langle s_1, \ldots, s_{n+1} \rangle$ be computed from these labels using the equations (5). Then $\langle s_1, \ldots, s_{n+1} \rangle$ is the index of C_4. Conversely if $\langle t_1, \ldots, t_n \rangle$ is any n-tuple of non-zero elements of F_3 then there is a coloring C_4 of $\mathcal{T}_{n+2}$ which induces $\langle t_1, \ldots, t_n \rangle$ as the labels on $\langle \Delta_1, \ldots, \Delta_n \rangle$ and having $\langle s_1, \ldots, s_{n+1} \rangle$ is its index, where the s_i and t_j are related by (5).*

If $\langle t_1, \ldots, t_n \rangle$ is any tuple of non-zero elements of F_3, and we use (5) to compute $\langle s_1, \ldots, s_{n+1} \rangle$, we will get an index of a four-coloring of a polygon $\mathcal{P}_{n+2}$, so by Theorem 6.1, $\mathsf{POL}_{n+1}(s_1, \ldots, s_{n+1})$ must be 0. It follows from the remarks in Section 3 that this polynomial, when expressed in terms of t_i as variables, must be in the ideal generated by $\{t_i^2 - 1 \mid i = 1, \ldots, n\}$, denoted $I_2(\{t_1, \ldots, t_n\})$. This gives us the following.

Figure 1: How local sums arise

Corollary 8.5 $\mathsf{red2}(\mathsf{POL}_{n+1}(s_1, \ldots, s_{n+1}), \{t_1, \ldots, t_n\}) = 0$.

With n triangles, there are 2^n possible labelings, hence we have the following.

Corollary 8.6 *$\mathcal{T}_{n+2}$ has 2^n four colorings.*

We now introduce our fundamental structure, the *triagon*. A triagon is a triangulated polygon with triangle labels and local sums assigned to it, connected by (5).

We have two ways of making a triagon from a triangulated polygon whose vertices are enumerated as $\{p_0, \ldots, p_{n+1}\}$ and whose triangles are enumerated as $\{\Delta_1, \ldots, \Delta_n\}$.

I(a) Assign labels $\langle t_1, \ldots, t_n \rangle$.

I(b) Compute local sums $\langle s_1, \ldots, s_n \rangle$ by $\mathsf{sqmat}[v] \cdot \{t_1, \ldots, t_n\}$.

or

II(a) Begin with local sums $\langle s_1, \ldots, s_n \rangle$ (which we intend to be copies of those of another triagon).

II(b) Compute labels $\langle t_1, \ldots, t_n \rangle$ by $\mathsf{sqmat}[v]^{-1} \cdot \{s_1, \ldots, s_n\}$.

9 Further Facts About POL_n

We present results that continue what began with Corollary 8.5.

Theorem 9.1 $\mathsf{POL}_n(s_1, \ldots, s_n) = \mathsf{POL}_n(s_n, \ldots, s_1)$.

Proof The following direct proof was provided by both Cormac O'Sullivan and Jonas Reitz (private communications). Define $\widehat{\mathsf{POL}}_n(u_1, \ldots, u_{n+1})$ to be $\mathsf{POL}_n(u_{n+1}, \ldots, u_1)$. Then show by induction on n and k that, if $1 \leq k \leq n$, $\mathsf{POL}_{n+1} = \widehat{\mathsf{POL}}_k \mathsf{POL}_{n+1-k} - \widehat{\mathsf{POL}}_{k-1} \mathsf{POL}_{n-k}$. The special case where $k = n$ almost immediately yields that $\mathsf{POL}_{n+1} = \widehat{\mathsf{POL}}_{n+1}$. ∎

The original proof of Theorem 9.1 was more indirect, but provided a non-recursive characterization of $\mathsf{POL}_n(s_1,\ldots,s_n)$. It is the sum of all monomials $a\cdot u_{i_1}\cdot u_{i_2}\cdot\ldots\cdot u_{i_k}$ whose subscripts are non-negative integers which obey the rules $1\le i_1<i_2\ldots<i_r<\ldots<i_k\le n$, each i_r is congruent to r mod (2) and whose greatest subscript i_k is congruent to n mod (2). The coefficient a of a monomial of POL_n is computed from the formula $a=(-1)^s$ where $s=(n-k)/2$. We take into account the case of a possible non-zero constant term in POL_n. It arises from $k=0$ and then only if n is even. Our formula for a gives the value for the constant term in this case. These rules can be deduced by a straightforward induction argument showing that the monomials of the left side of

$$\mathsf{POL}_{n+1}(u_1,\ldots,u_{n+1})=u_{n+1}\cdot\mathsf{POL}_n(u_1,\ldots,u_n)-\mathsf{POL}_{n-1}(u_1,\ldots,u_{n-1})$$

obey the proposed rules if those of the right side do so. We omit the details. We note that this description of the monomials of POL_n implies Theorem 9.1.

Theorem 9.1 justifies the use of the formula

$$\mathsf{POL}_{n+1}(s_1,\ldots,s_{n+1})=s_1\cdot\mathsf{POL}_n(s_2,\ldots,s_{n+1})-\mathsf{POL}_{n-1}(s_3,\ldots,s_{n+1})$$

in our next proof. We now examine the output when POL_n is evaluated at certain sets of local sums. It supplies the bridge from n to $n+1$ in the proof of our main theorem by mathematical induction.

Lemma 9.2 *Suppose we have a triagon, with e being the greatest integer for which p_e is an eye and $1\le e\le n$, that is, p_{e-1}, p_e, p_{e+1} are the vertices of Δ_n. We represent the local sums of the triagon at the points p_{e-1}, p_e, p_{e+1} as $s'_{e-1}+t_n$, t_n, $s'_{e+1}+t_n$. We assert the following.*

$$\begin{aligned}&\mathsf{POL}_n(s_1,\ldots,s'_{e-1}+t_n,t_n,s'_{e+1}+t_n,\ldots,s_n)=\\&\quad t_n\mathsf{POL}_{n-1}(s_1,\ldots,s'_{e-1},s'_{e+1},\ldots,s_n)-\\&\quad(1-t_n^2)\cdot\mathsf{POL}_{n-2}\{s_1,\ldots,s'_{e-1}+s'_{e+1},\ldots,s_n\}\end{aligned}$$

and consequently $\mathsf{POL}_n(s_1,\ldots,s'_{e-1}+t_n,t_n,s'_{e+1}+t_n,\ldots,s_n)$ is congruent to $t_n\cdot\mathsf{POL}_{n-1}(s_1,\ldots,s'_{e-1},s'_{e+1},\ldots,s_n)$ in $F_3[s_1,\ldots,s'_{e-1},t_n,s'_{e+1},\ldots,s_n]/(t_n^2-1)$.

Proof To prove our theorem we require the formulas $\mathsf{POL}_0=1$, $\mathsf{POL}_1(s_1)=s_1$, $\mathsf{POL}_2(s_1,s_2)=s_1\cdot s_2-1$, $\mathsf{POL}_3(s_1,s_2,s_3)=s_1\cdot s_2\cdot s_3-s_1-s_3$. Then, by direct substitution, we can confirm our assertion for the cases of local sums of triagons whose identifiers are

$$\langle 1,1\rangle,\langle 1,2\rangle,\langle 1,1,1\rangle,\langle 1,1,2\rangle,\langle 1,1,3\rangle,\langle 1,2,2\rangle \text{ and } \langle 1,2,3\rangle.$$

We extend these results to confirm it for all the remaining triagons, the triagons with local sums $[\ldots,s'_{e-1}+t_n,t_n,s'_{e+1}+t_n,\ldots]$, by using our recursion formulas for POL_n to refer back to these cases. ∎

Theorem 9.3 *Let $\langle t_1,\ldots,t_n\rangle$ be labels for the triangles of triagon $\mathcal{T}_{n+2}$ and let $\langle s_0,s_1,\ldots,s_{n+1}\rangle$ be its local sums at its vertices $[p_0,p_1,\ldots,p_{n+1}]$. Then over QR_n*

1. $\mathsf{red2}(\mathsf{POL}_n(s_1,\ldots,s_n),\{t_1,\ldots,t_n\})=t_1\cdot t_2\cdot\ldots\cdot t_n$
2. $\mathsf{red2}(\mathsf{POL}_{n+1}(s_1,\ldots,s_{n+1}),\{t_1,\ldots,t_n\})=0$
3. $\mathsf{red2}(\mathsf{POL}_{n+2}(s_0,\ldots,s_{n+1}),\{t_1,\ldots,t_n\})=-t_1\cdot t_2\cdot\ldots\cdot t_n$

Proof We use mathematical induction to prove our theorem. For $n = 1$ we see that the set of local sums for $\mathcal{T}_3$ at $\langle p_0, p_1, p_2\rangle$ is $\langle t_1, t_1, t_1\rangle$. We infer from $\mathsf{POL}_0 = 1$, from $\mathsf{POL}_1(t_1) = t_1$, from $\mathsf{POL}_2(t_1, t_1) = t_1^2 - 1$ and from $\mathsf{POL}_3(t_1, t_1, t_1) = t_1^3 - t_1 - t_1$ (which in QR_n is $-t_1$), that the assertions hold for the case $n = 1$.

Notice that Lemma 9.2 can be applied to the local sums $\{s_0, s_1, \ldots, s'_{e-1}, s'_{e+1}, \ldots, s_{n+1}\}$ of the triagon $\mathcal{T}_{n+1}$ which results from deleting the vertex p_e and the two edges which contain it from our original triagon $\mathcal{T}_{n+2}$. Assuming our theorem for $n - 1$ we deduce our theorem directly from Lemma 9.2. ∎

Observing that $f(\mathsf{expr}) = \mathsf{red2}(\mathsf{expr}, \{t_1, \ldots, t_n\})$ denotes the natural homomorphism from $F_3[t_1, \ldots, t_n]$ to $F_3[t_1, \ldots, t_n]/I_2\{t_1, \ldots, t_n\}$, we deduce our next result.

Corollary 9.4 *In QR_n,*

$$\mathsf{red2}(\mathsf{POL}_n(s_1, \ldots, s_n) \cdot \mathsf{POL}_{n+2}(s_0, \ldots, s_{n+1}), \{t_1, \ldots, t_n\}) = -1$$

10 An Algebraic Equivalent

Let u and v be a pair of triagons which triangulate the same polygon. Suppose we had a four-coloring C_4 of the two triagons that agreed on the polygon. Such a coloring would induce a labeling of the triangles of each of our triagons and the index of C_4 could be computed independently from each of these labelings. Let $\langle u_1, \ldots, u_n\rangle$ and $\langle v_1, \ldots, v_n\rangle$ be the induced labels. Then we can rewrite (6) for the two triagons as $\langle s_1, \ldots, s_n\rangle = \mathsf{sqmat}[u] \cdot \langle u_1, \ldots, u_n\rangle$ and $\langle s_1, \ldots, s_n\rangle = \mathsf{sqmat}[v] \cdot \langle v_1, \ldots, v_n\rangle$. We derive that $\langle v_1, \ldots, v_n\rangle = \mathsf{sqmat}[v]^{-1} \cdot \mathsf{sqmat}[u] \cdot \langle u_1, \ldots, u_n\rangle$. We shall use $\mathsf{matmat}[v, u]$ to stand for $\mathsf{sqmat}[v]^{-1} \cdot \mathsf{sqmat}[u]$, and thus we conclude

$$\langle v_1, \ldots, v_n\rangle = \mathsf{matmat}[v, u] \cdot \langle u_1, \ldots, u_n\rangle. \tag{7}$$

What we need is a condition that guarantees we can reverse the above steps.

In the ring QR_n we see that (7) expresses each v_i as a form in the u_i. We denote the product of these forms by $\mathsf{pr}[v, u]$. Consider the function mapping $\langle e_1, \ldots, e_n\rangle \in F_3^n$ to F_3 given by: evaluate $\mathsf{pr}[v, u]$ after assigning each variable u_i the value e_i. From the discussion in Section 3, this function is not identically 0 on arguments for which no e_i is 0 iff $\mathsf{red2}(\mathsf{pr}[v, u], \{u_1, \ldots, u_n\}) \neq 0$. Thus $\mathsf{red2}(\mathsf{pr}[v, u], \{u_1, \ldots, u_n\}) \neq 0$ is a necessary and sufficient condition to ensure that (7) has a solution in F_3 in which no u_i or v_i is 0.

Now assume that $\mathsf{red2}(\mathsf{pr}[v, u], \{u_1, \ldots, u_n\}) \neq 0$. Given a solution of (7) with no u_i or v_i being 0, the labels $\langle u_1, \ldots, u_n\rangle$ and $\langle v_1, \ldots, v_n\rangle$ induce four-colorings of the two triagons u and v. What we now show is that the four-colorings fit together, agreeing on the bounding polygon the triagons share. Part of this is easy. Since (7) is satisfied, if we compute local sums for the two triagons using $\langle s_1, \ldots, s_n\rangle = \mathsf{sqmat}[u] \cdot \langle u_1, \ldots, u_n\rangle$ and using $\langle s_1, \ldots, s_n\rangle = \mathsf{sqmat}[v] \cdot \langle v_1, \ldots, v_n\rangle$, we will get the same n-tuples. The remaining issue is to show that the two triagons share their values for s_0 and s_{n+1} as well, and for this we have the following argument. First, for each of the triagons, $\mathsf{POL}_{n+1}(s_1, \ldots, s_{n+1}) = 0$ by Theorem 6.1. Using the recurrence relation for POL_{n+1}, we then have the following, for each triagon.

$$s_{n+1} \cdot \mathsf{POL}_n(s_1, \ldots, s_n) - \mathsf{POL}_{n-1}(s_1, \ldots, s_{n-1}) = 0$$

It follows from part 1 of Theorem 9.3 that $\mathsf{POL}_n(s_1, \ldots, s_n) \neq 0$, which shows that s_{n+1} depends only on $s_1, \ldots, s_n$, which both triagons share, and thus both triagons have a common value for

s_{n+1}. In a similar way $\mathsf{POL}_{n+1}(s_0, \ldots, s_n) = 0$, so using Theorem 9.1, $\mathsf{POL}_{n+1}(s_n, \ldots, s_0) = 0$. Then by the recurrence relation for POL_{n+1} again,

$$s_0 \cdot \mathsf{POL}_n(s_n, \ldots, s_1) - \mathsf{POL}_{n-1}(s_n, \ldots, s_2) = 0.$$

Since $\mathsf{POL}_n(s_n, \ldots, s_1) = \mathsf{POL}_n(s_1, \ldots, s_n) \neq 0$, we conclude that s_0 also depends only on $s_1, \ldots, s_n$, which both triagons share, as noted above. Thus the two triagons share their values for both s_0 and s_{n+1}, and the four-colorings agree. We conclude that $\mathsf{red2}(\mathsf{pr}[v, u], \{u_1, \ldots, u_n\}) \neq 0$ is a valid test for the simulataneous four colorability of triagons v and u.

Theorem 10.1 (Algebraic Test) *Let u and v be a pair of triagons which triangulate the same polygon; the two triagons are simultaneously four-colorable if and only if* $\mathsf{red2}(\mathsf{pr}[v, u], \{u_1, \ldots, u_n\}) \neq 0$*. Equivalently stated, the two triagons are simultaneously four-colorable if and only if the polynomial $(v_1 \cdot v_2 \cdot \ldots \cdot v_n)$ is not in the ideal $I_2(\{u_1, \ldots, u_n\})$, where this is the polynomial in the variables u_1, u_2, ..., u_n in which the expressions v_i and the variables u_i are connected by (7).*

References

[1] Appel, K. and Haken, W. Every planar map is four colorable. *Bull. Amer. Math. Soc. 82* (1976), 711–712.

[2] Cox, D., Little, J., and O'Shea, D. *Ideals, Varieties, and Algorithms.* Springer-Verlag, New York, 1992.

[3] Whitney, H. A theorem on graphs. *Ann. Math. 32* (1931), 378–390.