

THE DINITZ PROBLEM

ALDEN MATHIEU

1. INTRODUCTION

The so-called "Dinitz conjecture," now the "Dinitz theorem," was a simple-sounding open question regarding the coloring of a graph or array, combining elements of combinatorics and graph theory. Proposed by Jeff Dinitz in the late Seventies, it wasn't solved until the early Nineties, when Fred Galvin produced a simple solution. We follow David Glynn's presentation of the material from [1, Chapter 36], which combines elements we recognize from earlier talks, including Daniel Matthews' discussion of the graph theory chapter "Five-coloring plane graphs" and Alden Mathieu's combinatorics presentation on "Three famous theorems on finite sets."

2. DINITZ'S PROBLEM

Question. *Given n^2 cells arranged in an $(n \times n)$ square array, let (i, j) denote the cell in row i and column j . For each cell (i, j) , we have a color set of $C(i, j)$ of n colors.*

Is it always possible to color the whole array, by selecting colors from $C(i, j)$ for each cell, such that the colors in each row and column are distinct?

It is easiest to visualize this array coloring as a Latin square.

Definition 2.1. A *Latin square* is an $(n \times n)$ array, where the n^2 cells are filled with the numbers $\{1, 2, \dots, n\}$ such that each number appears exactly once in each row and column.

For an $(n \times n)$ Latin square with a choice of the same n colors for each cell, this is **always** possible. Consider the Latin square below. We label the first row arbitrarily with distinct colors for each cell, and permute the row as we progress down the square. Coloring a Latin square in this way is identical to solving a sudoku puzzle.

1	2	3
3	1	2
2	3	1

TABLE 1. A Latin square with distinct coloring in each row and column.

Question. *What if the total set of colors has more than n colors? That is, is it still possible to color the $(n \times n)$ array if $C := \cup_{i,j} C(i, j) > n$ but each cell (i, j) has a set of n colors, $C(i, j)$, to choose from?*

Date: 14 November 2018.

Even with the simplest possible example, with a (2×2) Latin square with 3 colors to choose from, this is not always possible:

$$\begin{array}{|c|c|} \hline \{1,2\} & \{2,3\} \\ \hline \{1,3\} & \{2,3\} \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & 3 \\ \hline \end{array}$$

TABLE 2. The (2×2) array can fail to be distinctly colored with an awkward choice of just 3 colors.

To prove the Dinitz conjecture, we will restate it in terms of graph theory to make our proof easier.

3. GRAPH THEORY

3.1. Background.

Definition 3.1. A *simple graph* $G = (V, E)$ is an undirected graph without loops or multiple edges: that is, a simple graph is a collection of V vertices which are connected by E edges such that any pair of vertices in G is connected by at most 1 edge.

Definition 3.2. A *bipartite graph* $G = (V, E)$ is a graph whose vertex set V can be partitioned into two disjoint sets X, Y such that every edge has one endpoint in X and the other endpoint in Y .

Definition 3.3. A *coloring* (specifically, a *vertex coloring*) of a graph is an assignment of a color to each vertex from a given set of colors, such that it is *proper* (no pair of adjacent vertices share the same color).

Bipartite graphs can be equivalently defined as those graphs which are *2-colorable*.

Definition 3.4. The *chromatic number* of a graph G , denoted $\chi(G)$, is the smallest number of colors that one can assign to the vertices such that the coloring is proper.

The chromatic number of a graph can also be visualized as the minimal number of independent sets which partition the set of vertices V .

Definition 3.5. The *list coloring* is a coloring of the vertices v in V , denoted $L(v)$, where the list of allowed colors for each vertex is a subset of all colors in the graph.

Hence, the *list chromatic number* of a graph G , denoted $\chi_l(G)$, is the minimum k such that if every color set $C(v)$ has size k for all vertices v in V , then a list coloring exists. We note that $\chi(G) \leq \chi_l(G)$ always, since an ordinary coloring is a special case of list coloring.

3.2. Reframing the Dinitz problem. We reframe the Dinitz problem by assigning a vertex to each cell of our $(n \times n)$ array, changing our Latin square problem to a graph theory problem. Our graph, S_n , has n^2 vertices, with two vertices adjacent if and only if they are in the same row or column (see Figure 1).

Since we have n^2 vertices and any n cells in a row are pairwise adjacent, and recalling our Latin square (Table 1), we clearly require at least n colors. Since we have shown the existence of Latin squares, we infer that $\chi(S_n) = n$.

Given that this is the "Dinitz theorem," and not the "disproven Dinitz conjecture," we proceed to prove the restated Dinitz problem:

Theorem 3.6 (Dinitz theorem). $\chi_l(S_n) = n, \forall n$





FIGURE 1. Associating a graph to an array

Remarks 3.7. It is **not** true that $\chi(G) = \chi_l(G)$ for any graph G . Consider the exercise given in the book and depicted below (Figure 2). As a complete bipartite graph, $K_{2,4}$ has a chromatic number of 2 but is demonstrated not to be 2-list-colorable; $\chi_l(K_{2,4}) = 3$.

FIGURE 2. $K_{2,4}$ is 3-list-colorable

A 1979 paper by Erdős, Rubin and Taylor [2] demonstrates that there is no bound on how much the list chromatic number of a graph can exceed the chromatic number of the graph as the list of colors grows.

To prove the Dinitz theorem, we will need two results and a few further graph theory concepts.

Definition 3.8. An *induced subgraph* G_A is the subgraph of G which has the set of vertices $A \subseteq V$ and the set of all edges of G which connect the vertices in A .

G_A is the subgraph induced by A .

Definition 3.9. A *directed graph*, denoted $\vec{G}(V, E)$, is a graph on which each edge has an *orientation* or direction. The notation $u \rightarrow v$ denotes the directed edge starting at vertex u and ending at vertex v .

Definition 3.10. For a vertex v in a graph, we denote the *degree* of v by $d(v)$. The degree of v is the number of edges which connect to v .

We therefore define the *outdegree* $d^+(v)$ as the number of directed edges which start at v . Similarly, the *indegree* $d^-(v)$ is the number of directed edges which end at v .

We also have that $d(v) = d^+(v) + d^-(v)$.

Definition 3.11. A set of vertices S in a graph are *independent* if they are pairwise non-adjacent; that is, for every pair of vertices of S , there is no connecting edge between them.

Definition 3.12 (Kernel). Let $\vec{G} = (V, E)$ be a directed graph. A *kernel* $K \subseteq V$ is a subset of the vertices such that

- i) K is independent in G
- ii) for every vertex $u \in V \setminus K$, there is a vertex $v \in K$ such that $u \rightarrow v$. That is, all vertices not in K have an edge directed into K .

We now can state the first result.

4. FIRST RESULT

Lemma 4.1. Let $\vec{G}(V, E)$ be a directed graph, and suppose for all $v \in V$ there is a color set $C(v)$ such that $|C(v)| \geq d^+(v) + 1$. Then, if every induced subgraph of \vec{G} has a kernel, then G has a list coloring with $C(v)$ as the color set.

Proof. We use induction on the number of vertices. If there is only one vertex, the lemma is trivial; we consider only $|V| > 1$.

Let c be a color, $c \in C = \cup_{v \in V} C(v)$. Define $A(c) := \{v \in V : c \in C(v)\}$, the set of vertices with color c in the restricted list $C(v)$. Therefore $G_{A(c)}$ has a kernel: $K(c)$. Because all the vertices in $K(c)$ are independent (by definition 3.12(i)), we are able to color all the vertices in $K(c)$ with the color c . We then delete $K(c)$ from G and delete c from C .

Now we let G' be the induced subgraph on $V \setminus K(c)$ and let the reduced $C'(v) = C(v) \setminus \{c\}$. For all vertices u in $A(c) \setminus K(c)$, we have reduced the outdegree by at least 1 (by definition 3.12(ii)), so the condition $|C'(u)| \geq d^+(u) + 1$ still holds. Similarly, the condition holds for vertices outside $A(c)$ since the sets $C(v)$ were not changed. Thus, our new graph \vec{G}' has a list coloring.

By adding our removed colors back in, we can see that by induction we get a list coloring for \vec{G} . \square

4.1. Strategy. We now see that we need to find an orientation of edges of our associated graph S_n such that

- i) $d^+(v) \leq n - 1$, for all vertices v
- ii) every induced subgraph has a kernel

Hence, for our second result, we require concepts which should be somewhat familiar from an earlier presentation on Hall's marriage theorem.

5. SECOND RESULT

Definition 5.1. A *matching* M in a bipartite graph G is a set of edges such that no pair of edges share an endpoint.

Definition 5.2. For all vertices $v \in V$ of a graph $G(V, E)$, we consider the set $N(v)$, the set of adjacent vertices to v . We assume for all v there is a preferential ranking or ordering of the set $N(v) = \{z_1 > z_2 > \dots > z_{d(v)}\}$ where the highest preference for v is the vertex z_1 and the lowest is $z_{d(v)}$.

While the common analogy is to partners in marriages, a more apt and less sexist model might be a student's course choices in the CAO and a college's admissions offers.

Definition 5.3. A *stable matching* M of the bipartite graph $G(V, E)$ is a matching such that, if edge uv is **not** in the matching M , then either

- i) $uy \in M$ with $y > v$ in $N(u)$ (that is, u has a higher preference for y than for v); or
- ii) $xv \in M$ with $x > u$ in $N(v)$ (that is, v has a higher preference for x than for u); or
- iii) both.

Thus, in our CAO interpretation, a matching is stable if it **never** happens that a student s is **not** planning to attend a university u , but s prefers u to their current offer and u prefers student s to their current prospective cohort. (If student s would rather go to university u , and u would rather have s attend, then the current dissatisfactory matching obviously would not be stable.)

We consider a diagram of a bipartite graph, with a set of students $\{a, b, c, d\}$ and a set of universities $\{A, B, C, D\}$ with their respective ranking sets. While a largest matching exists (the thin lines), it is not stable. For example, student c prefers university A and vice versa; the bold lines indicate the smaller stable matching.




FIGURE 3. Tough luck for d .

We now state the second result.

Lemma 5.4. *A stable matching always exists.*

Proof. We consider an algorithm: in round 1, all universities offer an spot $u \in U$ to their top choice. If a student receives more than one offer, they select their highest preference; if they only receive one offer, they keep it. If a spot in a university isn't accepted, it goes into the pool of rejections ($R \subseteq U$) and is offered in the next round.

In round 2, all the universities with spots in R offer to their next choice. The students compare their offers (if any) to their first round offers (if any) and select their highest preference. The rest are rejected and form the new pool of rejections R .

And so on. A university who has offered to their last choice of student and is again turned down drops out from the pool R and from further consideration (our model algorithm doesn't account for last-minute still-open course listings on the CAO site). Eventually, R is empty, having seen every university either fill its spot or drop out of consideration, and the algorithm stops.

We claim that, when the algorithm stops, the universities that have successfully filled their spot and the corresponding students form a stable matching M .

Note that the universities accepted by a student are accepted in increasing preference (of the student), since in each round the student picks their highest preference between their current offer and the new offer. Hence, if su is in the set of edges but $su \notin M$, then either: u never offered a spot to s , implying that $ut \in M$ with $t > s$ in $N(u)$ (criterion 5.3(i)); or s rejected a spot at u , implying that $sc \in M$ with $c > u$ in $N(s)$ (criterion 5.3(ii)). But this exactly describes a stable matching. \square

With our two results, we now are able to formulate Galvin's proof of the Dinitz theorem.

6. BRINGING IT HOME

We restate the theorem 3.6 for convenience:

Theorem 6.1 (Dinitz theorem). $\chi_l(S_n) = n, \forall n$

Proof. Let (i, j) denote the vertices of S_n . We have (i, j) adjacent to (r, s) if and only if $i = r$ or $j = s$.

We construct \vec{S}_n . We consider an arbitrary Latin square L with elements $\{1, 2, \dots, n\}$ and denote by $L(i, j)$ the number in cell (i, j) . We then define an orientation horizontally by $(i, j) \rightarrow (i, j')$ if $L(i, j) < L(i, j')$ and vertically by $(i, j) \rightarrow (i', j)$ if $L(i, j) > L(i', j)$. Thus the Latin square and associated graph in Fig. 1 have the following orientation:




FIGURE 4. Defining an orientation on S_3

If $L(i, j) = k$ then there are $n - k$ cells in row i with entries larger than k and $k - 1$ cells in column j with entries smaller than k . Hence, $d^+(v) = n - k + k - 1 = n - 1$ and we have proved our first requirement.

It remains to show that every induced subgraph has a kernel.

We consider a subset $A \subseteq V$ of vertices and associate to A the bipartite graph $G = (X \cup Y, A)$ with $X = Y = \{1, 2, \dots, n\}$, and X the set of rows of L and Y the set of columns. Then $(i, j) \in A$ implies that $ij \in G$ is an edge, $i \in X, j \in Y$. In the diagram below, the cells of A are shaded and correspond to the matching on the associated bipartite graph G .




FIGURE 5. The cells in A are shaded and correspond to the matching on the graph.

We use the above described orientation on S_n to define a ranking on the vertices of G : $j' > j$ in $N(i)$ if $(i, j) \rightarrow (i, j') \in S_n$ and similarly $i' > i$ in $N(j)$ if $(i, j) \rightarrow (i', j) \in S_n$. By the lemma 5.4, G has a stable matching M . We claim that $M \subseteq A$ is a kernel. M is a subset of A since it is a set of edges in G and hence a set of vertices. Firstly, we know that M is independent in A because, as edges in G , they do not share vertices. Secondly, if $(i, j) \in A \setminus M$, then either $(i, j') \in M$ and $j' > j$ and thus $(i, j) \rightarrow (i, j') \in S_n$, or else $(i', j) \in M$ and $i' > i$ and thus $(i, j) \rightarrow (i', j) \in S_n$. Thus, M satisfies the conditions that define a kernel. \square

Thus we have finally shown that in fact, it is always possible to n -list-color the graph S_n and thus an $(n \times n)$ array.

REFERENCES

- [1] Martin Aigner and Günter M. Ziegler. *Proofs from The Book*, pages viii+308. Springer-Verlag, Berlin, fifth edition, 2014. Including illustrations by Karl H. Hofmann.
- [2] Paul Erdos, Arthur L Rubin, and Herbert Taylor. Choosability in graphs. *Congr. Numer.*, 26:125–157, 1979.