Some NP Complete Problems Based on Algebra and Algebraic Geometry

Paul Hriljac
Department of Mathematics, ERAU

Notation: Let p be a prime and let $\mathbb{F}$ be a finite field with $q = p^n \geq 4$ elements and algebraic closure $\bar{\mathbb{F}}$. For sets X, Y let $Hom(X, Y)$ denote the functions from X to Y. For numbers m, n let $Hom(m, n)$ denote $Hom(\{1, \dots m\}, \{1, \dots, n\})$. Let Γ be a connected graph with vertex set V and edge set E.

Definition a $\mathbb{F}$-*coloring* of Γ is an assignment $V \to \mathbb{F}^*$ so that no adjacent vertices share the same value.

From now on, suppose that each edge of Γ has be given an orientation, so one can speak of the initial vertex and terminal vertex of each edge.
Assign to Γ the array $A_\Gamma \in \mathbb{F}^{|V||E|(q-1)}$ given by

$$A_\Gamma(v, e, c) = \begin{cases} 1 \text{ if } v = e(0) \\ c \text{ if } v = e(1) \\ 0 \text{ else.} \end{cases}$$

Here v is vertex, e is an edge and c is an element of $\mathbb{F} - \{0\}$. Let $\mathbf{D}_{V,E}$ be the determintal variety corresponding to $|V|$ by $|E|$ matrices with rank $< min(|E|, |V|)$ (see [3]). We suppress the indices when the context is clear. Then $\mathbf{D}$ is a subscheme of $\mathbb{A}^{|V||E|}$. For $\alpha \in Hom(E, \mathbb{F} - \{0, -1\})$ let π_α be the projection $\mathbb{A}^{|V||E|(q-1)} \to \mathbb{A}^{|V||E|}$ given by $A \mapsto [A(v, e, \alpha(e))]_{v,e}$. Let $\mathbf{D}_\alpha = \pi_\alpha^* \mathbf{D}$ and let $\mathbf{C} = \mathbf{C}_{|V||E|(q-1)} = \cup_\alpha \mathbf{D}_\alpha$. Call $\mathbf{C}$ the *coloring variety* associated to graphs with $|E|$ edges, $|V|$ vertices and $q-1$ colors. Since $\mathbf{C}$ can be defined by cofactor relations, it can be realized as a closed subscheme of $\mathbb{A}^{|V||E|(q-1)}$.

Theorem: *Assuming* $|V| \leq |E|$ Γ *is* $\mathbb{F}$-*colorable if and only if* $A_\Gamma \in \mathbf{C}$.

Proof: Given an $\mathbb{F}$-*coloring* $\{c_v\}$ of Γ, Fermat's Theorem implies $c_v^{q-1} = 1$ for all vertices v. Therefore, if v and w are adjacent

$$\frac{c_v^{q-1} - c_w^{q-1}}{c_v - c_w} = 0.$$

Expanding this out gives

$$\prod_{\alpha \in \mathbb{F}-\{0,-1\}} c_{e(0)} + \alpha c_{e(1)} = 0.$$

Therefore $c_{e(0)} + \alpha c_{e(1)} = 0$ for some $\alpha = \alpha(e) \in \mathbb{F} - \{0, -1\}$, which implies $\mathbf{c}M_\Gamma = 0$ where

$$M_\Gamma = \pi_\alpha(A_\Gamma) = \begin{cases} 1 \text{ if } v = e(0) \\ \alpha(e) \text{ if } v = e(1) \ , \\ 0 \text{ else} \end{cases} \qquad \mathbf{c} = [c_v]_{v \in V}.$$

This shows that the rank of $\pi_\alpha(A_\Gamma)$ is less than $|V|$ and hence $A_\Gamma \in \mathbf{C}$.
On the other hand if $\pi_\alpha(A_\Gamma)$ has rank less than $|V|$ for some $\alpha \in Hom(E, \mathbb{F} - \{0, -1\})$ than a nonzero vector in the kernel space $\pi_\alpha(A_\Gamma)$ is easily seen to provide a procedure yielding a valid coloring for the graph.

From this proof one sees the

Theorem: *If* $|E| < |V|$ *then* Γ *is* $\mathbb{F}$*-colorable.*

Proof: Since $|E| < |V|$ the kernel space of the linear map corresponding to $\mathbf{x} \mapsto \mathbf{x}\pi_\alpha(A_\Gamma)$ is nonzero for any α. Using a nonzero element of the kernel space one can create a coloring.

Example 1: Try to 3-color the following graph:

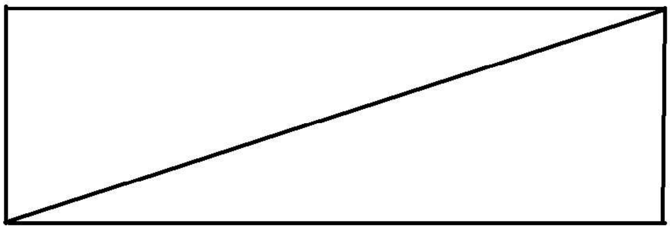

The theorem implies that the graph has a 3-coloring iff there is a map $\alpha : \{1, 2, 3, 4, 5\} \rightarrow$

$\mathbb{F}_4 - \{0, -1\}$ such that the rank of the matrix

$$\pi_\alpha(A_\Gamma) = \begin{bmatrix} \alpha_1 & 1 & 0 & 0 & 0 \\ 0 & \alpha_2 & 1 & 0 & \alpha_5 \\ 1 & 0 & 0 & \alpha_4 & 1 \\ 0 & 0 & \alpha_3 & 1 & 0 \end{bmatrix}.$$

is less than 4. Since the determinant of the matrix with the last column deleted is 1, the rank is never less than 4 so the graph cannot be 3-colored.

Definition: Given a family $\mathfrak{X} = (X, Y)$ of varieties X and subvarieities $Y \subset X$ over a field K *the membership problem for* $\mathfrak{X}$ to be the following: Given $P \in X(K)$, is $P \in Y(K)$?

Definition: Let $\mathfrak{C} = (\mathbf{C}_{m,n,q-1}, \mathbb{A}^{mn(q-1)})$ using coloring varities as defined above.

Theorem: *The membership problem for* $\mathfrak{C}$ *over a field* $\mathbb{F}$ *with* $q \geq 4$ *elements is NP Complete.*

Proof: To see that this problem is in NP, one just needs to check the rank of $\pi_\alpha(A_\Gamma)$ for a given α. To see that this problem is NP hard note that the NP complete problem of $\mathbb{F}$-coloring graphs is subsumed by this problem.

We now introduce another NP complete membership problem with the following example:

Example 2: Consider the graph in Example 1. Parametrize a line in $\mathbb{A}^{4,5,2}$ by

$$\gamma(t) = \begin{bmatrix} \alpha_1 & t & 0 & 0 & 1-t \\ 0 & \alpha_2 t & 1 & 0 & \alpha_5 \\ t & 0 & 0 & \alpha_4 & t \\ 0 & 0 & \alpha_3 & t & 0 \end{bmatrix}.$$

then $\gamma(1) = A_\Gamma$. Each choice of α yields a 4 by 5 array with coefficients in $\mathbb{F}[t]$. Taking the five major cofactors of this array gives five polynomials $f_{1,\alpha}, \ldots, f_{5,\alpha}$. Taking the greatest common divisor of these polynomials gives a polynomial f_α of degree ≤ 4. It is easy to see

$$f_\alpha(1) = 0 \text{ if and only if } \alpha \text{ gives a valid coloring of the graph.}$$

Next suppose $\mathcal{F} = \{f(x)_\alpha\}_{\alpha \in Hom(n,l)}$ is a set of polynomials in $\mathbb{F}[x]$ with degree $\leq n$ whose coefficients can be found from α in time polynomial in l, m, n. Let $Z(\mathcal{F})$ denote the roots of $\prod_\alpha f_\alpha$ and $\mathfrak{Z} = (Z(\mathcal{F}), \mathbb{A}^1)$.

Theorem: *The membership problem for $\mathfrak{Z}$ over a field $\mathbb{F}$ with $q \geq 4$ elements is NP Complete.*

Proof: Using elementary linear algebra it is easy to check that the variety $D_{n,m}$ is cut out set-theoretically by the equations $\delta_0 = \cdots = \delta_{m-n} = 0$ where δ_i is the determinant of the $n \times n$ matrix consisting of the first $n-1$ columns of a matrix as well as the column in position $n+1$. To decide if some $A \in \mathbb{A}^{m,n,l}$ is in $\mathbf{C}_{m,n,l}$ let L be a line in $\mathbb{A}^{m,n,l}$ parametrized by some linear function $\theta(t)$ with $\theta(t_0) = A$. Let $f_\alpha(t) = gcd[\delta_0(\pi_\alpha(\theta(t))), \ldots, \delta_{m-n}(\pi_\alpha(\theta(t)))]$. Then $A \in \mathbf{D}_\alpha$ iff $f_\alpha(t_0) = 0$ and so $A \in \mathbf{C}$ iff t_0 is a root of $\prod_\alpha f_\alpha$. This shows that the $\mathfrak{Z}$ membership problem is NP-hard. It is easy to see that the problem is in NP: Given α and t_0, just check that $f_\alpha(t_0) = 0$.

Definition: A *computation tree* over a finite field $\mathbb{F}$ is a rooted tree binary with a collection of field elements associated with each node (see [1]). Other than roots and leaves, each node is either a computation or a branch. Each computation node performs an algebraic calculation over that field that depends only on the field elements associated with earlier nodes. The branch nodes have a single associated field element and the two edges emanating from it are labeled " $=$ " and " $\neq$ ". A *decision tree* T for $X \subset \mathbb{F}$ is a computation tree that decides if some $x \in \mathbb{F}$ belongs in X. The leaves are labeled " $\in$ " and " $\notin$ ". Since the tree is operating only on elements of $\mathbb{F}$ we may assume that the root node has a unique field element. Let $|T|$ denote the length of the longest path through T.

Definition: A *computation sequence* for a polynomial $f(x) \in \mathbb{F}[x]$ is a sequence of polynomials $g_i(x) \in \mathbb{F}[x], i = 0, \ldots, r$ such that:

$$g_0 = x;$$

$$g_i = (\sum_{j=0}^{i-1} a_{i,j} g_j) * (\sum_{j=0}^{i-1} b_{i,j} g_j), \quad a_{i,j}, b_{i,j} \in \mathbb{F};$$

$$f(x) = f(0) + \sum_{i=0}^{r} c_i * g_i, \quad c_i \in \mathbb{F};$$

Let $\text{Ł}(f)$ denote the length of the shortest computation sequence for f. This is the *multiplicative complexity* of a polynomial $f(x) \in \mathbb{F}[x]$.

Suppose $\mathbb{F}$ is a finite field with q elements and $X \subset \mathbb{F}$ a subset with characteristic function χ. Then χ can be realized as a polynomial $\chi(x) \in \mathbb{F}[x]$. Suppose T is a decision tree for $X \subset \mathbb{F}$. Then

Theorem: $L(\chi) \leq 24 * log(q) * |T|$.

To prove the theorem we must introduce some new ideas and discuss their properties.

Definition: Given a decision tree for $X \subset \mathbb{F}$ let $\mathbf{p}_x$ denote the path T takes when input $x \in \mathbb{F}$ is provided. Let $|\mathbf{p}_x|$ denote the total number of internal nodes in $\mathbf{p}_x$. Let $\mathbf{p}_T$ be the *generic path* of T, corresponding to the decision " $\neq$ " at all decision nodes of T.

Definition: Say that a decision tree T is *pruned* if :
1. For every decision node ν in T there are inputs $x, y \in \mathbb{F}$ such that $\mathbf{p}_x, \mathbf{p}_y$ both pass through ν but take different forks there.
2. Every computation node has at least one decision node as a descendant.

Lemma 1: *Given any decision tree T there is a pruned decision tree T' deciding the same membership with $|T'| \leq |T|$.*

Proof: Simply remove the nodes of T that are not needed.

Definition: A decision tree T is *polynomial* if no computation node performs division.

Lemma 2: *Given a decision tree T there is a pruned polynomial decision tree T' deciding the same membership with $|T'| \leq 4|T|$*

Proof: If T is not pruned, prune it, assume henceforth that it is pruned. If it has no nodes performing division, do nothing. Otherwise proceed as follows:

1. Replace the input node ν with assignment $u_\nu \leftarrow x$ with an input node ν_1 with assignment $u_{\nu_1} \leftarrow x$ and and a 0-arity node ν_2 with $u_{\nu_2} \leftarrow 1$.
2. Replace every computation node ν with two nodes or four nodes as follows:
   (a) If ν is a 0-arity node $u_\nu \leftarrow c$ replace ν with nodes ν_1, ν_2 and let $u_{\nu_1} \leftarrow c, u_{\nu_2} \leftarrow 1$.
   (b) If ν is a 1-arity node $u_\nu \leftarrow c * u_\mu$ replace ν with nodes ν_1, ν_2 and let $u_{\nu_1} \leftarrow c * u_{\mu_1}, u_{\nu_2} \leftarrow u_{\mu_2}$.
   (c) If ν is a 2-arity node $u_\nu \leftarrow u_\mu / u_\lambda$ replace ν with nodes ν_1, ν_2 and let $u_{\nu_1} \leftarrow u_{\mu_1} * u_{\lambda_2}, u_{\nu_2} \leftarrow u_{\mu_2} * u_{\lambda_1}$.
   (d) If ν is a 2-arity node $u_\nu \leftarrow u_\mu + u_\lambda$ replace ν with nodes $\nu_1, \ldots \nu_4$ and let $u_{\nu_3} \leftarrow u_{\mu_1} * u_{\lambda_2}, u_{\nu_4} \leftarrow u_{\mu_2} * u_{\lambda_1}, u_{\nu_1} \leftarrow u_{\nu_3} + u_{\nu_4}, u_{\nu_2} \leftarrow u_{\nu_3} * u_{\nu_4}$.
   (e) If ν is a 2-arity node $u_\nu \leftarrow u_\mu - u_\lambda$ replace ν with nodes $\nu_1, \ldots \nu_4$ and let $u_{\nu_3} \leftarrow u_{\mu_1} * u_{\lambda_2}, u_{\nu_4} \leftarrow u_{\mu_2} * u_{\lambda_1}, u_{\nu_1} \leftarrow u_{\nu_3} - u_{\nu_4}, u_{\nu_2} \leftarrow u_{\nu_3} * u_{\nu_4}$.
   (f) If ν is a 2-arity node $u_\nu \leftarrow u_\mu * u_\lambda$ replace ν with nodes ν_1, ν_2 and let $u_{\nu_1} \leftarrow u_{\mu_1} * u_{\lambda_1}, u_{\nu_2} \leftarrow u_{\mu_2} * u_{\lambda_2}$.
3. For a decision node ν making the query "$u_\mu = u_\lambda$?" add two computation nodes ν_1, ν_2 with $u_{\nu_1} = u_{\mu_1} * u_{\lambda_2}, u_{\nu_2} = u_{\mu_2} * u_{\lambda_1}$ and then replace the decision node with a new node making the query "$u_{\nu_1} = u_{\nu_2}$?".

4. Connect all leaf nodes to the modifed decision nodes corresponding to the previous logic.

This procedure reflects the fact that $u_\nu = \frac{u_{\nu 1}}{u_{\nu 2}}$ and all operations using u_ν can be written in terms of u_{ν_1} and u_{ν_2}. Hence the value T calculates at a computation node with an input $x \in \mathbb{F}$ corresponds exactly to the values T' calculates at the corresponding nodes with the same input. Any decision node of T' is making a comparison on the same values as the corresponding node of T. It's clear that if T is pruned, so is T'. QED.

We introduce a new operation for modifying decision trees in a way that trades computation nodes for decsion nodes. Let T be a pruned polynomial decision tree for deciding membership of $X \subset \mathbb{F}$. Suppose λ is a leaf node of maximal distance from $\mathbf{p}_T$, as measured in the number of decisions nodes between λ and $\mathbf{p}_T$ and suppose $\lambda \notin \mathbf{p}_T$. Let μ the parent decision node of λ and suppose $\mu \notin \mathbf{p}_T$. Let ν be the parent decision node of μ. Since T is pruned there are two cases:

1. One child of ν is μ. One child of μ is λ. The other child of μ is another leaf node $\lambda 1$ with a label that is the opposite of the label of λ. The other child of ν is a leaf node $\lambda 2$. If necessary, switch the names of leaf nodes $\lambda, \lambda 1$ so that λ is the leaf node labeled " $\in$ ".

2. One child of ν is μ. The other child of ν is another decision node $\mu 1$. One child of μ is λ. The other child of μ is another leaf node $\lambda 1$ with a label which is the opposite of λ. The decision node $\mu 1$ has children, which are both leaf nodes and have opposite labels.

We describe a procedure RETRACT that replaces the entire segment of the tree which emanates from ν with a segment which has only one decision node ξ whose children are two leaf nodes. The new segment will have two leaves as children and will be labeled so that any input $x \in \mathbb{F}$ whose path enters that segment will now enter the new segment and terminate at a leaf with the same label as it had previously.

RETRACT Case I. The following table assumes that the leaf node λ is always labeled " $\in$ " and its sibling is always labeled " $\notin$ ". The label for leaf node varies as indicated. The middle column describes the path emanating from decision node ν and terminating at leaf node λ. The third column describes the decision function that the new node ξ is using.

$\lambda 2$ label	decisions for $\nu \to \lambda$	decision for ξ
$\in$	$u_\nu = 0, u_\mu = 0$	$(1 - u_\nu^{q-1})u_\mu$=0?
$\in$	$u_\nu = 0, u_\mu \neq 0$	$(1 - u_\nu^{q-1})(1 - u_\mu^{q-1}) = 0$?
$\in$	$u_\nu \neq 0, u_\mu = 0$	$u_\mu u_\nu = 0$?
$\in$	$u_\nu \neq 0, u_\mu \neq 0$	$(1 - u_\mu^{q-1})u_\nu = 0$?
$\notin$	$u_\nu = 0, u_\mu = 0$	$1 - (1 - u_\nu^{q-1})(1 - u_\mu^{q-1}) = 0$?
$\notin$	$u_\nu = 0, u_\mu \neq 0$	$1 - (1 - u_\nu^{q-1})u_\mu^{q-1} = 0$?
$\notin$	$u_\nu \neq 0, u_\mu = 0$	$1 - (1 - u_\mu^{q-1})u_\nu^{q-1} = 0$?
$\notin$	$u_\nu \neq 0, u_\mu \neq 0$	$1 - u_\nu^{q-1}u_\mu^{q-1} = 0$?

RETRACT Case II The following table describes the situation where the decision node has two children, both of which are decision nodes and each of these nodes in turn have two children, both of which are leaf nodes. The left column describes the labels of the four leaves numbered so that they correspond to the decisions first node is 0, second node is 0; first node is 0, second node is $\neq 0$; first node is $\neq 0$, second node is 0; first node is $\neq 0$, second node is $\neq 0$. The second column describes the decision function that the new node ξ is using.

labels on $\lambda, \lambda 1, \lambda 2, \lambda 3$	decision for ξ
$\notin\in\in\notin$	$(1 - u_\nu^{q-1})u_\mu + u_\mu u_\nu = 0$?
$\notin\in\in\notin$	$(1 - u_\mu^{q-1})(1 - u_\nu^{q-1}) + u_\mu u_\nu = 0$?
$\in\notin\notin\in$	$(1 - u_\nu^{q-1})u_\mu + (1 - u_\mu^{q-1})u_\nu = 0$?
$\notin\in\notin\in$	$(1 - u_\mu^{q-1})(1 - u_\nu^{q-1}) + (1 - u_\mu^{q-1})u_\nu = 0$?

In both cases the process of replacing the segment of T that emanates from the node ν and replacing it by a node ξ which has two children, both leaf nodes with opposite labels, is said to apply RETRACT on T at ν.

For a decision node ν, let $|\nu|$ denote the number of decision nodes that separate ν from the generic path $\mathbf{p}_T$. Let $\langle T \rangle = max_\nu |\nu|$ where ν ranges over all decision nodes of T.

Lemma 3: *If T is a decision tree for $X \subset \mathbb{F}$ that has $\langle T \rangle > 1$ then applying RETRACT on T at all decision nodes ν with $|\nu| = \langle T \rangle$ results in a new tree T' with $\langle T' \rangle = \langle T \rangle - 1$ and $|T'| \leq |T| + 6log(q)$.*

Proof: This is immediate from the construction.

Definition: Suppose T is a decision tree for $X \subset \mathbb{F}$ with $\langle T \rangle \leq 1$, so every leaf is a child of a decision node on $\mathbf{p}_T$ or a child of a decision node that is seperated from $\mathbf{p}_T$ only by computation nodes. Let D_0 be the set of decision nodes on $\mathbf{p}_T$. Let D_1 denote the decision nodes ν not on $\mathbf{p}_T$ such that when $x \in \mathbb{F}$ is entered $u_\nu = 0$ implies $x \in X$. Let D_2 denote the decision nodes ν not on $\mathbf{p}_T$ such that when $x \in \mathbb{F}$ is entered $u_\nu \neq 0$ implies $x \in X$. For $\nu \in D_1 \cup D_2$ let $\mu(\nu) \in \mathbf{p}_T$ denote the decision node that is the parent node (disregarding computation nodes).

Lemma 4: *Suppose T is a decision tree for $X \subset \mathbb{F}$ with $\langle T \rangle \leq 1$. Then*

$$\chi = \prod_{\nu \in D_0} u_\nu \times \prod_{\nu \in D_1} (1 - (1 - u_\nu^{q-1})(1 - u_{\mu(\nu)}^{q-1})) \times \prod_{\nu \in D_2} (1 - u_\nu(1 - u_{\mu(\nu)}^{q-1})$$

Proof of Theorem: First turn T into a polynomial pruned tree. Then apply RETRACT until one obtains a tree T' with $\langle T' \rangle \leq 1$. Now apply Lemma 4 to obtain the characteristic function of $X \subset \mathbb{F}$. Due to Lemmas 1 to 4, this gives the required upper bound on the complexity of χ.

We constrast this last theorem with the following result of Patterson and Stackmeyer (see [2]).

Theorem: *For each degree d there exists a hypersurface H_d in the space of polynomials $f(x) \in \bar{\mathbb{F}}[x]$ of degree d such that $f \notin H_d$ implies $L(f) \geq \sqrt{d} - 1$*

Defintion: We will call a polynomial $f(x) \in \bar{\mathbb{F}}[x]$ *generic* if $f \notin H_d$.

Definition: Given graphs Γ_1, Γ_2 of the same size with edge sets of cardinality n and a finite field $\mathbb{F}$ with $q = p^n$ elements let $\theta(x) = A_{\Gamma_1} + xA_{\Gamma_2}$,

$$\Phi_{\Gamma_1,\Gamma_2}(x) = \prod_\alpha f_\alpha(x)$$

where $\alpha : (1, \dots, n) \to \mathbb{F} - \{0, -1\}$ ranges over all possible maps and $f_\alpha(x) = gcd[\delta_0(\pi_\alpha(\theta(x))), \dots, \delta_{n-v}(\pi_\alpha(\theta(x)))]$ and δ_i is the determinant of the $v \times v$ matrix consisting of the first $v - 1$ columns of a matrix as well as the column in position $v + 1$ Let $\chi_{\Gamma_1,\Gamma_2} = {\Phi_{\Gamma_1,\Gamma_2}}^{q-1}$.

Conjecture: *For all n there exisits graphs Γ_1, Γ_2 of the same size with edge sets of cardinality n such that χ_{Γ_1,Γ_2} has degree exponential in n and is generic.*

We observe that the conjecture immediately implies $NP \neq P$ (see [4]).

References:

1. Bürgisser, Peter; Clausen, Michael; Shokrollahi, M. Amin (1997). Algebraic complexity theory. Grundlehren der Mathematischen Wissenschaften. Vol. 315. With the collaboration of Thomas Lickteig. Berlin: Springer-Verlag. ISBN 978-3-540-60582-9. Zbl 1087.68568.
2. Paterson, Michael S.; Stockmeyer, Larry J. (1973). "On the Number of Nonscalar Multiplications Necessary to Evaluate Polynomials". SIAM Journal on Computing. 2 (1): 60–66. doi:10.1137/0202007.
3. Bruns, Winfried; Vetter, Udo (1988). Determinantal rings. Lecture Notes in Mathematics. Vol. 1327. Springer-Verlag. doi:10.1007/BFb0080378. ISBN 978-3-540-39274-3.

4. Cook, Stephen (1971). "The complexity of theorem proving procedures". Proceedings of the Third Annual ACM Symposium on Theory of Computing. pp. 151–158. doi:10.1145/800157.805047. ISBN 9781450374644. S2CID 7573663.

Acknowlegements: The author wishes to thank RIMS for hosting this conference, Tsunekazu Nishinaka for running it, and Hisa Tsutsui for arranging my attendance.