HAMBURGER BEITRÄGE ZUR MATHEMATIK

Heft 123

Application of linear matroids to graph colorability

T.R. Jensen, Hamburg

Juli 2001

Application of linear matroids to graph colorability

Tommy R. Jensen
Fachbereich Mathematik
Bundesstraße 55
D-20146 Hamburg
Germany

July 13

2001

Abstract

The paper considers generalizing the concept of graph coloring to matroids having represention over finite fields. We describe a class of algorithms which apply elementary linear algebra to this generalized coloring problem. It is shown that these algorithms have polynomial running time, whereas they do not always terminate in a decided state. We apply these results to proving for fixed $k > 0$, that the non-k-colorability of certain polynomially recognizable classes of graphs can be decided in polynomial time.

1 Introduction

The generalization of graph coloring to linear matroids or finite projective spaces has deep connections with the study of the classical four color problem. Already early in the 20th century various directions were pointed out by O. Veblen (projective geometries), G.D. Birkhoff (counting the number of ways of coloring a map) and H. Whitney (on matroids and on coloring of graphs). A culmination has been reached with the work done up until the early 70s of W.T. Tutte (on the Tutte

Polynomial for matroids), H.H. Crapo and G.-C. Rota (combinatorial geometries), T.H. Brylawski (characterization of Tutte-Grothendieck invariants), among several others. The more recent survey [4] gives a succinct account of many of the results obtained in this area.

One of the main aims of the classical lines of research (see e.g. [5]) is to establish properties of the so-called Tutte Polynomial associated with a matroid. This polynomial generalizes the classical chromatic polynomial of a graph. These properties should eventually be used to deduce a generalized form of the four color theorem based on a Kuratowski-type excluded minor characterization of the relevant class of matroids. This would result in an almost purely algebraic, nearly non-combinatorial proof of the four color theorem. However, such a research program has yet to be carried out to success. In more recent times the activity directed towards this goal has been somewhat reduced following the combinatorial proof of the four color theorem by K. Appel, W. Haken and J. Koch in 1976 ([1]).

Since it is, from experience, extremely difficult to calculate the chromatic polynomial of a graph, there is not much promise of obtaining a practical method along these lines for deciding about colorability of any given planar graph, or of any given general graph for that matter.

The objective of this paper is to give a preliminary study of a class of algorithms for graph coloring, or more generally for coloring of linear matroids. The underlying principle of these algorithms is the familiar generalization of colorability from graphs to linear matroids. From the point of view of practical applicability the suggested method has the advantage, as opposed to the classical approach, that it does not require any calculations of polynomial invariants. Instead it applies direct reduction arguments to sequences of representation matrices. This leads to a class of algorithms that use only elementary techniques from linear algebra and in addition have polynomial time complexity.

On the negative side (as it should indeed be expected from the fact that this paper does not claim to prove that the two complexity classes $\mathcal{NP}$ and co-$\mathcal{NP}$ are equal) there is in general no guarantee given that such an algorithm will always decide correctly about colorability of any input. For the majority of these algorithms there are instances for which they terminate in an undecided state, neither giving a positive nor a negative answer to the colorability question. Although for the problem of deciding 2-colorability of a graph, or more generally a binary matroid, which is a problem known to be of polynomial complexity, it is observed that there is a simple algorithm of this class which always gives the correct answer for any input.

The main result of this paper is a sufficient condition to ensure that an algorithm of the investigated type has the property that its output, whether positive, negative or undecided, depends on the essential properties of its input, as opposed to depending on any particular choice of its presentation. (As an example, the edges and/or vertices of a graph might be presented to the algorithm in various different orders, and the states of the algorithm at termination could conceivably be quite different at each such new run.) This seemingly weak type of result implies a new type of classification of graphs with high chromatic number, as well as a corresponding classification of graphs colorable with a given number of colors. A particular example is a non-trivial classification of 4-chromatic graphs into two classes; one of these classes (for smaller graphs it is, by the way, far the largest of the two) is polynomially recognizable by a special algorithm $\mathcal{A}[3, U_4^2]$.

2 An example

A simple example will serve to illustrate some of the main ideas in an elementary setting, before they are formally developed in the following sections. The example demonstrates the primary reduction step of the algorithm $\mathcal{A}[3, U_4^2]$, to be defined later.

Example. Assume that a graph G is given as pictured in Figure 1, with vertices labeled from 1 to 8, and that it is required to test whether the vertices of G can be properly colored with three colors, that is, whether each vertex can be assigned one of a set of three colors in such a way that no two vertices are colored alike if they are joined by an edge in G.

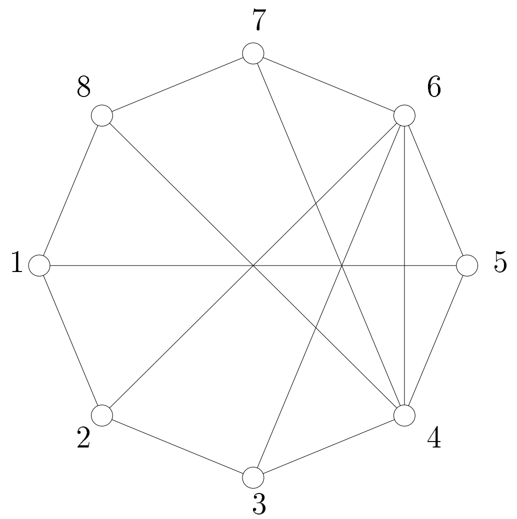

Figure 1. A graph G.

The following matrix $\boldsymbol{A}_G$ with entries in $\mathbb{F}_3$, the field of three elements $\{0, -1, +1\}$ will be associated with the graph G. The rows of $\boldsymbol{A}_G$ are labeled with the labels 1–8 of the vertices of G, and for all labels $i < j$ such that the corresponding vertices of G are neighbors, there is a column of $\boldsymbol{A}_G$ with $+1$ as its ith entry and -1 as its jth entry. In this way the columns of $\boldsymbol{A}_G$ correspond to the edges of G.

$$\boldsymbol{A}_G = \left[\begin{array}{cccccccccccccc}
+1 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & 0 & 0 & 0 & 0 & 0 \\
-1 & +1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & 0 & 0 & 0 & 0 \\
0 & -1 & +1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & 0 & 0 & 0 \\
0 & 0 & -1 & +1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & +1 \\
0 & 0 & 0 & -1 & +1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -1 & +1 & 0 & 0 & 0 & -1 & -1 & -1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & -1 & +1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & -1
\end{array}\right]
\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \end{array}$$

The following observation is almost from the construction of $\boldsymbol{A}_G$ and the fact that there is a proper 3-coloring of G precisely when there is a proper coloring using the set of colors $\{0, -1, +1\}$.

$$\begin{gathered}\text{G is 3-colorable}\\ \Updownarrow \\ \exists \boldsymbol{x} \in {\mathbb{F}_3}^{V(G)} : \boldsymbol{x}\boldsymbol{A}_G \in (\mathbb{F}_3 \setminus \{0\})^{E(G)}\end{gathered}$$

Defining that $\boldsymbol{y} \in {\mathbb{F}_3}^m$ is *nowhere-zero* if $\boldsymbol{y} \in (\mathbb{F}_3 \setminus \{0\})^m$ the above statement can be reformulated as follows.

$$\begin{gathered}\text{G is 3-colorable}\\ \Updownarrow \\ \text{The row-space of $\boldsymbol{A}_G$ contains a nowhere-zero element.}\end{gathered}$$

Since the latter property is just a property of the row-space of $\boldsymbol{A}_G$, the following is again equivalent.

$$\begin{gathered}\text{G is 3-colorable}\\ \Updownarrow \\ \text{The row space of $\boldsymbol{A}'$ contains a nowhere-zero element,}\\ \text{where $\boldsymbol{A}'$ is any matrix obtained from $\boldsymbol{A}_G$ by successive}\end{gathered}$$

(i) elementary row operations

(ii) insertion or deletion of linearly dependent rows

(iii) scaling of columns by nonzero scalars.

The following matrix $\boldsymbol{A}'$ is easily obtained from $\boldsymbol{A}_G$ by a succession of elementary row operations followed by deletion of a row of all zeroes, hence the 3-colorability of G can be decided by inspection of $\boldsymbol{A}'$ rather than $\boldsymbol{A}_G$.

$$
\boldsymbol{A}' = \begin{bmatrix}
+1 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & 0 & 0 & 0 & 0 & 0 \\
0 & +1 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & +1 & 0 & 0 & 0 & 0 \\
0 & 0 & +1 & 0 & 0 & 0 & 0 & +1 & +1 & +1 & +1 & 0 & 0 & 0 \\
0 & 0 & 0 & \underline{+1} & \underline{0} & 0 & 0 & +1 & +1 & +1 & +1 & \underline{+1} & +1 & +1 \\
0 & 0 & 0 & \underline{0} & \underline{+1} & 0 & 0 & +1 & 0 & +1 & +1 & \underline{+1} & +1 & +1 \\
0 & 0 & 0 & 0 & 0 & +1 & 0 & +1 & 0 & 0 & 0 & 0 & +1 & +1 \\
0 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & 0 & 0 & 0 & 0 & 0 & +1
\end{bmatrix}
$$

The underlined entries of $\boldsymbol{A}'$, and the fact that the remaining entries of the corresponding columns are exclusively zeroes, supply the information about the fourth and fifth coordinates x'_4 and x'_5 of $\boldsymbol{x}'$ that if $\boldsymbol{x}'\boldsymbol{A}'$ is nowhere-zero, then

$$
\begin{array}{ccccc}
x'_4 & & & \neq & 0 \\
 & & x'_5 & \neq & 0 \\
x'_4 & + & x'_5 & \neq & 0
\end{array}
$$

It is easily checked that the only solutions in $\mathbb{F}_3$ are

$$
(x'_4, x'_5) \in \{(-1,-1), (+1,+1)\},
$$

and in particular $x'_4 = x'_5$ holds for any $\boldsymbol{x}'$ such that $\boldsymbol{x}'\boldsymbol{A}'$ is nowhere-zero. This easily implies that the problem reduces to deciding whether the row-space of the matrix $\boldsymbol{A}''$ contains a nowhere-zero element, where $\boldsymbol{A}''$ is obtained by replacing the fourth and fifth row of $\boldsymbol{A}'$ by their sum:

$$
\boldsymbol{A}'' = \begin{bmatrix}
+1 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & 0 & 0 & 0 & 0 & 0 \\
0 & +1 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & +1 & 0 & 0 & 0 & 0 \\
0 & 0 & \underline{+1} & \underline{0} & 0 & 0 & 0 & +1 & +1 & +1 & \underline{+1} & 0 & 0 & 0 \\
0 & 0 & \underline{0} & \underline{+1} & +1 & 0 & 0 & -1 & +1 & -1 & \underline{-1} & -1 & -1 & -1 \\
0 & 0 & 0 & 0 & 0 & +1 & 0 & +1 & 0 & 0 & 0 & 0 & +1 & +1 \\
0 & 0 & 0 & 0 & 0 & 0 & +1 & +1 & 0 & 0 & 0 & 0 & 0 & +1
\end{bmatrix}
$$

The similar argument may now be applied with respect to the third and fourth row of $\boldsymbol{A}''$: If $\boldsymbol{x}''\boldsymbol{A}''$ is nowhere-zero, then

$$\begin{array}{rcrcl} x_3'' & & & \neq & 0 \\ & & x_4'' & \neq & 0 \\ x_3'' & - & x_4'' & \neq & 0 \end{array}$$

and it follows now that $x_3'' = -x_4''$, and $\boldsymbol{A}''$ may be replaced by another matrix obtained by replacing the third and fourth rows by their difference.

It is easy to continue this process, and to verify that it eventually terminates, in the case of this particular graph G, when the final matrix $\boldsymbol{A}^*$ has but just a single row left, a row which contains at least one zero entry. This trivially impies that the row-space of $\boldsymbol{A}^*$ contains no nowhere-zero elements, and hence it proves algebraically that G cannot be 3-colored. $\square$

3 Terminology and notation

For the purpose of making the presentation as self-contained as possible, this section contains a number of standard definitions mainly from the area of matroid theory. For more thorough introductions to this topic, one may consult [8, 10, 11].

The finite q-element field is denoted by $\mathbb{F}_q$, and the n-dimensional vector space ${\mathbb{F}_q}^n$ by $\mathcal{W}_n$, $n > 0$. In the following, all scalars and entries of vectors and matrices belong to $\mathbb{F}_q$, unless explicitly stated otherwise. Unless otherwise specified, vectors of dimension n are $n \times 1$ matrices (column vectors) and are written with bold-face small letters $\boldsymbol{e}, \boldsymbol{f}, \boldsymbol{x}, \boldsymbol{y}, \ldots$ etc. General matrices are written with bold-face capital letters $\boldsymbol{A}, \boldsymbol{B}, \ldots$ etc. Transposes are denoted by $\boldsymbol{e}^t$, $\boldsymbol{A}^t$, etc. If $\boldsymbol{A}_1$ and $\boldsymbol{A}_2$ are matrices with the same number of rows, then $\boldsymbol{A}_1|\boldsymbol{A}_2$ denotes the matrix whose columns are those of $\boldsymbol{A}_1$ followed by those of $\boldsymbol{A}_2$ from left to right; in particular $\boldsymbol{A}|\boldsymbol{e}$ is obtained by adjoining a new rightmost column $\boldsymbol{e}$ to $\boldsymbol{A}$. The set of *nowhere-zero* row vectors $\{\boldsymbol{e}^t : \boldsymbol{e} \in (\mathbb{F}_q \setminus \{0\})^m\}$ is denoted by $\mathcal{Y}_m$, $m > 0$.

Let $\boldsymbol{A} \in \mathbb{F}_q{}^{n\times m}$ be an $n \times m$ matrix, where either both n and m are positive integers, or $m = 0$ (in which case $\boldsymbol{A}$ is called *empty* and $n \geq 0$ may have any value). The *linear matroid* $M = M(\boldsymbol{A})$ *represented by* $\boldsymbol{A}$ is defined as the pair $(E, \mathcal{I})$, where the set of M's *elements* E is the m-set of columns in $\boldsymbol{A}$ and the set $\mathcal{I}$ of *independent sets* of M is the set of subsets of E that are linearly independent in $\mathcal{W}_n$. A minimal dependent subset of E is a *circuit* of M (note that the pair $(E, \mathcal{I})$ is uniquely determined by specifying the set E together with the set of circuits). A *loop* is an element $\boldsymbol{e} \in E$ such that $\{\boldsymbol{e}\}$ is a circuit. It is easily checked that the relation $\boldsymbol{e} \sim \boldsymbol{f}$, for $\boldsymbol{e}, \boldsymbol{f} \in M$, of belonging to a common circuit in M define an equivalence relation on the elements of M. The equivalence classes with respect to $\sim$ are the *components* of M, and M is *connected* if M has precisely one component. The *rank* of M, denoted by $r(M)$, is the largest cardinality of any member of $\mathcal{I}$; this is of course precisely the rank of $\boldsymbol{A}$. The *empty matroid* is the matroid $(\emptyset, \{\emptyset\})$ represented by the empty matrix. In the general axiomatic setting of matroid theory, see [10, 11], there exist matroids which do not originate from matrices, that is, matroids that are not linear. We will not be concerned with the general setting here, and proceed to use the term 'matroid' synonymously with 'linear matroid'.

$M_1 = (E_1, \mathcal{I}_1)$ and $M_2 = (E_2, \mathcal{I}_2)$ are *isomorphic*, $M_1 \simeq M_2$, if there exists a bijection $\varphi : E_1 \to E_2$ such that

$$S \in \mathcal{I}_1 \;\Leftrightarrow\; \varphi(S) \in \mathcal{I}_2 \;\text{ for all } S \subseteq E_1.$$

The notation $\boldsymbol{A}_1 \approx \boldsymbol{A}_2$ is used when $M_1 = M(\boldsymbol{A}_1)$, $M_2 = M(\boldsymbol{A}_2)$, and the isomorphism φ can be chosen as the bijection that maps the ith column of $\boldsymbol{A}_1$ to the ith column of $\boldsymbol{A}_2$, for $i = 1, 2, \ldots, m$. It is not difficult to show using arguments from elementary linear algebra (albeit performed over $\mathbb{F}_q$) that $\boldsymbol{A}_1$ and $\boldsymbol{A}_2$ satisfy $\boldsymbol{A}_1 \approx \boldsymbol{A}_2$ if one of the matrices can be obtained from the other by left-multiplying with an $n \times n$ nonsingular matrix, or by replacing any column by some nonzero scalar multiple of itself; in particular $M(\boldsymbol{A}_1)$ and $M(\boldsymbol{A}_2)$ will then be isomorphic.

By elimination, if necessary, of linearly dependent rows from $\boldsymbol{A}$ it may be assumed that $\boldsymbol{A}$ is chosen as a representation of M with a minimal number of rows, that is, either the number is equal to the rank of $\boldsymbol{A}$, or it is equal to 1, in which case it may happen that $\boldsymbol{A}$ has only zero entries.

The *dual matroid* M^* of M is defined as the matroid $M(\boldsymbol{A}^*)$, where $\boldsymbol{A}^*$ is a $n^* \times m$ matrix, $n^* \geq 0$, satisfying $\boldsymbol{A}^* \boldsymbol{A}^t = \boldsymbol{0}$ and of maximal rank subject to this. It is not difficult to see that a circuit of M^* corresponds precisely to a minimal subset S of columns in $\boldsymbol{A}$ such that there exists a nonzero linear combination of the rows in $\boldsymbol{A}$ having nonzero coefficients precisely corresponding to the elements of S. Thus the set of circuits for M^* is uniquely determined, and in particular M^* is well-defined. Using the natural bijection between the element sets of M and M^* we speak of a subset of E corresponding to a circuit of M^* as a *cocircuit* of M.

For $M = M(\boldsymbol{A}) = (E, \mathcal{I})$ and $\boldsymbol{e} \in E$, there is a new matroid $M - \boldsymbol{e} = (E \setminus \{\boldsymbol{e}\}, \mathcal{I}(M - e))$, obtained by the *deletion* of $\boldsymbol{e}$ from M, where

$$\mathcal{I}(M - \boldsymbol{e}) = \{I \in \mathcal{I} \mid \boldsymbol{e} \notin I\}.$$

It is clear that $M - \boldsymbol{e}$ is equal to the matroid represented by $\boldsymbol{A}'$, obtained by deleting the column $\boldsymbol{e}$ from $\boldsymbol{A}$. For any $X \subseteq E$, the matroid $M - X$ arises from M by successive deletion of the elements in X (where if $X = E$, then $M - X = \emptyset$). For $X \subseteq E$ (possibly $X = \emptyset$) and $M' = M - X$ we call M' a *submatroid* of M, written $M' \subseteq M$.

Another matroid $M/\boldsymbol{e}$ obtained by *contraction* of $\boldsymbol{e} \in E$ is defined by $M/\boldsymbol{e} = (E \setminus \{\boldsymbol{e}\}, \mathcal{I}(M/\boldsymbol{e}))$, where, if $\boldsymbol{e}$ is not a loop of M,

$$\mathcal{I}(M/\boldsymbol{e}) = \{I \in \mathcal{I} \mid \boldsymbol{e} \notin I \text{ and } I \cup \{\boldsymbol{e}\} \in \mathcal{I}\},$$

and $\mathcal{I}(M/\boldsymbol{e}) = \mathcal{I}(M - \boldsymbol{e})$ otherwise. It follows immediately from this definition that the order of contractions is immaterial: $\mathcal{I}((M/\boldsymbol{e})/\boldsymbol{f}) = \mathcal{I}((M/\boldsymbol{f})/\boldsymbol{e})$ for all $\boldsymbol{e}, \boldsymbol{f} \in M$. Hence for $X \subseteq E$ the matroid M/X obtained by successively contracting the elements of X is well-defined. The matroid $M(\boldsymbol{A}_1|\boldsymbol{A}_2)/X$ where X is the set of columns for $\boldsymbol{A}_2$ is conveniently written as $M(\boldsymbol{A}_1|\boldsymbol{A}_2)/\boldsymbol{A}_2$ etc.

A matrix representing $M/\boldsymbol{e}$ in the case when $\boldsymbol{e}$ is not a loop may be obtained as follows. Let $\boldsymbol{B}$ be any non-singular matrix such that $\boldsymbol{Be}$ has precisely one nonzero coordinate (such a matrix exists when $\boldsymbol{e}$ is not a loop), say in the coordinate numbered i. Then it is not difficult to show that $M/\boldsymbol{e}$ is isomorphic to the matroid represented by the matrix obtained from $\boldsymbol{BA}$ by removing row i (unless row i is the unique row of $\boldsymbol{BA}$, in which case row i is replaced by the all-zero row of the same length) and deleting the column $\boldsymbol{Be}$.

The following is an easy consequence of the definitions, or can be proved from the representations of matroids obtained by deletions and contractions.

Proposition 3.1 *For all $\boldsymbol{e}, \boldsymbol{f} \in M$,*

$$(M - \boldsymbol{e})/\boldsymbol{f} = (M/\boldsymbol{f}) - \boldsymbol{e}.$$

$\square$

4 Coloring of linear matroids

A *graph* may be defined as a mapping $\psi : E \to V \times V$, where V and E are finite sets of *vertices* and *edges*, respectively.† As it is customary for notational reasons, we will refer to the triple $G = (V, E, \psi)$ as 'the graph G'. For $u, v \in V$ such that there exists $e \in E$ with $\psi(e) = (u, v)$, we say that u and v are mutually *adjacent*. For further graph terminology see [2].

A *k-coloring* of G, where k is a nonnegative integer, is a mapping $\varphi : V \to K$, where $|K| = k$, such that if u and v are adjacent, then $\varphi(u) \neq \varphi(v)$. The determination of the *chromatic number* $\chi(G)$, the smallest k for which a k-coloring of G exists, is a well-known $\mathcal{NP}$-hard combinatorial problem. In fact, for every fixed $k \geq 3$ it is an $\mathcal{NP}$-hard problem to decide whether $\chi(G) \leq k$. For definitions related to the class $\mathcal{NP}$ and the complexity of decision problems see [6].

†In normal usage, this actually defines a *directed graph*, whose associated undirected graph arises after replacing each pair $(u, v) \in V \times V$ by the set $\{u, v\} \in V^2 \cup V^1$. We will ignore this distinction here.

Let $\boldsymbol{A}_G$ be the matrix with set of rows V and set of columns E, such that the (u,e)-entry is equal to

$$a_{ue} = \begin{cases} 1 & \text{if } \psi(e) = (u,v) \text{ for some } v \in V \setminus \{u\} \\ -1 & \text{if } \psi(e) = (v,u) \text{ for some } v \in V \setminus \{u\} \\ 0 & \text{otherwise.} \end{cases}$$

Then the *matroid for* G, denoted by $M(G)$, may be defined as the matroid $M(\boldsymbol{A}_G)$. It is not hard to check that any subset $X \subseteq E$ is dependent in $M(G)$, if and only if the subgraph $G[X]$ induced by X contains a cycle.

Proposition 4.1 *$G = (V, E, \psi)$ with $E \neq \emptyset$ is q-colorable if and only if the row space of $\boldsymbol{A}_G$ intersects $\mathcal{Y}_{|E|}$.*

Proof: Since $\mathbb{F}_q$ has q elements, G is by definition q-colorable if and only if there exists $\varphi : V \to \mathbb{F}_q$ such that $\varphi(u) - \varphi(v) \neq 0$ for all $(u,v) \in \psi(E)$. This is equivalent to $(\boldsymbol{\varphi})^t \boldsymbol{A}_G \in \mathcal{Y}_{|E|}$, where $\boldsymbol{\varphi} = (\varphi(u) : u \in V)$. $\square$

Motivated by Proposition 4.1, we define that the matroid M is *q-colorable* (or simply *colorable*) if M is empty, or if there exists some matrix $\boldsymbol{A}$ such that $M = M(\boldsymbol{A})$ and such that the row space of $\boldsymbol{A}$ intersects $\mathcal{Y}_m$, where $m = |E|$. In particular, by Proposition 4.1, if $\chi(G) \leq q$ holds, then $M(G)$ is q-colorable. It is an interesting fact that the converse is true; if $M(G)$ is q-colorable, then this implies $\chi(G) \leq q$.

This is a consequence of the following more general result.

Theorem 4.1 (Crapo & Rota, [5]) *For all* $\boldsymbol{A}_1, \boldsymbol{A}_2 \in \mathbb{F}_q{}^{n\times m}$, $n, m > 0$, *such that* $M(\boldsymbol{A}_1) = M(\boldsymbol{A}_2)$,

$$|\text{Row-Space}(\boldsymbol{A}_1) \cap \mathcal{Y}_m| = |\text{Row-Space}(\boldsymbol{A}_2) \cap \mathcal{Y}_m|.$$

□

By Theorem 4.1, in order to decide whether M is colorable, one may consider any arbitrary matrix $\boldsymbol{A}$ that represents M. This fact will be used silently throughout the remainder of this paper.

The following is an important example of a matroid that is not q-colorable. Let U^2_{q+1} denote a matroid whose element set has size $q+1$, and whose independent sets are all its subsets of size at most 2 (this is the *uniform matroid* of order $q+1$ and rank 2). It may be represented over $\mathbb{F}_q$ by the matrix

$$\begin{pmatrix} 1 & 0 & 1 & 1 & 1 & \dots & 1 \\ 0 & 1 & 1 & \alpha_2 & \alpha_3 & \dots & \alpha_{q-1} \end{pmatrix},$$

when listing the elements of $\mathbb{F}_q$ as $0, 1 = \alpha_1, \alpha_2, \dots, \alpha_{q-1}$. For every $\boldsymbol{x} \in \mathcal{W}_2$ there is some column $\boldsymbol{e}$ of this matrix such that $\boldsymbol{x}^t\boldsymbol{e} = 0$, thus in particular U^2_{q+1} is not colorable. It is not hard to check that U^2_{q+1} is in fact the unique loopless uncolorable matroid of rank at most 2.

It may seem at first that the concept of a q-coloring is too restrictive compared with the concept of a k-coloring, in which k may range freely over the non-negative integers. But this is only superficially so, since if q is a prime power such that $q > k \geq 0$, then the k-colorings of a graph G can essentially be identified with the q-colorings of a graph obtained by joining each vertex of G to $q - k$ mutually adjacent new vertices.

5 Reductions for matroid colorability

A few immediate consequences of the definition of matroid coloring are collected in the following Proposition.

Proposition 5.1 *For all M,*
a) If M contains a loop, then M is not colorable,
b) if M is colorable and $M' \subseteq M$, then M' is colorable,
c) if $\{\boldsymbol{e}, \boldsymbol{f}\}$ is a 2-element circuit of M, then M is colorable if and only if $M - \boldsymbol{f}$ is colorable.

Proof: If $\boldsymbol{e}$ is a loop of $M = M(\boldsymbol{A})$, then $\boldsymbol{A}$ has a column of all zeroes, and the row space of $\boldsymbol{A}$ cannot intersect $\mathcal{Y}_m$. This proves a).

Assume $M = M(\boldsymbol{A})$, $\boldsymbol{A} = \boldsymbol{A}_1|\boldsymbol{A}_2$ and $M' \simeq M(\boldsymbol{A}_1)$, where $\boldsymbol{A}_1 \in \mathbb{F}_q{}^{n\times m_1}$; we may assume that $\boldsymbol{A}_2$ is nonempty, otherwise the statement is trivially true. Since M is colorable there exists $\boldsymbol{x} \in \mathcal{W}_n$ such that $\boldsymbol{x}^t\boldsymbol{A} \in \mathcal{Y}_m$. Then clearly $\boldsymbol{x}^t\boldsymbol{A}_1 \in \mathcal{Y}_{m_1}$, and M' is colorable.

Finally, assume that $\{\boldsymbol{e}, \boldsymbol{f}\}$ is a circuit of $M = M(\boldsymbol{A})$ of size 2. If M is colorable, then $M - \boldsymbol{f}$ is colorable by b). Assume therefore that $M - \boldsymbol{f}$ is colorable, and that $\boldsymbol{x}^t\boldsymbol{A}' \in \mathcal{Y}_{m-1}$, where $\boldsymbol{A}'$ is the matrix obtained from $\boldsymbol{A}$ after deleting column $\boldsymbol{f}$. Since $\{\boldsymbol{e}, \boldsymbol{f}\}$ is a minimal linearly dependent set, $\boldsymbol{f}$ must be a nonzero scalar multiple of $\boldsymbol{e}$. Thus $\boldsymbol{x}^t\boldsymbol{e} \neq 0$ implies $\boldsymbol{x}^t\boldsymbol{f} \neq 0$, and $\boldsymbol{x}^t\boldsymbol{A} \in \mathcal{Y}_m$ follows. □

We will now consider two 'reduction tools' used in formulating a matroid coloring algorithm. The first is a generalization of a well-known statement concerning the reducibility of small edge cuts in graphs.

Proposition 5.2 *If C^* is a cocircuit of M with $|C^*| < q$, then M is colorable if and only if $M - C^*$ is colorable.*

Proof: If M is colorable, then so is $M - C^*$, by Proposition 5.1 b). Assume therefore that $M - C^*$ is colorable.

With $M = M(\boldsymbol{A})$, where $\boldsymbol{A} = \boldsymbol{A}_1|\boldsymbol{A}_2$ and $M - C^* = M(\boldsymbol{A}_1)$, there exists by assumption $\boldsymbol{x} \in \mathcal{W}_n$ satisfying $\boldsymbol{x}^t\boldsymbol{A}_1 \in \mathcal{Y}_{m-k}$. Since C^* is a cocircuit of M, there exists $\boldsymbol{y} \in \mathcal{W}_n$ such that $\boldsymbol{y}^t\boldsymbol{A}_1 = \boldsymbol{0}$ and $\boldsymbol{y}^t\boldsymbol{A}_2 \in \mathcal{Y}_{|C^*|}$. Since $|C^*| < q = |\mathbb{F}_q|$ it follows easily that at least one of the vectors $(\boldsymbol{x} + \alpha\boldsymbol{y})^t\boldsymbol{A}$, $\alpha \in \mathbb{F}_q$, belongs to $\mathcal{Y}_m$, and M is colorable.

$\square$

The following observation is also quite well-known. It can be viewed as related both to the recursive deletion/contraction procedure for calculating the Tutte Polynomial for a matroid (see [10, 11]), as well as to the identification step of the classical Hajós-construction of all graphs with a given chromatic number (see e.g. [7]).

Proposition 5.3 *Assume that M is not colorable and let $\boldsymbol{e} \in M$. Then $M - \boldsymbol{e}$ is colorable if and only if $M/\boldsymbol{e}$ is colorable.*

Proof: Let $M = M(\boldsymbol{A})$, where $\boldsymbol{A} = \boldsymbol{A}'|\boldsymbol{e} \in \mathbb{F}_q{}^{n\times m}$ and $M - \boldsymbol{e} = M(\boldsymbol{A}')$. First, if $\boldsymbol{e}$ is a loop, then $M/\boldsymbol{e} = M - \boldsymbol{e}$ and there is nothing to prove. So we may assume that there exists an $n \times n$ invertible matrix $\boldsymbol{B}$ such that $\boldsymbol{B}\boldsymbol{e} = \boldsymbol{e}_n$, where $\boldsymbol{e}_i$ denotes the ith element of the standard basis for $\mathcal{W}_n$.

By definition, $M - \boldsymbol{e}$ is colorable if and only if there exists $\boldsymbol{x} \in \mathcal{W}_n$ such that $\boldsymbol{x}^t\boldsymbol{A}' \in \mathcal{Y}_m$, and (since M is not colorable) such that $\boldsymbol{x}^t\boldsymbol{e} = 0$.

On the other hand, considering the representation of $M/\boldsymbol{e}$ obtained by deleting the nth row of $\boldsymbol{BA}$ and deleting the column corresponding to $\boldsymbol{e}$, $M/\boldsymbol{e}$ is colorable if and only if there exists $\boldsymbol{y} \in \mathcal{W}_n$ such that

$$\boldsymbol{y}^t(\boldsymbol{BA}') \in \mathcal{Y}_m, \text{ and } \boldsymbol{y}^t\boldsymbol{e}_n = 0.$$

Hence if either such an $\boldsymbol{x}$ or such a $\boldsymbol{y}$ exists, then the other also exists, and it may be found by solving $\boldsymbol{x} = \boldsymbol{B}^t\boldsymbol{y}$. □

6 An algorithm for coloring linear matroids

The following is a description of a simple general scheme for a polynomial time matroid coloring algorithm. The algorithms will be *inexact*, that is, whereas they will decide correctly the question of colorability for some input matroids (given by matrix representations over $\mathbb{F}_q$), the question will be left undecided for certain other matroids. Whether the problem is decided or not for a given input matroid may in general depend on which matrix is chosen as its representation. But as we will see in the next section, this behavior may be avoided.

Assume that a fixed finite set of matroids $\mathcal{M} = \{M_1, M_2, \ldots, M_k\}$ has been chosen, such that each M_i is representable over $\mathbb{F}_q$, contains no loops and is uncolorable. The set $\mathcal{M}$ will be called the *basis* for the algorithm, and the algorithm itself will be denoted by $\mathcal{A}[q, \mathcal{M}]$, or shorter $\mathcal{A}[\mathcal{M}]$ when there is no danger of confusion.

Let $M = M(\boldsymbol{A})$ be the matroid whose colorability is about to be tested.

Algorithm: $\mathcal{A}[\mathcal{M}]$

Input: $\boldsymbol{A} \in \mathbb{F}_q{}^{n \times m}$.

Phase 1:

1.0 If $\boldsymbol{A}$ is empty, then M is colorable by definition. If not, and if $\boldsymbol{A}$ contains a column of all zeroes, then M is uncolorable by Proposition 5.1 a). In either of the two cases, the algorithm terminates. Otherwise it proceeds to Step 1.1.

1.1 Determine whether it is possible to adjoin a new nonzero column $\boldsymbol{e}$ to $\boldsymbol{A}$ such that there exists $M_i \in \mathcal{M}$ where M_i is isomorphic to a submatroid of $M(\boldsymbol{A}|\boldsymbol{e})$.

Assuming there exists such an $\boldsymbol{e}$, then by Propositions 5.1 b) and 5.3, $M = M(\boldsymbol{A}|\boldsymbol{e}) - \boldsymbol{e}$ is colorable if and only if $M(\boldsymbol{A}|\boldsymbol{e})/\boldsymbol{e}$ is colorable. Now calculate a matrix $\boldsymbol{A}'$ such that $M(\boldsymbol{A}|\boldsymbol{e})/\boldsymbol{e} \simeq M(\boldsymbol{A}')$. Note that the rank of $\boldsymbol{A}'$ is strictly smaller than the rank of $\boldsymbol{A}$ (since $\boldsymbol{e}$ is not a loop of $M(\boldsymbol{A}|\boldsymbol{e})$). Return to Step 1.0 with $\boldsymbol{A}'$ in place of $\boldsymbol{A}$.

If no such $\boldsymbol{e}$ exists, then continue with Step 2.1.

We may assume that the first stage terminates with a nonempty matroid $\tilde{M} = M(\tilde{\boldsymbol{A}})$ without loops, and with the property that no nonzero column $\boldsymbol{e}$ can be adjoined to $\tilde{\boldsymbol{A}}$ resulting in a submatroid of $M(\tilde{\boldsymbol{A}}|\boldsymbol{e})$ which is isomorphic to some $M_i \in \mathcal{M}$.

Phase 2:

2.1 If $\tilde{M}$ contains a 2-circuit $\{\boldsymbol{e}, \boldsymbol{f}\}$, then $\tilde{M}$ is colorable if and only if $\tilde{M} - \boldsymbol{f}$ is colorable, by Proposition 5.1 c); replace $\tilde{M}$ by $\tilde{M} - \boldsymbol{f}$ and repeat Step 2.1. If no 2-circuit exists, go to Step 2.2.

2.2 If $\tilde{M}$ contains a cocircuit C^* of size less than q, then by Proposition 5.2, $\tilde{M} - C^*$ is colorable if and only if $\tilde{M}$ is colorable; replace

$\tilde{M}$ with $\tilde{M} - C^*$ and return to Step 2.1. If no such cocircuit exists, let $\hat{M} = \tilde{M}$ and proceed to Step 2.3.

<u>**2.3**</u> If $\hat{M}$ is empty, then M is colorable. If not, then no decision can be made about colorability of M. In either case, the algorithm stops.

This completes the description of the algorithm. $\square$

It should be noted that since every application of Proposition 5.1 c) or Proposition 5.2 consists in the deletion of a subset of elements of $\tilde{M}$, no matroid with a loop will ever occur during the second phase of the algorithm, and similarly no matroid occurs in this phase that would be subject to the reduction of Step 1.1.

Proposition 6.1 $\mathcal{A}[\mathcal{M}]$ *is a polynomial algorithm.*

Proof: It is straightforward to prove that the running time of $\mathcal{A}[\mathcal{M}]$ is bounded as a polynomial function of $nm \log q$, which is the natural measure of the complexity of representing the input matrix $\boldsymbol{A}$, with a leading term exponent crudely bounded by the maximum cardinality of the element sets of all basis matroids $M_i \in \mathcal{M}$.

Instead of giving all the details, the following is a sketch of a proof of the statement that each application of Step 1.1 can be performed in polynomial time. Since the number of times this part of the algorithm is carried out is at most the rank of the input matrix, this shows that the total contribution of Step 1.1 to the running time is polynomially bounded. What is presented is a crude argument which does not in any way pretend to indicate how Step 1.1 may be implemented efficiently.

For the remaining parts of the algorithm, the proofs are easy, and they will be omitted.

Assume that $\boldsymbol{A}$ is an $n \times m$ matrix and that $M_i \in \mathcal{M}$ has rank r and s elements, $n, m, r, s > 1$. The dimensions r and s of M_i are considered constants, and the size of $\boldsymbol{A}$ may be assumed equal to $nm \log q$.

Let $m' = s - 1$ and consider in turn each of the $\binom{m}{m'}$ (which is bounded by a polynomial of degree m' as a function of the size of $\boldsymbol{A}$) m'-subsets of the m-set of columns of $\boldsymbol{A}$; such a set naturally induces an $n \times m'$ matrix $\boldsymbol{A}'$. Let $\boldsymbol{A}'' = \boldsymbol{B}\boldsymbol{A}'$, where $\boldsymbol{B}$ is a suitable $n' \times n$ matrix of rank n' equal to the rank of $\boldsymbol{A}'$; this B can be found and $\boldsymbol{A}''$ can be calculated in a small degree polynomial number of steps. Then $\boldsymbol{A}'' \approx \boldsymbol{A}'$ holds. For $n' > r$, there is clearly no $\boldsymbol{e} \in \mathcal{W}_n$ with $M(\boldsymbol{A}'|\boldsymbol{e}) \simeq M_i$. Assuming $n' \leq r$, the search for such an $\boldsymbol{e}$ may proceed. Consider each vector $\boldsymbol{e}'$ in $\mathcal{W}_{n'} \setminus \{\boldsymbol{0}\}$ in turn; their number is bounded by the constant $q^r - 1$. For each of the $s!$ (a constant) bijections π from the column set of $\boldsymbol{A}''|\boldsymbol{e}'$ to the element set of M_i it suffices to check for each of the constant number 2^s of subsets X of the column set of $\boldsymbol{A}''|\boldsymbol{e}'$ whether X is linearly independent in $\boldsymbol{A}''|\boldsymbol{e}'$ if and only if $\pi(X)$ is independent in M_i (this is done in constant time since the number of rows of $\boldsymbol{A}''$ is bounded by the constant r and the size of X by the constant s, and since the size of M_i is constant). Whenever $M(\boldsymbol{A}''|\boldsymbol{e}') \simeq M_i$, then the desired $\boldsymbol{e} \in \mathcal{W}_n$ is chosen such that $\boldsymbol{e}' = \boldsymbol{B}\boldsymbol{e}$, which may easily be done in polynomial time. Hence the total number of steps is bounded by a polynomial in the size of $\boldsymbol{A}$ for each M_i. Since $\mathcal{M}$ is finite, the same is also true for each application of Step 1.1 by $\mathcal{A}[\mathcal{M}]$. $\square$

The simplest example from the described class of algorithms is $\mathcal{A}[\emptyset]$, which effectively only employs the second phase of the algorithm. Restricted to graphs (more precisely matroids $M(G)$ where G is a graph) $\mathcal{A}[\emptyset]$ essentially reduces to the classical procedure of coloring by *smallest last order*; see Section 1.3 of [7] for details and references. This statement is not quite accurate, since the typical version of the graph coloring algorithm only considers cocircuits (identifiable as edge-cuts of G) of $M(G)$ corresponding to sets of edges incident with the same vertex of G. The algorithm $\mathcal{A}[\emptyset]$ will naturally consider all cocircuits of $M(G)$ and will consequently be successful in certain instances for which the coloring by smallest last order will fail to color with the prescribed number q of colors.

A less trivial example of an algorithm is obtained when setting $\mathcal{M} = \{U^2_{q+1}\}$, called $\mathcal{A}[q, U^2_{q+1}]$ or $\mathcal{A}[U^2_{q+1}]$ for short. Because of the uniqueness of U^2_{q+1} as uncolorable loopless matroid of rank at most 2, this can be considered the unique simplest algorithm of the class having a nonempty basis.

The results of a few computational experiments concerning the performance of $\mathcal{A}[3, U^2_4]$ applied to 3-coloring of graphs are presented in a later section of this paper.

7 Consistency

It is not explicitly required that the sequence of steps performed by $\mathcal{A}[\mathcal{M}]$ should be uniquely determined by the input matrix $\boldsymbol{A}$; this is indeed a natural requirement to pose on any practical implementation, and we will assume without further discussion that it is met. However, it is clear that a different sequence of steps may be traced out if $\boldsymbol{A}$ is replaced by, say, an input matrix obtained by left-multiplication of

$\boldsymbol{A}$ with a non-singular square matrix, or by permutation of rows or columns, or by scaling some of the columns of $\boldsymbol{A}$ by nonzero factors. It is conceivable that replacement with such a different input matrix may change the output from $\mathcal{A}[\mathcal{M}]$ notwithstanding the fact that the matroids represented by the different matrices are in these cases isomorphic. This section deals with conditions on $\mathcal{M}$ to ensure dependence of the output from $\mathcal{A}[\mathcal{M}]$ solely on M, the matroid represented by the input matrix $\boldsymbol{A}$.

We will say that $\mathcal{A}[\mathcal{M}]$ *fails* for input $\boldsymbol{A}$ if it terminates without deciding whether $M(\boldsymbol{A})$ is colorable or not; otherwise it *succeeds* for input $\boldsymbol{A}$. $\mathcal{A}[\mathcal{M}]$ is called *consistent* if its success or failure for input $\boldsymbol{A}$ depends only on $M(\boldsymbol{A})$ up to isomorphism; $M(\boldsymbol{A}_1) \simeq M(\boldsymbol{A}_2)$ implies that $\mathcal{A}[\mathcal{M}]$ succeeds for input $\boldsymbol{A}_1$ if and only if it succeeds for input $\boldsymbol{A}_2$.

In the special case $q = 2$ one may observe

Theorem 7.1 *$\mathcal{A}[2, U_3^2]$ succeeds for every input.*

Proof: If $\boldsymbol{A}$ is any matrix over $\mathbb{F}_2$ with a pair of independent columns $\boldsymbol{e}$ and $\boldsymbol{f}$, then U_3^2 is clearly isomorphic to a submatroid of $M(\boldsymbol{A}|\boldsymbol{e}+\boldsymbol{f})$. It follows that Step 1.1 applies to give a reduction of $\boldsymbol{A}$. Thus the first phase of the algorithm $\mathcal{A}[2, U_3^2]$ always leaves a matroid $\tilde{M} = M(\tilde{\boldsymbol{A}})$ without pairwise independent elements. That is, either $\tilde{M}$ contains a loop, or all the columns of $\tilde{\boldsymbol{A}}$ are parallel. In the former case the algorithm terminates with the information that M is uncolorable. In the latter case, Step 2.1 leaves a matroid with only one element, which, since it is not a loop, is in fact a cocircuit of size 1 and thus is subsequently deleted in Step 2.2. The algorithm then terminates with the information that M is colorable. $\square$

It immediately follows that

Corollary 7.1 $\mathcal{A}[2, U_3^2]$ *is consistent.* $\square$

A theorem along the lines of Theorem 7.1 is fairly unlikely to be true for any $\mathcal{A}[q, \mathcal{M}]$ with $q > 2$. Indeed, for every $q > 2$ it is quite easy to find matroids M for which $\mathcal{A}[U_{q+1}^2]$ fails on M. In fact, since determining whether the chromatic number of a graph is at most q is an $\mathcal{NP}$-complete problem for $q > 2$, there exists for every finite $\mathcal{M}$ a colorable matroid M_1 such that $\mathcal{A}[\mathcal{M}]$ fails for at least one input matrix $\boldsymbol{A}$ for which $M_1 = M(\boldsymbol{A})$ (unless $\mathcal{P} = \mathcal{NP}$). Similarly, there must exist an uncolorable matroid M_2 such that $\mathcal{A}[\mathcal{M}]$ fails on every input $\boldsymbol{A}$ with $M_2 = M(\boldsymbol{A})$ (unless $\mathcal{NP} = \text{co-}\mathcal{NP}$).

In the remaining part of this section, we will consider the general problem of choosing $\mathcal{M}$ such that $\mathcal{A}[\mathcal{M}]$ is consistent. This problem is related to the following seemingly difficult question. *Given an algorithm $\mathcal{A}[\mathcal{M}]$ is it possible to characterize the class of matroids M for which $\mathcal{A}[\mathcal{M}]$ is successful on every input matrix $\boldsymbol{A}$ with $M = M(\boldsymbol{A})$?* For this class there is no need to be concerned that the output from the algorithm might possibly depend on the particular representation chosen as input; it would only depend on the matroid itself. If $\mathcal{A}[\mathcal{M}]$ is consistent, then the problem is easy to solve: The desired class of matroids contains precisely such M for which the algorithm succeeds on any arbitrarily chosen matrix $\boldsymbol{A}$ with $M = M(\boldsymbol{A})$. Since the algorithm is polynomial, it decides for any given M (providing M has a known representation over $\mathbb{F}_q$) in polynomial time whether M belongs to the class or not.

As a first step towards ensuring consistency of $\mathcal{A}[\mathcal{M}]$ it will be natural to require that whenever $M_i \in \mathcal{M}$ in Step 1.1 is identified as a matroid represented by a fixed subset of the columns in $\boldsymbol{A}|\boldsymbol{e}$, then the matroid of the matrix $\boldsymbol{A}'$ resulting from the reduction step should be uniquely determined. This condition is captured by the following formal definition.

A matroid $M \simeq M(\boldsymbol{A}|\boldsymbol{e})$, where $\boldsymbol{A} \in {\mathbb{F}_q}^{n\times m}$ and $\boldsymbol{e} \in \mathcal{W}_n$, is *recoverable* (over $\mathbb{F}_q$) if for every $\boldsymbol{A}'$ satisfying $\boldsymbol{A} \approx \boldsymbol{A}'$, there exists $\boldsymbol{e}' \in \mathcal{W}_n$ such that $\boldsymbol{A}|\boldsymbol{A}_1 \approx \boldsymbol{A}'|\boldsymbol{A}_1'$ implies $\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e} \approx \boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}'$ for all matrices $\boldsymbol{A}_1, \boldsymbol{A}_1'$ on n rows.

Low connectivity may cause a lack of recoverability. In the extreme case, when every subset of the element set of M is independent, and the components of M are singletons, it is easy to check that M is not recoverable over any field. One can easily check that U_3^2 is recoverable over $\mathbb{F}_2$, but not over any larger field. Thus the property also depends on the field of representation. It is not clear in general how to establish whether a given matroid is recoverable, but in some important cases this can be done, as we will see.

The following technical lemma is a first application of the recoverability concept.

Lemma 7.1 *Let $M = M(\boldsymbol{A})$, where $\boldsymbol{A}$ is an $n \times m$ matrix, and let $M' = M(\boldsymbol{A}')$ where $\boldsymbol{A}'$ is an $n' \times m'$ matrix. Assume that $M' \simeq M(\boldsymbol{A}|\boldsymbol{A}_0)/\boldsymbol{A}_0$ for some matrix $\boldsymbol{A}_0$. If $M(\boldsymbol{A}'|\boldsymbol{e}')$ has a recoverable submatroid, for some $\boldsymbol{e}' \in \mathcal{W}_{n'}$, which contains $\boldsymbol{e}'$ as one of its elements,* *then there exists* $\boldsymbol{e} \in \mathcal{W}_n$ *such that*

$$M(\boldsymbol{A}'|\boldsymbol{e}') \simeq M(\boldsymbol{A}|\boldsymbol{A}_0|\boldsymbol{e})/\boldsymbol{A}_0,$$

where the isomorphism can be chosen to map $\boldsymbol{e}'$ *to* $\boldsymbol{e}$*.*

Proof: The proof is by induction on $r(M) - r(M') \geq 0$.

First, if $r(M) = r(M')$ then every entry of $\boldsymbol{A}_0$ is a zero and $M \simeq M'$ follows. It may now be assumed without loss of generality that $\boldsymbol{A}' \approx \boldsymbol{A}$. If $M(\boldsymbol{A}'|\boldsymbol{e}')$ has a recoverable submatroid containing $\boldsymbol{e}'$, we may assume $\boldsymbol{A}'|\boldsymbol{e}' \approx \boldsymbol{A}'_1|\boldsymbol{A}'_2|\boldsymbol{e}'$ where $M(\boldsymbol{A}'_1|\boldsymbol{e}')$ is recoverable. Hence there exists $\boldsymbol{e} \in \mathcal{W}_n$ such that

$$\boldsymbol{A}'|\boldsymbol{e}' \approx \boldsymbol{A}'_1|\boldsymbol{A}'_2|\boldsymbol{e}' \approx \boldsymbol{A}|\boldsymbol{e},$$

and clearly $M(\boldsymbol{A}'|\boldsymbol{e}')$ allows the desired isomorphism to $M(\boldsymbol{A}|\boldsymbol{e})$.

For $k > 0$ assume that the statement of the lemma is true whenever $r(M) - r(M') < k$, and let M and M' satisfy $r(M) - r(M') = k$. There exists at least one column $\boldsymbol{f}$ of $\boldsymbol{A}_0$ such that $\boldsymbol{f} \neq \boldsymbol{0}$; we may assume $\boldsymbol{A}_0 = \boldsymbol{A}'_0|\boldsymbol{f}$. Defining $M'' = M(\boldsymbol{A}|\boldsymbol{f})/\boldsymbol{f}$ it follows that $r(M'') = r(M) - 1$. Moreover M'' may be represented by a matrix $\boldsymbol{A}''$ obtained by multiplying $\boldsymbol{A}|\boldsymbol{f}$ from the left with a non-singular matrix $\boldsymbol{B}$ such that $\boldsymbol{B}\boldsymbol{f} = \boldsymbol{e}_n$, the n'th standard basis vector, followed by deleting from $\boldsymbol{B}\boldsymbol{A}|\boldsymbol{B}\boldsymbol{f}$ the n'th row and the column $\boldsymbol{B}\boldsymbol{f}$.

Now M' satisfies

$$M' \simeq M(\boldsymbol{A}|\boldsymbol{A}'_0|\boldsymbol{f})/\boldsymbol{A}'_0/\boldsymbol{f} \simeq M(\boldsymbol{B}\boldsymbol{A}|\boldsymbol{B}\boldsymbol{A}'_0|\boldsymbol{B}\boldsymbol{f})/\boldsymbol{B}\boldsymbol{A}'_0/\boldsymbol{B}\boldsymbol{f} \simeq M(\boldsymbol{A}''_1|\boldsymbol{A}''_0)/\boldsymbol{A}''_0, \tag{1}$$

where $\boldsymbol{A}''_1$ and $\boldsymbol{A}''_0$ denote the matrices obtained from $\boldsymbol{B}\boldsymbol{A}$ and $\boldsymbol{B}\boldsymbol{A}'_0$ respectively by removing their nth rows.

By the induction hypothesis applied to $M(\boldsymbol{A}_1'')$ and M' with $\boldsymbol{A}_0''$ in place of $\boldsymbol{A}_0$, where $r(M(\boldsymbol{A}_1'')) - r(M') = k-1$, there exists $\boldsymbol{e}'' \in \mathcal{W}_{n-1}$ such that

$$M(\boldsymbol{A}'|\boldsymbol{e}') \simeq M(\boldsymbol{A}_1''|\boldsymbol{A}_0''|\boldsymbol{e}'')/\boldsymbol{A}_0''. \tag{2}$$

Let $\boldsymbol{e} \in \mathcal{W}_n$ be any vector such that $\boldsymbol{Be}$ agrees with $\boldsymbol{e}''$ in its first $n-1$ coordinates. Then $M(\boldsymbol{A}|\boldsymbol{A}_0|\boldsymbol{e})/\boldsymbol{f}$ is isomorphic to $M(\boldsymbol{A}_1''|\boldsymbol{A}_0''|\boldsymbol{e}'')$, and analogously to (1), it follows that

$$M(\boldsymbol{A}|\boldsymbol{A}_0|\boldsymbol{e})/\boldsymbol{A}_0 \simeq M(\boldsymbol{A}_1''|\boldsymbol{A}_0''|\boldsymbol{e}'')/\boldsymbol{A}_0''. \tag{3}$$

Combining (2) with (3), and observing that the two isomorphisms can each be chosen with the desired property, the result follows by induction. □

We next want to ensure that if a reduction can be potentially performed in Step 2.1, but it is not, then a potential reduction is inherited for the algorithm to carry out at a later stage.

A set $\mathcal{M}$ of matroids is *A/C-closed* if, for every $M \in \mathcal{M}$, every $\boldsymbol{A}$ with $M \simeq M(\boldsymbol{A})$ and every $\boldsymbol{A}'$ having the same number of rows as $\boldsymbol{A}$ and such that $M(\boldsymbol{A}|\boldsymbol{A}')/\boldsymbol{A}'$ is loopless, there exists $M_0 \in \mathcal{M}$ such that $M_0 \subseteq M(\boldsymbol{A}|\boldsymbol{A}')/\boldsymbol{A}'$. (Here A/C stands for 'Addition/Contraction'). Note that Lemma 5.3 implies that $M(\boldsymbol{A}|\boldsymbol{A}')/\boldsymbol{A}'$ cannot be colorable unless $M(\boldsymbol{A})$ is colorable. Hence if $\mathcal{M}$ is not already A/C-closed then it is easy to modify $\mathcal{M}$ by adding new loopless matroids of the form $M(\boldsymbol{A}|\boldsymbol{A}')/\boldsymbol{A}'$ where $M \simeq M(\boldsymbol{A})$ and $M \in \mathcal{M}$ until the resulting new basis is A/C-closed. Thus the property is not difficult to achieve.

Now the main theorem can be stated.

Theorem 7.2 *If $\mathcal{M}$ is an A/C-closed set of recoverable matroids over $\mathbb{F}_q$, then $\mathcal{A}[q,\mathcal{M}]$ is consistent.*

Proof: Let any matrix $\boldsymbol{A} \in \mathbb{F}_q{}^{n\times m}$ be input to $\mathcal{A}[q,\mathcal{M}]$, where it may be assumed that $\boldsymbol{A}$ is nonempty. We will show that either the first phase of $\mathcal{A}[q,\mathcal{M}]$ decides that $M = M(\boldsymbol{A})$ is not colorable, or the matroid $\tilde{M}$ with which the first phase finishes depends only on M, not on the representation $\boldsymbol{A}$. Since each step of the second phase of $\mathcal{A}[q,\mathcal{M}]$ consists in the deletion of either one of a pair of elements that form a 2-circuit, or the deletion of an entire cocircuit, it is not difficult to see that the matroid $\hat{M}$ with which $\mathcal{A}[q,\mathcal{M}]$ terminates in Step 2.3 is uniquely determined by the isomorphism class of $\tilde{M}$. The proof of the theorem then follows.

The execution of the first phase of $\mathcal{A}[q,\mathcal{M}]$ is not affected by initially removing from $\boldsymbol{A}$ one from each pair of parallel columns, and repeating this until no such pair remains. Thus it may be assumed that $\boldsymbol{A}$ has no pair of parallel columns.

Define sequences of sets $\overline{F_1}, \overline{F_2}, \ldots \subseteq \mathcal{W}_n$, $F_1, F_2, \ldots \subseteq \mathcal{W}_n$, and of matrices $\overline{\boldsymbol{A}_1}, \overline{\boldsymbol{A}_2}, \ldots$ as follows. For $i = 1, 2, \ldots$ let

$$\overline{F_i} = \bigcup_{j<i} F_j,$$

let $\overline{\boldsymbol{A}_i}$ be any matrix whose columns (in some arbitrary but fixed order) are precisely the elements of $\overline{F_i}$. Finally let

$$F_i = \{\boldsymbol{e} \in \mathcal{W}_n \setminus \overline{F_i} \mid \exists M_0 \in \mathcal{M} \;:\; M_0 \subseteq M(\boldsymbol{A}|\boldsymbol{e}|\overline{\boldsymbol{A}_i})/\overline{\boldsymbol{A}_i}\}$$

The sets F_i are clearly disjoint for all i, and the sets $\overline{F_i}$ and the matrices $\overline{\boldsymbol{A}_i}$ eventually coincide with their limits $\overline{F_\infty}$ and $\overline{\boldsymbol{A}_\infty}$ respectively, as i becomes large.

Suppose that $\boldsymbol{A}'$ satisfies $M(\boldsymbol{A}') \simeq M(\boldsymbol{A})$, and let F_i', $\overline{F_i}'$ and $\overline{\boldsymbol{A}_i}'$, $i > 0$, be sequences defined similarly with respect to $\boldsymbol{A}'$ instead of $\boldsymbol{A}$.

Claim 1: If $M(\boldsymbol{A}|\overline{\boldsymbol{A}_\infty})/\overline{\boldsymbol{A}_\infty}$ is loopless, then for every $i > 0$

$$M(\boldsymbol{A}|\overline{\boldsymbol{A}_i})/\overline{\boldsymbol{A}_i} \simeq M(\boldsymbol{A}'|\overline{\boldsymbol{A}_i}')/\overline{\boldsymbol{A}_i}'.$$

The claim is proved by induction on i. For $i = 1$ it is true, since $\overline{\boldsymbol{A}_1} = \overline{\boldsymbol{A}_1}' = \emptyset$ and $M(\boldsymbol{A}) \simeq M(\boldsymbol{A}')$ by assumption.

Assume $M(\boldsymbol{A}|\overline{\boldsymbol{A}_\infty})/\overline{\boldsymbol{A}_\infty}$ is loopless, $i > 1$ and $M(\boldsymbol{A}|\overline{\boldsymbol{A}_{i-1}})/\overline{\boldsymbol{A}_{i-1}} \simeq M(\boldsymbol{A}'|\overline{\boldsymbol{A}_{i-1}}')/\overline{\boldsymbol{A}_{i-1}}'$. Suppose the claim fails to be true for i. Then there exist partitionings $\overline{\boldsymbol{A}_i} = \boldsymbol{A}_1|\boldsymbol{A}_2$ and $\overline{\boldsymbol{A}_i}' = \boldsymbol{A}_1'|\boldsymbol{A}_2'$ (possibly after reordering of columns) such that

$$M(\boldsymbol{A}|\boldsymbol{A}_1)/\boldsymbol{A}_1 \simeq M(\boldsymbol{A}'|\boldsymbol{A}_1')/\boldsymbol{A}_1'$$

but

$$M(\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e})/\boldsymbol{A}_1/\boldsymbol{e} \not\simeq M(\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}')/\boldsymbol{A}_1'/\boldsymbol{e}'$$

for all columns $\boldsymbol{e}$ of $\boldsymbol{A}_2$ and $\boldsymbol{e}'$ of $\boldsymbol{A}_2'$. By induction hypothesis it can be assumed that the leftmost columns of $\boldsymbol{A}_1$ (resp. $\boldsymbol{A}_1'$) are precisely those of $\overline{\boldsymbol{A}_{i-1}}$ (resp. $\overline{\boldsymbol{A}_{i-1}}'$). It may be assumed without loss of generality that $\boldsymbol{A}_2$ has at least as many columns as $\boldsymbol{A}_2'$, in particular $\boldsymbol{A}_2$ is not empty.

Let $\boldsymbol{e}$ be any column of $\boldsymbol{A}_2$. Then $\boldsymbol{e} \in F_{i-1}$, since $\boldsymbol{e}$ is a column of $\overline{\boldsymbol{A}_i}$ but not $\overline{\boldsymbol{A}_{i-1}}$. By definition of F_{i-1} there exists $M_0 \in \mathcal{M}$ with $M_0 \subseteq M(\boldsymbol{A}|\overline{\boldsymbol{A}_{i-1}}|\boldsymbol{e})/\overline{\boldsymbol{A}_{i-1}}$. Then $\boldsymbol{e}$ is an element of this submatroid, since $M_0 \subseteq M(\boldsymbol{A}|\overline{\boldsymbol{A}_{i-1}})/\overline{\boldsymbol{A}_{i-1}}$ would imply $\overline{F_i} = \mathcal{W}_n$ and $M(\boldsymbol{A}|\overline{\boldsymbol{A}_\infty})/\overline{\boldsymbol{A}_\infty}$ would consist of m loops, contrary to the assumption that it is loopless.

By Lemma 7.1 there exists $\boldsymbol{e}'$ such that

$$M(\boldsymbol{A}|\overline{\boldsymbol{A}_{i-1}}|\boldsymbol{e})/\overline{\boldsymbol{A}_{i-1}} \simeq M(\boldsymbol{A}'|\overline{\boldsymbol{A}_{i-1}}'|\boldsymbol{e}')/\overline{\boldsymbol{A}_{i-1}}',$$

where the isomorphism may be chosen such that it maps $\boldsymbol{e}$ to $\boldsymbol{e}'$. Hence M_0 is isomorphic to a submatroid containing $\boldsymbol{e}'$ of the latter matroid. Moreover, we have

$$M(\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e})/\boldsymbol{A}_1 \simeq M(\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}')/\boldsymbol{A}_1',$$

and it is easy to decude that $M(\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}')/\boldsymbol{A}_1'$ is loopless, since any loop of $M(\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e})/\boldsymbol{A}_1$ would be a loop of $M(\boldsymbol{A}|\overline{\boldsymbol{A}_\infty})/\overline{\boldsymbol{A}_\infty}$. Since $\mathcal{M}$ is A/C-closed, there exists $M_0' \in \mathcal{M}$ such that $M_0' \subseteq M(\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}')/\boldsymbol{A}_1'$.

It remains to show that $\boldsymbol{e}'$ is a column of $\boldsymbol{A}_2'$; this would contradict the choice of $\boldsymbol{A}_1$ and $\boldsymbol{A}_1'$, since $M(\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e})/\boldsymbol{A}_1/\boldsymbol{e} \simeq M(\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}')/\boldsymbol{A}_1'/\boldsymbol{e}'$ holds, and thus the proof of the claim would follow.

Since $M_0 \subseteq M(\boldsymbol{A}'|\overline{\boldsymbol{A}_{i-1}}'|\boldsymbol{e}')/\overline{\boldsymbol{A}_{i-1}}'$, it is clear that $\boldsymbol{e}' \in \overline{F_i}'$, so that $\boldsymbol{e}'$ is indeed a column of $\overline{\boldsymbol{A}_i}'$. Suppose $\boldsymbol{e}'$ is a column of $\boldsymbol{A}_1'$. Then $\boldsymbol{e}'$ is a loop of $M(\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}')/\boldsymbol{A}_1'$ which as before leads to contradiction. Thus there only remains the possibility that $\boldsymbol{e}'$ is a column of $\boldsymbol{A}_2'$ and the proof of Claim 1 follows by induction.

<u>Claim 2</u>: $\tilde{M} \simeq M(\boldsymbol{A}|\overline{\boldsymbol{A}_\infty})/\overline{\boldsymbol{A}_\infty}$.

Using that the elements of $\mathcal{M}$ are recoverable combined with Lemma 7.1, Phase 1 of $\mathcal{A}[q,\mathcal{M}]$ produces a matroid $\tilde{M} \simeq M(\boldsymbol{A}|\tilde{\boldsymbol{A}})/\tilde{\boldsymbol{A}}$ for some $\tilde{\boldsymbol{A}}$, where it may be assumed that $\tilde{M}$ has no loops. It follows by a straighforward inductive argument that each column of $\tilde{\boldsymbol{A}}$ belongs to $\overline{F_\infty}$. By adding, if necessary, more columns to $\tilde{\boldsymbol{A}}$, it can be assumed that if any vector in $\mathcal{W}_n$ is parallel to some column of $\tilde{\boldsymbol{A}}$, then it is itself a column of $\tilde{\boldsymbol{A}}$. The addition of such a new column does not change $M(\boldsymbol{A}|\tilde{\boldsymbol{A}})/\tilde{\boldsymbol{A}}$, and since it is easy to see that each F_i consists of complete parallel classes of elements of $\mathcal{W}_n$, it does not change that every column of $\tilde{\boldsymbol{A}}$ belongs to $\overline{F_\infty}$. By a similar argument, now applied to removing redundant columns, the columns of $\tilde{\boldsymbol{A}}$ may be assumed distinct elements of $\mathcal{W}_n$.

Suppose that some element of $\overline{F_\infty}$ is not parallel to any column of $\tilde{\boldsymbol{A}}$, and choose $\boldsymbol{e} \in F_i$ with i as small as possible such that $\boldsymbol{e}$ has this property. By the choice of $\boldsymbol{e}$ there exists $M_0 \in \mathcal{M}$ such that $M_0 \subseteq M(\boldsymbol{A}|\boldsymbol{e}|\overline{\boldsymbol{A}_i})/\overline{\boldsymbol{A}_i}$. Since the choice of i implies that every element of $\overline{F_i}$ is parallel to some column of $\tilde{\boldsymbol{A}}$, there exists $\boldsymbol{A}'$ such that (w.l.o.g.) $\tilde{\boldsymbol{A}} \approx \overline{\boldsymbol{A}_i}|\boldsymbol{A}'$.

Contracting first the columns of $\overline{\boldsymbol{A}_i}$ from $M(\boldsymbol{A}|\boldsymbol{e}|\overline{\boldsymbol{A}_i}|\boldsymbol{A}')/(\overline{\boldsymbol{A}_i}|\boldsymbol{A}')$, thereby calculating a specific representation matrix for $M(\boldsymbol{A}|\boldsymbol{e}|\overline{\boldsymbol{A}_i}|\boldsymbol{A}')/\overline{\boldsymbol{A}_i}$, it is clear that there exist $\boldsymbol{A}_0$, $\boldsymbol{A}_0'$ and $\boldsymbol{A}_0''$ such that

$$\begin{aligned}
&M_0 \simeq M(\boldsymbol{A}_0),\\
&M(\boldsymbol{A}|\boldsymbol{e}|\overline{\boldsymbol{A}_i})/\overline{\boldsymbol{A}_i} \simeq M(\boldsymbol{A}_0|\boldsymbol{A}_0') \text{ and}\\
&M(\boldsymbol{A}|\boldsymbol{e}|\overline{\boldsymbol{A}_i}|\boldsymbol{A}')/(\overline{\boldsymbol{A}_i}|\boldsymbol{A}') \simeq M(\boldsymbol{A}_0|\boldsymbol{A}_0'|\boldsymbol{A}_0'')/\boldsymbol{A}_0''.
\end{aligned}$$

It follows that $M(\boldsymbol{A}|\boldsymbol{e}|\tilde{\boldsymbol{A}})/\tilde{\boldsymbol{A}}$ contains $M(\boldsymbol{A}_0|\boldsymbol{A}_0'')/\boldsymbol{A}_0''$ as a submatroid. In particular the latter matroid is loopless. Since $\mathcal{M}$ is A/C-closed, there exists $M' \in \mathcal{M}$ with $M' \subseteq M(\boldsymbol{A}|\boldsymbol{e}|\tilde{\boldsymbol{A}})/\tilde{\boldsymbol{A}}$. This contradicts the definition of $\tilde{M}$. Hence every $\boldsymbol{e} \in \overline{F_\infty}$ is parallel to some column of $\tilde{\boldsymbol{A}}$, and in fact the columns of $\tilde{\boldsymbol{A}}$ are precisely the columns of $\overline{\boldsymbol{A}_\infty}$. This implies $M(\boldsymbol{A}|\tilde{\boldsymbol{A}})/\tilde{\boldsymbol{A}} \simeq M(\boldsymbol{A}|\overline{\boldsymbol{A}_\infty})/\overline{\boldsymbol{A}_\infty}$ as claimed.

Claim 1 and 2 together imply that $\tilde{M}$ depends only on M, which concludes the proof of the theorem. $\square$

Theorem 7.2 can be applied to show the consistency of $\mathcal{A}[q, U^2_{q+1}]$. First an easy lemma is needed.

Lemma 7.2 *Let $\boldsymbol{A}' \in \mathbb{F}_q{}^{n\times q}$ be a matrix satisfying $M(\boldsymbol{A}') \simeq U^2_q$.*

Then there exist non-singular matrices $\boldsymbol{B} \in \mathbb{F}_q{}^{2\times n}$ and $\boldsymbol{P}, \boldsymbol{D} \in \mathbb{F}_q{}^{q\times q}$, where $\boldsymbol{P}$ is a permutation matrix and $\boldsymbol{D}$ is diagonal, such that

$$\boldsymbol{B}\boldsymbol{A}'\boldsymbol{P}\boldsymbol{D} = \begin{bmatrix} 1 & 0 & 1 & 1 & \dots & 1 \\ 0 & 1 & 1 & \alpha_2 & \dots & \alpha_{q-2} \end{bmatrix}$$

(where the elements of $\mathbb{F}_q$ are $0, 1 = \alpha_1, \alpha_2, \dots, \alpha_{q-1}$).

Proof: This follows easily from the fact that the rank of $\boldsymbol{A}'$ is equal to 2, and that no two of its q columns are parallel. $\square$

Theorem 7.3 *$\mathcal{A}[q, U^2_{q+1}]$ is consistent.*

Proof: Suppose $\boldsymbol{A}_0 \in \mathbb{F}_q{}^{n\times(q+1)}$ satisfies $M(\boldsymbol{A}_0) \simeq U^2_{q+1}$. It is easy to see that either $M(\boldsymbol{A}_0|\boldsymbol{e})/\boldsymbol{e} \simeq U^2_{q+1}$ (if $\boldsymbol{e} = \boldsymbol{0}$, or if $\boldsymbol{e}$ is not a linear combination of columns of $\boldsymbol{A}_0$), or $M(\boldsymbol{A}_0|\boldsymbol{e})/\boldsymbol{e}$ contains a loop. Thus $\{U^2_{q+1}\}$ is A/C-closed.

For fixed $\boldsymbol{e} \in M(\boldsymbol{A}_0)$ we may assume without loss of generality that $\boldsymbol{A}_0 = \boldsymbol{A}|\boldsymbol{e}$ for some $\boldsymbol{A}$. Clearly $M(\boldsymbol{A}) \simeq U^2_q$. Let $\boldsymbol{A}' \in \mathbb{F}_q{}^{n\times q}$ be any matrix satisfying $M(\boldsymbol{A}') \simeq U^2_q$. Then there exists, by an easy calculation using Lemma 7.2, a vector $\boldsymbol{e}'$ such that $M(\boldsymbol{A}'|\boldsymbol{e}') \simeq U^2_{q+1}$ and hence such that $\boldsymbol{A}|\boldsymbol{e} \approx \boldsymbol{A}'|\boldsymbol{e}'$.

Suppose there exist matrices $\boldsymbol{A}_1, \boldsymbol{A}'_1$, each with n rows, such that $\boldsymbol{A}|\boldsymbol{A}_1 \approx \boldsymbol{A}'|\boldsymbol{A}'_1$ but $\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e} \not\approx \boldsymbol{A}'|\boldsymbol{A}'_1|\boldsymbol{e}'$. In particular, without loss of generality, there exists a circuit X of $M(\boldsymbol{A}|\boldsymbol{A}_1|\boldsymbol{e})$ such that the corresponding columns X' of $\boldsymbol{A}'|\boldsymbol{A}_1'|\boldsymbol{e}'$ are linearly independent. Clearly $\boldsymbol{e} \in X$ and $\boldsymbol{e}' \in X'$ hold. There is at most one column of $\boldsymbol{A}$ in X, since otherwise X would have size 3 and be a circuit of $M(\boldsymbol{A}|\boldsymbol{e})$, so that X' would be a subset of size 3 in $M(\boldsymbol{A}'|\boldsymbol{e}') \simeq U^2_{q+1}$ and could not be independent.

If X contains precisely one column $\boldsymbol{a}_1$ from $\boldsymbol{A}$, then let $\boldsymbol{a}_2$ be a column of $\boldsymbol{A}$ different from $\boldsymbol{a}_1$. It is clear that $X \setminus \{\boldsymbol{e}\} \cup \{\boldsymbol{a}_1, \boldsymbol{a}_2\}$ is a linearly dependent set, and the corresponding set $X' \setminus \{\boldsymbol{e}'\} \cup \{\boldsymbol{a}_1', \boldsymbol{a}_2'\}$ is therefore also linearly dependent. Since $\boldsymbol{a}_2'$ can be written as a linear combination of $\boldsymbol{a}_1'$ and $\boldsymbol{e}'$, and X' is linearly independent, this implies $X' = \{\boldsymbol{a}_1', \boldsymbol{e}'\}$. But then $X = \{\boldsymbol{a}_1, \boldsymbol{e}\}$ contradicts that X is linearly dependent.

Assume next that X contains no column of $\boldsymbol{A}$. Again $Y = X \setminus \{\boldsymbol{e}\} \cup \{\boldsymbol{a}_1, \boldsymbol{a}_2\}$ is a linearly dependent set for any two distinct columns $\boldsymbol{a}_1, \boldsymbol{a}_2$ of $\boldsymbol{A}$, and similarly for the corresponding set $Y' = X' \setminus \{\boldsymbol{e}'\} \cup \{\boldsymbol{a}_1', \boldsymbol{a}_2'\}$ of columns of $\boldsymbol{A}'|\boldsymbol{A}_1'$. There exists a nontrivial linear combination of the elements of Y' whose value is zero, and by the linear independence of X' the coefficients of $\boldsymbol{a}_1'$ and $\boldsymbol{a}_2'$ are not both zeroes. But then there is a linear combination of elements of $Y' \setminus \{\boldsymbol{a}_1', \boldsymbol{a}_2'\} = X' \setminus \{\boldsymbol{e}'\}$ parallel to a column $\boldsymbol{z}'$ of $\boldsymbol{A}'|\boldsymbol{e}'$. Since X' is linearly independent, $\boldsymbol{z}'$ is not parallel to $\boldsymbol{e}'$, hence $\boldsymbol{z}'$ may be assumed a column of $\boldsymbol{A}'$. Now $X \setminus \{\boldsymbol{e}\} \cup \{\boldsymbol{z}\}$ is linearly dependent, where $\boldsymbol{z}$ is the column of $\boldsymbol{A}$ that corresponds to $\boldsymbol{z}'$. This implies that $\boldsymbol{z}$ may be written as a linear combination of $X \setminus \{\boldsymbol{e}\}$. Any column of $\boldsymbol{A}$ may be written as a linear combination of $\boldsymbol{e}$ and $\boldsymbol{z}$, therefore it may also be written as a linear combination of $X \setminus \{\boldsymbol{e}\}$, and consequently any column of $\boldsymbol{A}'$ may be written as a linear combination of $X' \setminus \{\boldsymbol{e}'\}$. But $\boldsymbol{e}'$ is a linear combination of any two such columns, contradicting the linear independence of X'. Therefore we can conclude that $\boldsymbol{A}|\boldsymbol{A}_1\boldsymbol{e} \approx \boldsymbol{A}'|\boldsymbol{A}'_1|\boldsymbol{e}'$ for all $\boldsymbol{A}_1, \boldsymbol{A}'_1$ such that $\boldsymbol{A}|\boldsymbol{A}_1 \approx \boldsymbol{A}'|\boldsymbol{A}'_1$, and hence that U^2_{q+1} is recoverable.

The proof follows by Theorem 7.2. $\square$

It is worth noting that among the algorithms of the form $\mathcal{A}[q, U^2_{q+1}]$ only $\mathcal{A}[2, U^2_3]$ and $\mathcal{A}[3, U^2_4]$ are directly applicable to graphs. More precisely, any algorithm $\mathcal{A}[q, U^2_{q+1}]$ is equivalent to $\mathcal{A}[q, \emptyset]$ for $q > 3$ when restricted to graphs. This is since deleting an element from U^2_{q+1} yields the matroid U^2_q which contains U^2_4 as a submatroid for $q > 3$. It is well-known and easy to see that U^2_4 has no representation over $\mathbb{F}_2$, whereas $M(G)$ can clearly be represented over $\mathbb{F}_2$ (by $\boldsymbol{A}_G$). This observation implies that Step 1.1 of the algorithm will be fruitless when applied to graphs for $q > 3$.

It is clear by the same reasoning that $\mathcal{A}[3, U^2_4]$ fails for every 3-edge-connected graph containing no triangles (i.e. cycles of length 3), since no reduction in Phase 1 or 2 is possibles for such a graph. Thus the class of 3-edge-connected trianglefree graphs is an easily described class of graphs for which this particular algorithm fails.

Any consistent algorithm $\mathcal{A}[q, \mathcal{M}]$ partitions the set of all graphs into three polynomially recognizable classes:

1. q-colorable graphs for which $\mathcal{A}[q, \mathcal{M}]$ succeeds,
2. non-q-colorable graphs for which $\mathcal{A}[q, \mathcal{M}]$ succeeds, and
3. graphs for which $\mathcal{A}[q, \mathcal{M}]$ fails.

It would seem quite interesting to obtain a purely structural graph-theoretical characterization of either class, even in the simple special case $\mathcal{A}[3, U_4^2]$. There are many open problems in graph theory that deal with the structure of highly chromatic graphs. If it is possible to choose $\mathcal{M}$ such that the second class contains a very large fraction of all non-q-colorable graphs, while at the same time allowing a manageable structural characterizaton of each class, then many such problems should reduce to the, by assumption easily recognizable, graphs of the relatively small third class.

8 3-Coloring

We will now restrict attention to the special case $q = 3$, that is, to *ternary* matroids. It is a particularly interesting special case, since as already noted, 3-colorability of graphs remains an $\mathcal{NP}$-complete problem, whereas it is clear (e.g. from Theorem 7.1) that 2-colorability of a graph is polynomially decidable. Moreover, ternary matroids are in a sense better understood than the matroids representable over finite fields of larger cardinality. Accordingly, one can say a little more about the behavior of algorithms of the form $\mathcal{A}[3, \mathcal{M}]$. The following noteworthy theorem characterizes completely the set of matrices over $\mathbb{F}_3$ which may represent isomorphic matroids. As observed in [3], it has no analogue for larger finite fields.

Theorem 8.1 (Brylawski & Lucas 1976, [3]) *Let $\boldsymbol{A}_1, \boldsymbol{A}_2 \in \mathbb{F}_3{}^{n\times m}$ be two matrices with $\boldsymbol{A}_1 \approx \boldsymbol{A}_2$, Then there exists a non-singular $n \times n$ matrix $\boldsymbol{B}$ and a non-singular diagonal $m \times m$ matrix $\boldsymbol{D}$ such that $\boldsymbol{A}_1 = \boldsymbol{B}\boldsymbol{A}_2\boldsymbol{D}$.* □

The following consequence of Theorems 7.2 and 8.1 states that the consistency of an algorithm over $\mathbb{F}_3$ is ensured under fairly mild conditions, including an easily checked connectivity condition for each basis matroid.

Theorem 8.2 *Let $\mathcal{M}$ be an A/C-closed set of ternary matroids without loops and satisfying the condition that $M - \boldsymbol{e}$ is connected for all $\boldsymbol{e} \in M$ and $M \in \mathcal{M}$. Then $\mathcal{A}[3, \mathcal{M}]$ is consistent.* □

Proof: By Theorem 7.2 it suffices to show that a ternary matroid M without loops is recoverable if it satisfies the above connectivity condition.

For any $\boldsymbol{e} \in M$ assume that $M - \boldsymbol{e} \simeq M(\boldsymbol{A}_1) \simeq M(\boldsymbol{A}_2)$, and assume without loss of generality $\boldsymbol{A}_1 \approx \boldsymbol{A}_2$. Choose $\boldsymbol{B}$ and $\boldsymbol{D}$ as in Theorem 8.1 such that $\boldsymbol{A}_1 = \boldsymbol{B}\boldsymbol{A}_2\boldsymbol{D}$. First we show that the choice of $\boldsymbol{B}$ and $\boldsymbol{D}$ is unique up to a sign: whenever $\boldsymbol{A}_1 = \boldsymbol{B}'\boldsymbol{A}_2\boldsymbol{D}'$ holds for an $n \times n$ non-singular matrix $\boldsymbol{B}'$ and an $m \times m$ nonsingular diagonal matrix $\boldsymbol{D}'$, then $(\boldsymbol{B}', \boldsymbol{D}') \in \{(\boldsymbol{B}, \boldsymbol{D}), (-\boldsymbol{B}, -\boldsymbol{D})\}$.

For the uniqueness proof we may assume, possibly after left-multiplication of $\boldsymbol{A}_2$ by a suitable non-singular matrix, that $\boldsymbol{A}_2$ contains n columns which together form an $n \times n$ identity matrix $\boldsymbol{I}_n$. Suppose that $\boldsymbol{A}_1 = \boldsymbol{B}\boldsymbol{A}_2\boldsymbol{D} = \boldsymbol{B}'\boldsymbol{A}_2\boldsymbol{D}'$. Then $\boldsymbol{A}_2 = \boldsymbol{B}^{-1}\boldsymbol{B}'\boldsymbol{A}_2\boldsymbol{D}'\boldsymbol{D}^{-1} = \boldsymbol{B}''\boldsymbol{A}_2\boldsymbol{D}''$ where $\boldsymbol{B}''$ and $\boldsymbol{D}''$ are again non-singular $n \times n$ and $m \times m$ matrices, respectively, and $\boldsymbol{D}''$ is diagonal. Since $\boldsymbol{B}''\boldsymbol{I}_n\boldsymbol{D}'' = \boldsymbol{I}_n$, it is clear that $\boldsymbol{B}''$ is a diagonal matrix. Any column $\boldsymbol{f}$ of $\boldsymbol{A}_2$ contains at least one nonzero entry since M has no loop. Hence $\boldsymbol{B}''\boldsymbol{f} = \alpha^{-1}\boldsymbol{f}$, where α denotes the diagonal entry corresponding to $\boldsymbol{f}$ of $\boldsymbol{D}''$, and it follows that all diagonal entries of $\boldsymbol{B}''$ corresponding to the nonzero entries of $\boldsymbol{f}$ have the value $\alpha^{-1} = \alpha$. If $\boldsymbol{e}_i$ and $\boldsymbol{e}_j$ are the columns of $\boldsymbol{I}_n$ corresponding to the ith and jth standard basis vectors of $\mathcal{W}_n$, and $\boldsymbol{e}$ is a column of $\boldsymbol{A}_2$ but not $\boldsymbol{I}_n$, then there exists, by the connectivity of $M - \boldsymbol{e}$, a circuit C in $M - \boldsymbol{e}$ containing $\boldsymbol{e}_i$ and $\boldsymbol{e}_j$. It is easy to deduce that the columns of $\boldsymbol{A}_2$ that correspond to C and do not belong to $\boldsymbol{I}_n$ all have at least two nonzero entries, that at least one of them has a nonzero entry in the ith coordinate, and at least one has a nonzero entry in the jth coordinate. Hence, by the above argument applied to each such $\boldsymbol{f}$ in C, the ith and jth diagonal entries of $\boldsymbol{B}''$ must be identical, and their common value agrees with the value of each diagonal entry of $\boldsymbol{D}''$ corresponding to a column belonging to C. This shows that all diagonal entries of $\boldsymbol{B}''$ and of $\boldsymbol{D}''$ have the same common value. Thus the choice for $\boldsymbol{B}$ and $\boldsymbol{D}$ is indeed unique, up to a simultaneous change of sign.

To show that M is recoverable, consider any $\boldsymbol{e}_1 \in \mathcal{W}_n$ such that $M = M(\boldsymbol{A}_1|\boldsymbol{e}_1)$, and consider any two matrices $\boldsymbol{A}'_1, \boldsymbol{A}'_2$ such that $\boldsymbol{A}_1|\boldsymbol{A}'_1 \approx \boldsymbol{A}_2|\boldsymbol{A}'_2$. By Theorem 8.1 there exist $\boldsymbol{B}$ and $\boldsymbol{D}$ as in the statement of the theorem such that $\boldsymbol{A}_1 = \boldsymbol{B}\boldsymbol{A}_2\boldsymbol{D}$, and as just shown, $\boldsymbol{B}$ and $\boldsymbol{D}$ are uniquely determined up to a choice of sign. Theorem 8.1 applied to $\boldsymbol{A}_1|\boldsymbol{A}'_1 \approx \boldsymbol{A}_2|\boldsymbol{A}'_2$ combined with the uniqueness of B and D implies

$$\boldsymbol{A}_1|\boldsymbol{A}'_1 = \boldsymbol{B}(\boldsymbol{A}_2|\boldsymbol{A}'_2)\boldsymbol{D}.$$

It is clear that by choosing $\boldsymbol{e}_2 = \boldsymbol{B}^{-1}\boldsymbol{e}_1$ (independently of $\boldsymbol{A}'_1$ and $\boldsymbol{A}'_2$) $\boldsymbol{A}_1|\boldsymbol{A}'_1|\boldsymbol{e}_1 \approx \boldsymbol{A}_2|\boldsymbol{A}'_2|\boldsymbol{e}_2$ holds, and this finishes the proof. $\square$

As noted earlier, the consistency of $\mathcal{A}[3, U_4^2]$ implies that $\mathcal{A}[3, U_4^2]$ is itself a polynomial recognition algorithm for the class of 4-chromatic graphs whose non-3-colorability it can be used to establish. The algorithm further has the property that if it succeeds for a 4-chromatic graph G then it succeeds for any 4-chromatic graph H that contains G as a subgraph. This follows from the consistency of $\mathcal{A}[3, U_4^2]$, and the fact that any sequence of steps performed while starting from the input matrix $\boldsymbol{A}_G$ and terminating in Phase 1 corresponds to an analogous sequence of steps starting from input $\boldsymbol{A}_H = \boldsymbol{A}_G | \boldsymbol{A}_{H-E(G)}$ and also terminating in Phase 1, simply by disregarding the columns of $\boldsymbol{A}_{H-E(G)}$.

This observation makes it particularly interesting to investigate the behavior of $\mathcal{A}[3, U_4^2]$ when applied to the class of minimal 4-chromatic graphs, commonly known as 4-*critical* graphs. Table 1 shows the result of implementing the algorithm and applying it to the 4-critical graphs on at most 11 vertices.

# vertices	# 4-critical graphs	# failures
4	1	0
5	0	0
6	1	0
7	2	0
8	5	0
9	21	1
10	150	4
11	1221	53

Table 1. Performance of the algorithm $\mathcal{A}[3, U_4^2]$ on small 4-critical graphs.†

It is interesting to note that $\mathcal{A}[3, U_4^2]$ succeeds for the majority of the smallest 4-critical graphs. Figure 2 shows the unique 4-critical graph

†The complete list of small 4-critical graphs was kindly supplied by Gordon F. Royle, who has made it publically available on http://www.cs.uwa.edu.au:80/gordon/graphs/

on 9 vertices for which the algorithm fails.

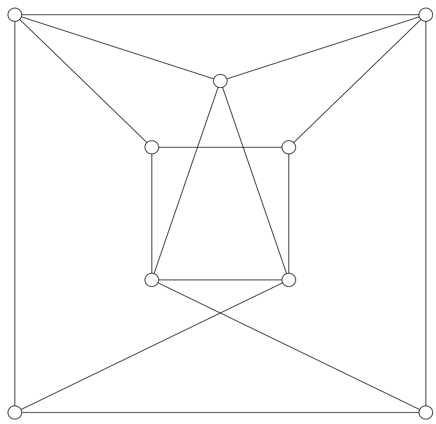

Figure 2. A 9-vertex 4-critical graph for which $\mathcal{A}[3, U_4^2]$ fails.

As noted, $\mathcal{A}[3, U_4^2]$ fails for all 3-edge-connected graphs that are trianglefree. The graph of Figure 2 shows that it may fail also for 3-edge-connected graphs with triangles. It seems an interesting open problem to describe the exact structure of the failure graphs for this particular algorithm.

Another problem with relevance for the four color theorem may be posed: Is it possible to find a basis $\mathcal{M}$ of ternary matroids having the property that $\mathcal{A}[3, \mathcal{M}]$ succeeds for every line-graph of a planar 3-regular graph? By the well-known Tait's Theorem the statement that the line-graph of a planar 3-regular bridgeless graph is 3-colorable is equivalent to the four color theorem (e.g. see [2]). It is perhaps encouraging that already the very simple algorithm $\mathcal{A}[3, U_4^2]$ succeeds for the line-graph shown in Figure 3 of the famous Petersen graph, which may be verified straightforwardly with some effort. This fact seems especially noteworthy given that there exists a widespread belief that there is a close relation between the non-3-colorability of the graph in Figure 3 and the four color theorem (e.g. see Problem 12.15 of [7]).

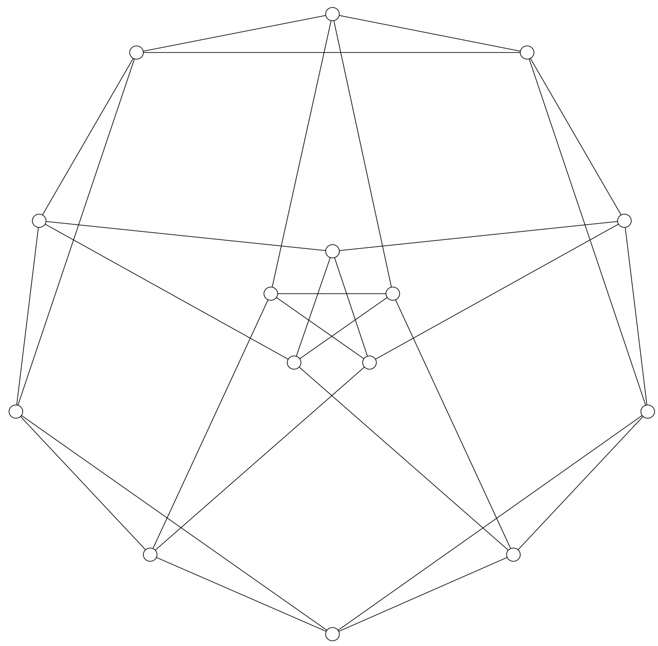

Figure 3 The line-graph of the Petersen graph.

9 Concluding remarks

It has been explained above how Theorem 7.2 provides conditions sufficient to deduce that $\mathcal{A}[q, \mathcal{M}]$ is a polynomial recognition algorithm for three classes of graphs, or more generally of matroids representable over $\mathbb{F}_q$, one with the property that all its members are q-colorable (by $\mathcal{A}[q, \mathcal{M}]$), another with the property that its members are not q-colorable (which may be demonstrated by $\mathcal{A}[q, \mathcal{M}]$), and the third with the property that $\mathcal{A}[q, \mathcal{M}]$ fails on each of its members. Two pertinent questions arise.

A. Given $\mathcal{M}$, is there a purely structural characterization of the elements of each of these classes?

B. What is the best possible choice for $\mathcal{M}$? In terms of algorithmic efficiency, or in terms of obtaining the best possible descriptions of the classes for which $\mathcal{A}[q, \mathcal{M}]$ succeeds or fails, assuming its consistency.

For an answer to question B a couple of directions may be pointed out.

It is fairly clear that any $M \in \mathcal{M}$ is redundant if there exists for every $\boldsymbol{e} \in M$ a submatroid $M' - \boldsymbol{e}' \subseteq M - \boldsymbol{e}$ with $\boldsymbol{e}' \in M' \in \mathcal{M}$. Every application of Step 1.1 of $\mathcal{A}[\mathcal{M}]$ which involves M and an $\boldsymbol{e} \in M$ can then be replaced by an application of Step 1.1 involving M' and $\boldsymbol{e}'$ instead. Assuming $\mathcal{A}[\mathcal{M}]$ is consistent, this implies that $\mathcal{A}[\mathcal{M} \setminus \{M\}]$ always terminates with precisely the same output as $\mathcal{A}[\mathcal{M}]$. As an example, $\mathcal{A}[3, \{U_4^2, K_4\}]$ can easily be proved consistent from Theorem 7.2, but despite its larger basis, it can by this argument be no more efficient than $\mathcal{A}[3, U_4^2]$, since the removal of any edge from K_4 leaves a graph K_4^- containing a 3-circuit, that is, $M(U_3^2) \subseteq M(K_4^-)$.

It seems reasonable then, that one should choose $\mathcal{M}$ in such a way that the matroids in $\mathcal{M}$ all have as few elements as possible (if for no other reason then for reasons of computational efficiency), and in addition such that there is for no $M \in \mathcal{M}$ and $\boldsymbol{e} \in M$ any $M' \in \mathcal{M}$ and $\boldsymbol{e}' \in M'$ such that $M' - \boldsymbol{e}' \subseteq M - \boldsymbol{e}$. Thus the study of sets

$\mathcal{M}$ with this property would seem interesting. It could perhaps be expected that this property combined with the assumption that $\mathcal{M}$ is A/C-closed already suffices to make $\mathcal{A}[\mathcal{M}]$ consistent.

A different approach may be suggested. A loopless matroid M representable over $\mathbb{F}_q$ is known as a *tangential 1-block* over $\mathbb{F}_q$ if M is uncolorable, but every proper loopless *minor* is colorable, that is, any loopless matroid $(M - X)/Y$, where $X, Y \subseteq M$, $X \cap Y = \emptyset$ and $X \cup Y \neq \emptyset$. It is clear that U^2_{q+1} is a tangential 1-block over $\mathbb{F}_q$ and easy to see that U^2_3 is the unique tangential 1-block over $\mathbb{F}_2$, and it was shown in [9] that U^2_4 and $M(K_4)$ are precisely the tangential 1-blocks over $\mathbb{F}_3$. This may suggest to choose as a natural candidate for $\mathcal{M}$ the set of tangential 1-blocks over $\mathbb{F}_q$ (or in the case $q = 3$ choose $\mathcal{M} = \{U^2_4\}$, for the reason already explained). Unfortunately this has the drawback in practice that the precise set of tangential 1-blocks over $\mathbb{F}_q$ is unknown for all $q \geq 4$, and it even seems unclear whether the set is finite in general (further discussion can be found in [4]). It seems nonetheless that a study of algorithms $\mathcal{A}[q, \mathcal{M}]$ where $\mathcal{M}$ is a set of tangential 1-blocks over $\mathbb{F}_q$ might still be illuminating.

10 Acknowledgments

I would like to thank Jens Clausen, Reinhard Diestel, Jack Edmonds, Volker Mehrmann and Bjarne Toft for their interest in this project and for helpful discussions, and also to thank Gordon F. Royle for supplying a catalogue of the 4-critical graphs on at most 11 vertices.

References

[1] K. Appel, W. Haken and J. Koch, Every planar map is four colorable, Part I and II, *Illinois J. Math.* **21** (1977) 429–567.

[2] J.A. Bondy and U.S.R. Murty, Graph Theory with Applications, Macmillan 1976.

[3] T.H. Brylawski and D. Lucas, Uniquely representable combinatorial geometries, In: *Teorie Combinatorie, Proc. 1973 Internat. Colloq.* pp. 83–104. Acc. Naz. dei Lincei, Roma 1976.

[4] T.H. Brylawski and J.G. Oxley, The Tutte Polynomial and Its Applications, in: *Matroid Applications. Encyclopedia of Mathematics and Its Applications.* (Ed. N. White), pp. 123–225. Cambridge University Press 1992.

[5] H.H. Crapo and G.-C. Rota, *On the Foundations of Combinatorial Theory: Combinatorial Geometries*. M.I.T. Press 1970.

[6] M.R. Garey and D.S. Johnson, *Computers and Intractability*, Freeman 1979.

[7] T.R. Jensen and B. Toft *Graph Coloring Problems*, Wiley 1995.

[8] J.G. Oxley, *Matroid Theory*, Oxford University Press 1992.

[9] P.N. Walton and D.J.A. Welsh, Tangential 1-blocks over GF(3), *Discrete Math.* **40** (1982) 319–320.

[10] D.J.A. Welsh, *Matroid Theory*, Academic Press 1976.

[11] D.J.A. Welsh, Matroids: Fundamental Concepts, in: *Handbook of Combinatorics*, pp. 481–526, Elsevier 1996.