CHOOSABILITY IN GRAPHS

by

Paul Erdös

Hungarian Institute of Sciences
Budpest, Hungary

Arthur L. Rubin*

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91103

Herbert Taylor**

Department of Electrical Engineering
University of Southern California
Los Angeles, California 90007

*The results of one phase of research carried out at the Jet Propulsion Laboratory, California Institute of Technology, under Contract NAS 7-100, sponsored by National Aeronautics and Space Administration.

**Research sponsored in part by the Army Research Office under Grant number DAAG29-79-C-0054.

Let the nodes be named 1,2,...,n in a graph G. Given a function f on the nodes which assigns a positive integer f(j) to node j, we'll require that the adversary put f(j) distinct letters on node j, for each j from 1 to n. Now we'll say that G is f-choosable if, no matter what letters the adversary puts, we can always make a choice consisting of one letter from each node, with distinct letters from adjacent nodes.

Using the constant function f(j)=k, we'll say that the choice number of G is equal to k if G is k-choosable but not (k-1)-choosable.

For the complete graph with $n \geq 1$, it is always true that the choice number of K_n is equal to n.

Since one of the things the adversary may do is put the same k-set of letters on every node of G, it follows that the choice number = choice #G $\geq \chi G$ = the chromatic number of G. Here is an example of a graph

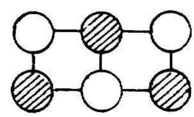

which is 2-colorable (therefore has chromatic number ≤ 2) but not 2-choosable (therefore has choice number > 2). If the adversary uses the pattern pictured below, then no choice of one letter from each node can have distinct letters from every adjacent pair of nodes.

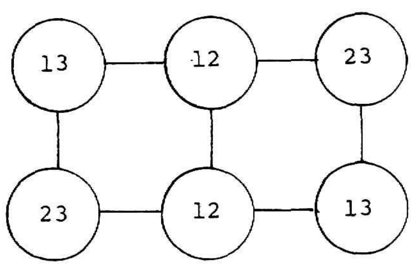

EXAMPLE SCHEME SHOWING HOW CHOICE #G EXCEEDS χG

There is no bound on how much choice #G can exceed χG, as n increases. The complete bipartite graph $K_{m,m}$ is 2-colorable.

But if $m=\binom{2k-1}{k}$, then $K_{m,m}$ is not k-choosable.

The adversary can construct a pattern to prove it as follows. Recall that $m=\binom{2k-1}{k}$ represents the number of k-subsets of a (2k-1)-set. Picture $K_{m,m}$ with m nodes in the top row, and m nodes in the bottom row, having an edge between two nodes iff one is in the top row, and the other is in the bottom row. Let the letters be the elements of a (2k-1)-set. Put each k-subset of letters on one node of the top row, and on one node of the bottom row.

When we try to make a choice, we find it must include k distinct letters from nodes of the top row - otherwise a k-set consisting of letters not chosen from any node would be the k-subset of letters on some node in the top row. But now the attempted choice must fail in the bottom row because some set of k distinct letters, already chosen from nodes in the top row, will be exactly the k-subset of letters put on some node of the bottom row.

Thus $K_{m,m}$ is not k-choosable, when $m=\binom{2k-1}{k}$.

Here is the picture when k=3, $m=\binom{5}{3}=10$, and the set of letters is {1,2,3,4,5}. The dashed line is meant to suggest the 100 edges which connect nodes of the top row with nodes of the bottom row.

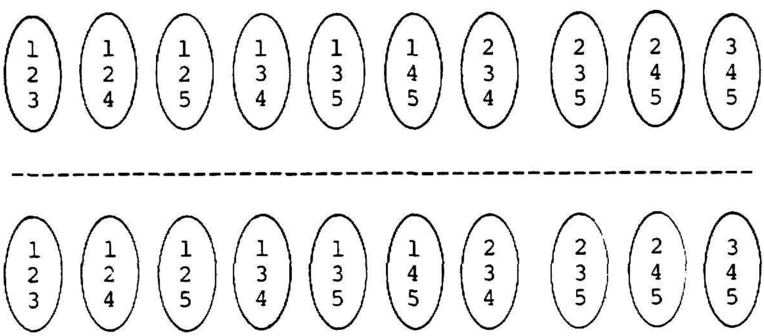

OPEN QUESTION

What is the minimum number $N(2,k)$ of nodes in a graph G which is 2-colorable but not k-choosable?

BOUNDS FOR $N(2,k)$

A family F of sets has property B iff there exists a set B which meets every set in F but contains no set in F. In other words F has property B iff there exists a set B such that

$$\left.\begin{array}{ll} 1. & X \cap B \neq \Phi \\ 2. & X \not\subseteq B \end{array}\right\} \quad \text{for every } X \in F.$$

M_k is defined as the cardinality of a smallest family of k-sets which does not have property B. Although M_k is only known exactly for $k \le 3$, there are bounds for it. The crude bounds

$$2^{k-1} < M_k < k^2 2^{k+1}$$

will suffice here. Sharper bounds can be found in P. Erdos, "On a Combinatorial Problem III", Canad. Math. Bull., vol. 12, no. 4, 1969.

In what follows we shall prove that

$M_k \le N(2,k) \le 2M_k$.

To establish the upper bound, we argue that $K_{m,m}$ is not k-choosable when $m \ge M_k$. For the k-sets of letters on nodes of the top row the adversary can use a family F which does not have property B, and use the same F on the bottom row. If C is any set of letters chosen one from each node of the top row, then of course $X \cap C \neq \Phi$ for every $X \in F$, and consequently there must exist $W \in F$ such that $W \subseteq C$. But then in the bottom row no letter can be chosen from the node which has W.

To establish the lower bound, we argue that $K_{b,t}$ is k-choosable when $b+t < M_k$. With t nodes in the top row, and b nodes in the bottom row, let F be the family of k-sets of letters assigned to the nodes. F will have property B because $|F| < M_k$, and we can use B to make our choices. First choose a letter of B from each node in the top row - the choice exists because B meets each of them. Then choose a letter not in B from each node of the bottom row - that choice exists because B does not contain any of them. Two nodes are adjacent only when one is in the top row, and the other in the bottom row, and their chosen letters are distinct because one is in B, and the other is not in B.

That completes the proof. The following theorem summarizes the above discussion.

THEOREM

$$2^{k-1} < M_k \le N(2,k) \le 2M_k < k^2 2^{k+2}.$$

Here is all we know regarding exact evaluation of N(2,k).

$M_1=1$	$N(2,1) = 2$
$M_2=3$	$N(2,2) = 6$
$M_3=7$	$12 \le N(2,3) \le 14$

Although it is most likely that N(2,3)=14, it would be quite a surprise if $N(2,k)=2M_k$ were to persist for large k.*

$K_{7,7}$ is pictured below with the adversary's assignment which shows it is not 3-choosable. Again the dashed line indicates the 49 edges.

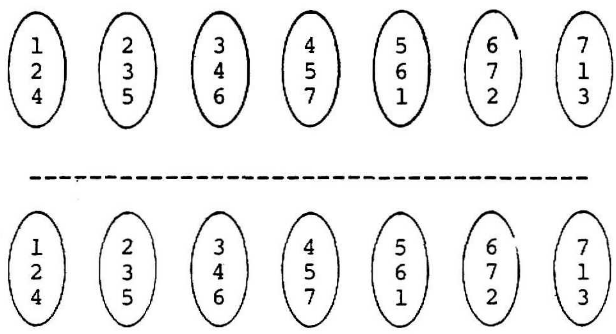

*We know that $M_k+1 < N(2,k)$, for $k > 1$.

CHARACTERIZATION OF 2-CHOOSABLE GRAPHS

A graph is 2-choosable iff each connected component is 2-choosable, so we restrict our attention to connected graphs. To start the investigation of which graphs are 2-choosable, consider a node of valence 1. We can always choose one of its two letters <u>after</u> deciding which letter to choose from the one node adjacent. The obvious thing to do is prune away nodes of valence 1, successively until we reach the <u>core</u>, which has no nodes of valence 1. A graph is 2-choosable iff its core is 2-choosable.

By definition let's say a Θ graph consists of two distinguished nodes i and j together with three paths which are node disjoint except that each path has i at one end, and j at the other end. Thus a Θ graph can be specified by giving the three paths' lengths. Here are some examples.

$\Theta_{1,2,2}=$ 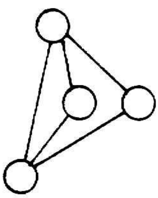 , $\Theta_{1,2,3}=$ 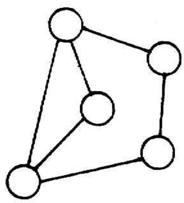 , $\Theta_{1,3,3}=$ 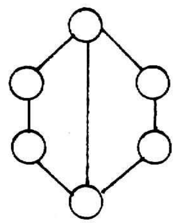

$\Theta_{2,2,2}=$ 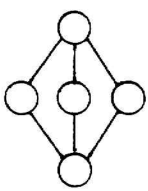 , $\Theta_{2,2,4}=$ 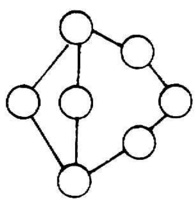 , $\Theta_{2,2,6}=$ 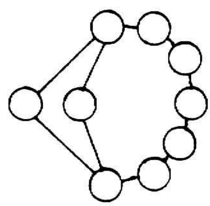

Here is a proof that $\Theta_{2,2,2m}$ is 2-choosable, for $m \ge 1$. Let the assigned 2-sets be named as in the picture.

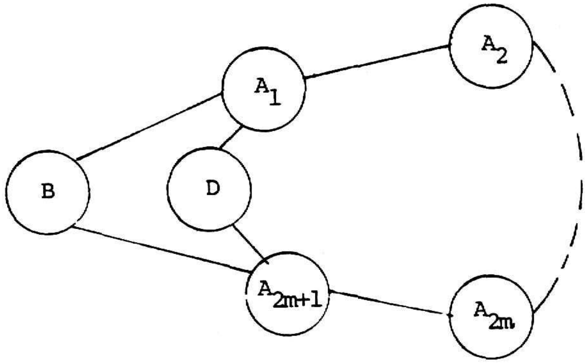

CASE I: Suppose $A_1=A_2=\dots=A_{2m+1}=\{x,y\}$. From A_i choose x when i is odd, y when i is even, so that x is chosen from both A_1 and A_{2m+1}. Complete the choice with a letter from $B-\{x\}$, and a letter from $D-\{x\}$.

CASE II: Suppose the A_j's are not all equal. Find one particular adjacent pair $A_i \neq A_{i+1}$. Tentatively choose $x_i \in A_i - A_{i+1}$, and go in sequence choosing $x_{i-1} \in A_{i-1}-\{x_i\}$, $x_{i-2} \in A_{i-2}-\{x_{i-1}\}$, ... until $x_1 \in A_1-\{x_2\}$. At this point we look ahead to $A_{2m+1}=\{b,d\}$, and look at B and D. If $\{B,D\} \neq \{\{x_1,b\},\{x_1,d\}\}$, then there will exist a choice of $x \in A_{2m+1}$ such that $B-\{x_1,x\} \neq \Phi$ and $D-\{x_1,x\} \neq \Phi$, and so we can continue choosing $x_{2m} \in A_{2m}-\{x\}$, $x_{2m-1} \in A_{2m-1}-\{x_{2m}\}$, ... until $x_{i+1} \in A_{i+1}-\{x_{i+2}\}$, thereby completing the choice. But if $\{B,D\}=\{\{x_1,b\},\{x_1,d\}\}$, then we go back to $A_i \neq A_{i+1}$ and start the other way. Start by choosing $y_{i+1} \in A_{i+1}-A_i$, and go in sequence choosing $y_{i+2} \in A_{i+2}-\{y_{i+1}\}$,... until $y \in A_{2m+1}-\{y_{2m}\}$. Here $y \neq x_1$ so we can choose $x_1 \in B$ and $x_1 \in D$, and continue with $y_1 \in A_1-\{x_1\}$, $y_2 \in A_2-\{y_{i-1}\}$... until we complete the whole choice at $y_i \in A_i-\{y_{i-1}\}$.

That completes the proof that $\Theta_{2,2,2m}$ is 2-choosable.

Since an even cycle C_{2m+2} is a subgraph of $\Theta_{2,2,2m}$, we also know that all even cycles are 2-choosable.

At this point in the investigation, every 2-choosable graph we know about has as its core a subgraph of some $\Theta_{2,2,2m}$. The remarkable fact that no others exist will be told as follows.

Let $\{K_1, C_{2m+2}, \Theta_{2,2,2m} : m \geq 1\} = T$.

<u>THEOREM (A.L. Rubin)</u>

A graph G is 2-choosable if, and only if, the core of G belongs to T.

<u>PROOF</u>

Let G be the core of a connected graph.

The idea of the proof is to show, by exhausting the possibilities, that either G is in T, or else G contains a subgraph belonging to one of the following five types.

1. An odd cycle.
2. Two node disjoint even cycles connected by a path.
3. Two even cycles having exactly one node in common.
4. $\Theta_{a,b,c}$ where $a \neq 2$ and $b \neq 2$.
5. 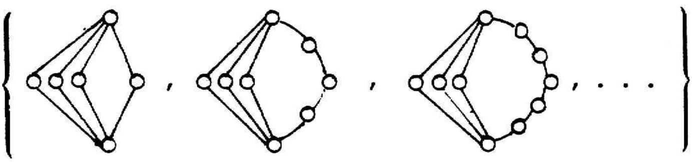

We start by assuming that G is not in T.

If G contains an odd cycle we are done. Thus we proceed on the assumption that G is bipartite.

Let C_1 be a shortest cycle. Note that there must exist an edge of G not in C_1, because otherwise G would <u>be</u> an even cycle.

If there is a cycle C_2 having at most one node in common with C_1, then we will be in case (2.) or (3.), and be done.

Let P_1 be a shortest path, edge disjoint from C_1, and connecting two distinct nodes of C_1. (This is now known to exist.)

If $C_1 \cup P_1$ is not in T, then it must be in case (4.), in which case we are done.

Now suppose $C_1 \cup P_1$ is in T, so it must be a $\Theta_{2,2,2m}$, and C_1 must be a 4-cycle. Observing that there must be more to G, we can say the following.

Let P_2 be a shortest path, edge disjoint from $C_1 \cup P_1$, connecting two distinct nodes of $C_1 \cup P_1$.

Next we examine six cases to see what the end nodes of P_2 might be. It will help to name the nodes of C_1 as shown in this picture of $C_1 \cup P_1$.

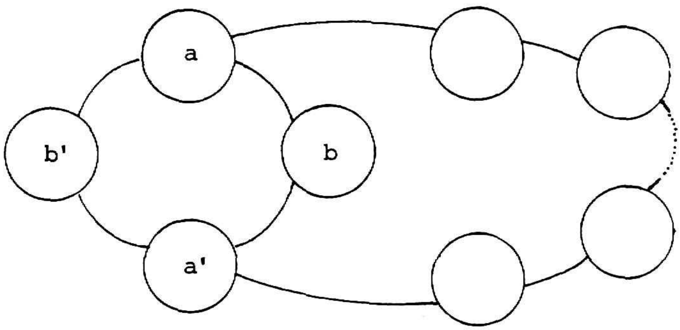

Case (i). If the ends of P_2 are two interior nodes of P_1, then we have a cycle disjoint from C_1, and are in case (2.) again.

Case (ii). If the ends of P_2 are a and an interior node of P_1, then we have a cycle with exactly one node in common with C_1, and are in case (3.).

Case (iii). If the ends of P_2 are b and an interior node of P_1, then we have a path from a to b edge disjoint from C_1, which puts us in case (4.).

Case (iv). If the ends of P_2 are a and b, we are put in case (4.) again, as we were in case (iii).

Case (v). If the ends of P_2 are a and a', and P_1 is of length 2, then we are in case (5.). If P_1 is of length >2, then we are in case (4.).

Case (vi). If the ends of P_2 are b and b', then by removing any edge of C_1 we find a Θ graph which puts us in case (4.).

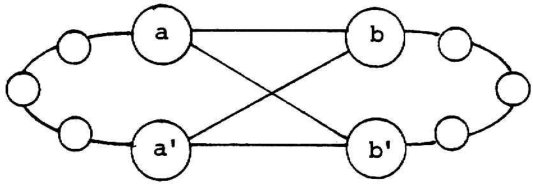

We now know that if G is not in T, then G contains one of the five types. Thus it only remains to show that any graph of type 1., 2., 3., 4., or 5. is not 2-choosable.

Type (1.) is not even 2-colorable.

To deal with 2., 3., 4., 5. we can use the following reduction. Remove a node b, and merge the nodes that were adjacent to b. Any multiple edges that result can be made single, and no "loops" will appear, because the graph remains bipartite. If the reduced graph G' is not 2-choosable, then G is not 2-choosable.

To prove it, suppose G' is not 2-choosable. Unmerge, and assign the same {x,y} to b as to all the nodes adjacent to b. If, say, x is the letter chosen from b, then y will have to be chosen from all the nodes adjacent to b, and therefore a choice for G would have worked just as well for G'. It is worth special notice that this proof would not have worked for 3-choosability.

After repeated application of this reduction process, we will only need to verify that each of the following four particular graphs is not 2-choosable.

2.

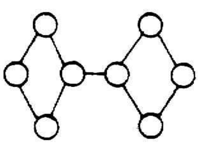

3.

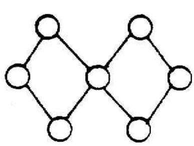

4.

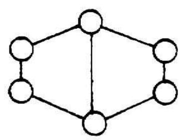

5.

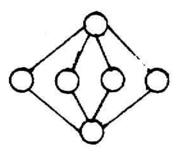

No choice exists for the

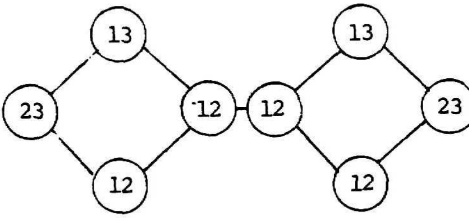

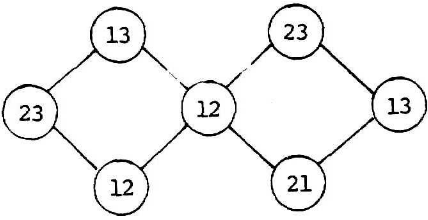

That completes the proof

2-choosable graphs.

assignments shown.

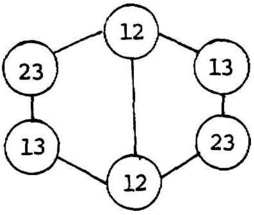

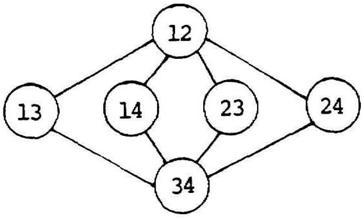

of the theorem characterizing

A THEOREM ON GRAPH STRUCTURE

The following theorem is due to Arthur Rubin. It will lead to a characterization of D-choosability, and consequently to a generalization of Brooks' theorem. But, apart from choosability considerations, here is a remarkable theorem.

THEOREM R. If there is no node which disconnects G, then G is an odd cycle, or $G=K_n$, or G contains, as a node induced subgraph, an even cycle without chord or with only one chord.

PROOF. (by exhaustion, and induction on n)

Assume no node disconnects G, G is not an odd cycle, and $G \neq K_n$. Observe that a Θ graph either contains an even cycle as a (node) induced subgraph, or consists of an even cycle with only one chord. Thus each subcase will be settled when we find an induced even cycle in G, or find an induced Θ graph in G.

CASE I. There is a node of valence 2. Call it N. Remove N, and prune nodes of valence 1 successively. Now look at what is left.

I.1 One node. G must have been a cycle (not odd).

I.2 An odd cycle. G must have been a Θ graph.

I.3 K_m, where $m \geq 4$. We find an induced $\Theta_{1,2,p}$, where p is the length of the pruned off path.

I.4 If I.1, I.2, I.3 do not hold, and still the graph that remains after pruning has no node which disconnects it, then we're done by the inductive hypothesis.

I.5 What remains has a node X which disconnects it. Name the end nodes of the pruned off path A and B. First we argue that A could not disconnect what remains, because contrariwise it would have to have done so before pruning as well. Thus we know $A \neq X \neq B$.

What if A and B were connected by some path not through X? If this were so, then X would have disconnected G before pruning. Thus all paths from A to B go through X.

Let α be a shortest path from A to B. The picture should look something like the one below. Naturally a shortest path cannot have any chords.

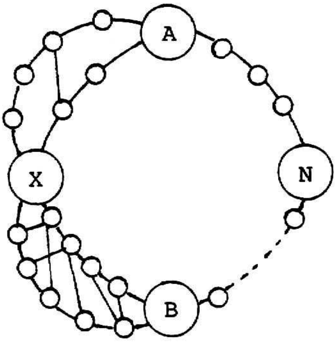

Let β be a shortest path from a node U adjacent to A (U not on α, A not on β), to a node Z adjacent to α-A.

If Z is adjacent to more than one node of α-A, let Y_1 and Y_2 be the two such closest to A along α. Then the nodes on the arc of α from A to Y_2, and on β, induce a cycle with only one chord, that is, a Θ graph.

If Z is adjacent to only one node of α-A, then the nodes of α, β, and the path through N induce a Θ graph.

CASE II. There is no node of valence 2. Delete one node N, and look at what is left.

II.1 It cannot be just one node.

II.2 An odd cycle γ. Note first that N must have been adjacent to every node of γ. If γ were a 3-cycle, then G would have been K_4. Thus γ is a larger cycle, and we find the "diamond", $\Theta_{2,1,2}$, induced in G. It will look like this.

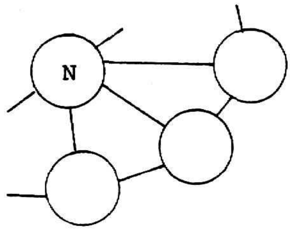

II.3 If G-N is a complete graph, then since G is not K_n, there must be some node Y of G-N which is not adjacent to N. In this case we find a diamond induced in G.

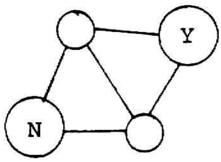

II.4 If not II.2, not II.3, and G-N has <u>no</u> node which disconnects it, then we're done by the inductive hypothesis.

II.5 Otherwise the graph G-N has a node X which disconnects it.

First we observe that the subgraph induced on nodes adjacent to N cannot be a complete graph. If it were, then the node X which disconnects G-N would also disconnect G.

Let α be a shortest path in G-N between two nodes A and B which are adjacent to N, but not themselves adjacent.

If the number of edges of α is equal to 2, then we have C_4 or $\Theta_{2,1,2}$, in the form:

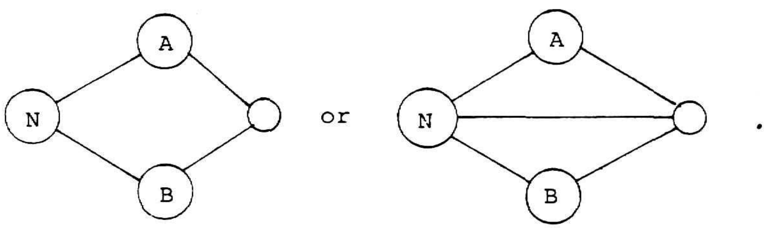

Otherwise α has more than two edges, and we construct as follows.

Let β be a shortest path in G-N, from a node C which is different from A and B but adjacent to N, to a node Z which is adjacent to α. (C=Z is possible)

In case Z is adjacent to two or more nodes of α, we can identify two more nodes, as follows.

Let Y_A be adjacent to Z, along α, closest to A.

Let Y_B be adjacent to Z, along α, closest to B.

The picture below may help remember the above adjacencies.

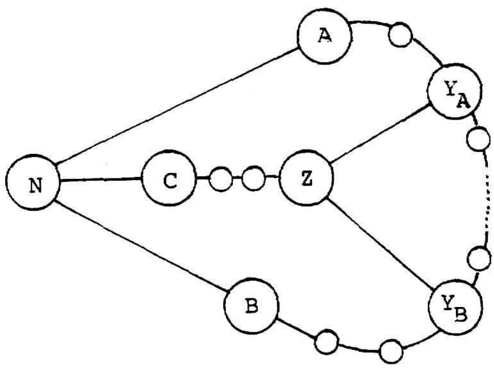

If Y_A is not adjacent to Y_B, then the Θ graph we find is the induced subgraph on N, β , the arc of α from A to Y_A, and the arc of α from B to Y_B.

If Y_A is adjacent to Y_B, and $Y_B \neq B$, then our Θ graph is induced on N, β , and the arc of α from A to Y_B.

If Y_A is adjacent to Y_B, and $Y_B = B$, then $Y_A \neq A$, so it is symmetric with the previous case.

Finally, if Z is adjacent to only one node of α, then our Θ graph is induced by N, α , and β.

The proof of theorem R is complete.

CHARACTERIZATION OF D-CHOOSABILITY

To define a function D on the nodes of G, let D(j) = the valence of node j.

Thus the question of whether G is D-choosable, or not, is posed by specifying that the number of letters assigned to a node shall be equal to the number of edges on that node. We start by exploring graphs which are not D-choosable.

Supposing G and H are two separate graphs, take any node i of G, and any node j of H, and merge them into a single node (ij) to produce a new graph G(ij)H It goes understood that the node (ij) disconnects G(ij)H.

Generate a family non D as follows. For every integer $n \geq 1$, put K_n into non D. Put all odd cycles into non D. Whenever $G \in$ non D and $H \in$ non D, put G(ij)H into non D. A typical member of the family non D will look like this.

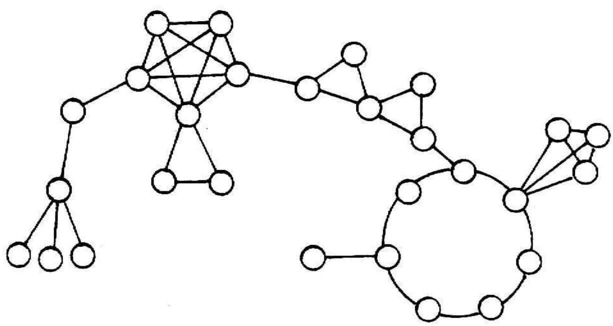

Since all complete graphs and odd cycles are not D-choosable, it will become apparent that every graph in non D is not D-choosable, after we prove a quick lemma.

LEMMA

If G and H are both not D-choosable, then G(ij)H is not D-choosable.

PROOF

Presume the adversary's assignments used different letters on G and H. Let A be the set of letters put on node i of G, and let B be the set of letters put on node j of H. Since D(ij)=D(i)+D(j), the adversary can assign $A \cup B$ to the node ij of G ij H, and keep the other assignments as before the merger. When we try to choose a letter from ij, our choice will fail in G if we take a letter from A, and fail in H if we take a letter from B.

Next we explore graphs which are D-choosable, starting with Θ graphs.

Consider an arbitrary $\Theta_{a,b,c}$ with say, $c \geq 2$. Let the nodes be named $1,2,\ldots,n$ as shown in the picture.

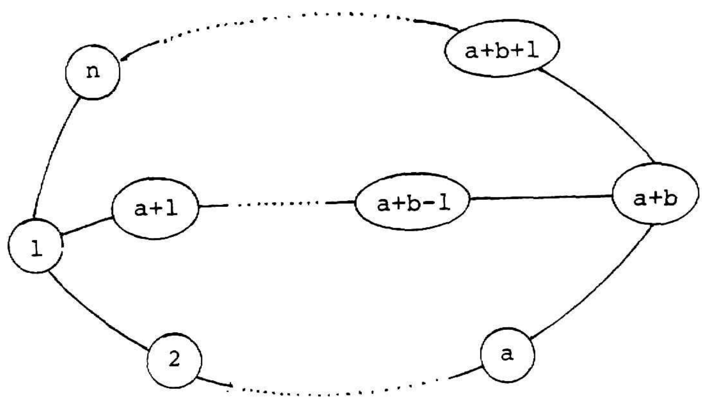

Make the choices in sequence, starting at node 1. Node 1 has three letters, so we can choose a letter not in node n. At each node in sequence there will be more letters than adjacent earlier nodes, until we reach node n. Node n is adjacent to two earlier nodes, but neither of its two letters is excluded by the choice we made at node 1. Thus every Θ graph is D-choosable. Also, recalling our discussion of 2-choosability, we know that every even cycle is D-choosable.

LEMMA

If G is connected, and G has an induced subgraph H which is D-choosable, then G is D-choosable.

PROOF

Assuming G-H is not empty, find a node x, of G-H, which is at maximal distance from H. This guarantees that G-x will be connected. Start the choice with any letter from x, and then erase that letter from all nodes adjacent to x. The choice can be completed because G-x is an earlier case.

THEOREM

Assume G is connected. G is not D-choosable iff $G \in$ non D.

PROOF

Take G and look at parts not disconnected by a node. If every such part is an odd cycle or a complete graph, then $G \in$ non D, and therefore G is not D-choosable.

If some such part is neither an odd cycle nor a complete graph, then Theorem R tells us that G must contain, as a node induced subgraph, an even cycle or a particular kind of Θ graph. By the preceding lemma this means that if $G \notin$ non D, then G is D-choosable.

SAME THEOREM

Assume G is connected. G is D-choosable iff G contains an induced even cycle or an induced Θ graph.

COMMENT

As a consequence of this characterization, we can prove that, for large n, almost all graphs are D-choosable.

DIGRESSION $\rightarrow$ INFINITE GRAPHS

Consider the infinite asterisk

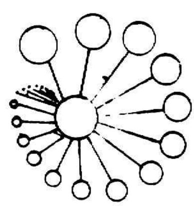

It is not D-choosable, because the adversary can use Z^+, the set of positive integers, thus:

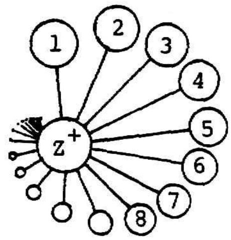

On the other hand, if we disallow infinitely many edges on any one node, we get the following.

<u>THEOREM</u>

Let G be a countably infinite connected graph with finite valence. Then G is D-choosable.

<u>PROOF</u>

Let the nodes of G be named with the positive integers. At each node the number of letters put there by the adversary will be no less than the number of edges on that node. Choose letters by the following rules - with i the smallest named node from which a letter has not yet been chosen.

1. If erasing i does not leave a finite component disconnected from the rest of the graph, choose a letter x from node i. Erase x from the nodes adjacent to i, and remove i from further consideration.

2. If erasing i would disconnect a finite component, deal with each such finite component before dealing with i. In each finite component start with a node at maximal distance from i, to be sure it will not disconnect the component.

By following rules 1. and 2. we can choose a letter from node j for every $j \in Z^+$, and never take the same letter from two adjacent nodes.

COROLLARY: BROOKS' THEOREM

The infinite case of Brooks' Theorem is an immediate consequence of the theorem just proved. The finite case is a consequence of Rubin's characterization of D-choosability. Refer to R.L. Brooks, "On Colouring the Nodes of a Network", Proc. Cambridge Philosophical Soc., vol. 37 (1941).

Here is the statement of his original theorem, verbatim.

"Let N be a network (or linear graph) such that at each node not more than n lines meet (where $n > 2$), and no line has both ends at the same node. Suppose also that no connected component of N is an n-simplex. Then it is possible to colour the nodes of N with n colours so that no two nodes of the same colour are joined.

An n-simplex is a network with n+1 nodes, every pair of which are joined by one line.

N may be infinite, and need not lie in a plane."

Of course for D-choosable graphs, Brooks' theorem holds a fortiori.

Now consider $G \in$ non D. Pick one node j of G, and define a new function ${}^{j}D$ thus: let ${}^{j}D(j)=1+D(j)$, and let ${}^{j}D(i)=D(i)$ if $i \neq j$. We can see that G is ${}^{j}D$-choosable by attaching an infinite tail at j, as in the picture.

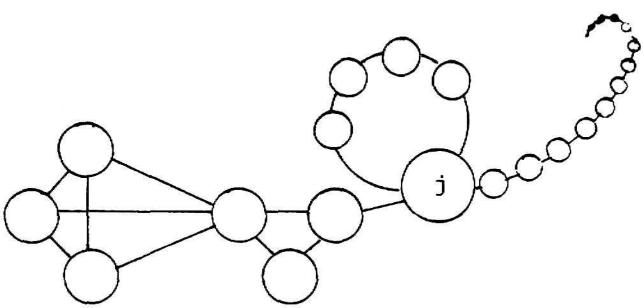

Lastly, with the observation that the only regular (D(i)= constant) graphs in non D are complete graphs or odd cycles, we have a choice version which covers the finite case.

THEOREM

If a connected graph G is not K_n, and not an odd cycle, then choice $\#G \leq \max D(j)$.

END DIGRESSION

According to a well known result of Nordhaus and Gaddum, $XG+X\overline{G} \leq n+1$. Before proving the choice version, we state a lemma which may prove useful elsewhere.

A CHOOSING FUNCTION LEMMA

Let the nodes of G be labeled $1,2,\ldots,n$, as usual. In that order define a choosing function g, as follows.

$$g(j)=1+|\{i:1 \leq i < j \leq n, \text{ and } \{i,j\} \text{ is an edge of } G\}|.$$

A choosing function has four immediate properties.

1. G is g-choosable.
2. choice $\#G \leq \max g(j)$
3. $g(j) \leq j$ $\qquad 1 \leq j \leq n$
4. $g(j) \leq 1+G$ valence j

THEOREM

Choice $\#G$ + choice $\#\overline{G} \leq n+1$

PROOF

Label the nodes $1,2,\ldots,n$ in such a way that G valence $i \geq$ G valence j if $i \leq j$. Let g be the choosing function which results from that labeling. Let $\overline{g}$ be the reverse in $\overline{G}$ defined by $\overline{g}(i) = 1+|\{j:1 \leq i < j \leq n, \text{ and } \{i,j\} \text{ is an edge of } \overline{G}\}|$.

Properties 1., 2., and 4. still hold for $\overline{g}$ and $\overline{G}$, while property 3. becomes $\overline{3}$. $\overline{g}(i) \leq n+1-i$.

Observe that, because of the special labeling, G valence j + $\overline{G}$ valence $i \leq n-1$, whenever $j \geq i$. When $j \leq i$, we have $g(j)+\overline{g}(i) \leq j+n+1-i \leq n+1$. When $j \geq i$, we have $g(j)+\overline{g}(i) \leq 1+G$ valence j + 1 + $\overline{G}$ valence $i \leq n+1$.

Hence $\max g(j) + \max \overline{g}(i) \leq n+1$.

The proof is finished by property 2. .

OPEN QUESTION

Does there exist $\xi > 0$ such that for large n,

$n^{\frac{1}{2}+\xi} <$ choice $\#G$ + choice $\#\overline{G}$?

GRAPH CHOOSABILITY IS NP-HARD

This result is due to A.L. Rubin. For background, and the terminology of Π_2^p-completeness, please refer to the book, Computers and Intractability by Michael R. Garey and David S. Johnson.

To show that graph choosability is Π_2^p-complete, we now describe a logical statement which is prototypical in a Π_2^p-complete class. This will be encoded into a graph G with a function f, such that the logical statement is true iff G is f-choosable.

The logical statement is this:

$$\forall_{U_1} \cdots \forall_{U_k} \exists_{U_{k+1}} \cdots \exists_{U_r} (C_1 \wedge C_2 \wedge \cdots \wedge C_m)$$

where each C_i is of the form $(X_{i1} \vee X_{i2} \vee X_{i3})$ and each X_{ij} is U_s or $\overline{U}_s$.

The basic ideas of constructs for the graph involve "propagators", "half-propagators", "multioutput propagators", and "initial graphs", with some nodes designated as input nodes, and some nodes designated as output nodes. In the following pictures a number on a node will be the value f takes on that node when G is formed. The value on an in node will be acquired when it gets merged with an out node.

HALF-PROPAGATOR

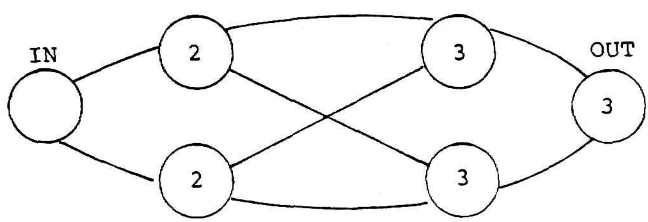

PROPERTIES

0. A 2-coloration will give the out node opposite color to the in node.

1. For any choice of a letter from the in node, and no matter what letters are put on nodes other than the in node, there is a compatible choice of letters from the remaining nodes of the half-propagator.

2. For any assignment of letters to nodes other than the in node, for any choice of a letter from the out node, there is at most one choice of letter incompatible with it on the in node. (This is a direct consequence of $\Theta_{2,2,2}$ being 2-choosable)

3. There is an assignment of letters, and a choice of in letter, such that only one choice of letter from the out node is compatible with it.

PROPAGATOR

It can be made by merging the out node of any half-propagator with the in node of any other half-propagator.

PROPERTIES

0. A 2-coloration will give the out node the same color as the in node.

1. Same as for half-propagator.

2. For any assignment of letters, there is at most one choice of letter for the in node which is incompatible with some letter on the out node. (Every other letter on the in node is compatible with all letters on the out node.)

3. Same as for half-propagator.

MULTIOUTPUT PROPAGATOR

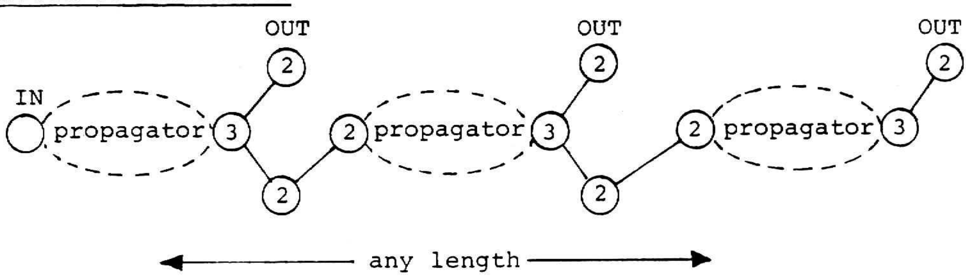

PROPERTIES

0. A 2-coloration will give the out nodes opposite color to the in node.

1. Same as for half-propagator.

2. For any assignment of letters, there is at most one choice of letter for the in mode which is incompatible with some combination of letters from the out nodes. (Every other letter in is compatible with all letters out.)

3. There is an assignment of letters, and a choice of in letter, such that only one combination of choices from the out nodes is incompatible with it.

INITIAL GRAPHS

A "∃-graph" is:

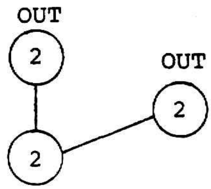

PROPERTIES

0. A 2-coloration gives all out nodes the same color.

1. For all assignments of letters, take any letter from any out node, and find that a compatible choice exists for the rest of the ∃-graph.

2. There is an assignment of letters such that any letter from any out node is compatible with only one letter from the other out node.

A "∀-graph" is:

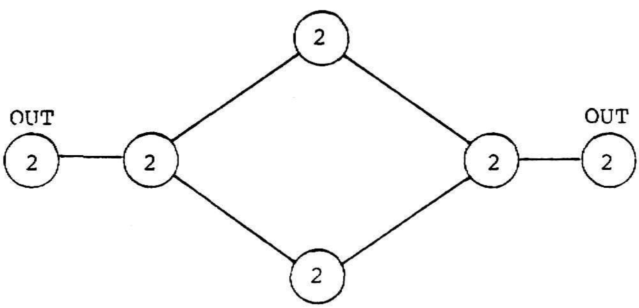

PROPERTIES

0. A 2-coloration gives out nodes the same color.

1. For any assignment of letters, at least one of the out nodes has the property that for either choice of letter from the out node, there is a compatible choice for the rest of the graph.

2. For either out node, there is a choice of letters such that only one choice of letter from that out node is compatible with a choice from the rest of the graph.

The graph G consists of the following.

For each i from 1 to k, we have a $\forall$-graph, with the out nodes named U_i and $\overline{U}_i$.

For each i from k+1 to r, we have a $\exists$-graph, with the out nodes named U_i and $\overline{U}_i$.

We think of the C_i's as clauses, and think of u_s and $\overline{u}_s$ as literals. For each literal v we connect a multioutput propagator to the node named V, identifying the in node of the propagator with V. All the multioutput propagators look alike having 3m output nodes, one for each ij where $1 \le i \le m$ and $1 \le j \le 3$.

Now we add m new nodes (each with $f(C_i)=3$) named $C_1, C_2, \ldots C_m$. For each i from 1 to m, and each j from 1 to 3, connect C_i to the ij node of the multioutput propagator attached to the node named X_{ij}.

That describes the graph. Now the graph can be pruned by taking off nodes which have valence less than f value. The final G will have valences only 2,3, or 4, and it will be bipartite (that is, 2-colorable). The f value at any node will be 2 or 3.

It should be verifiable from the properties of the constructs, that G is f-choosable iff the logical statement is true.

THE RANDOM BIPARTITE CHOICE NUMBER

Now we present a theorem which tells that there exist constants C_1 and C_2 such that an $m \times m$ random bipartite graph will have choice number between $C_1 \log m$ and $C_2 \log m$. The proof will be self contained, with the aid of a lemma.

Having fixed m top nodes, and m bottom nodes, let $R_{m,m}$ denote any one of the bipartite graphs whose edges constitute a subset of the m^2 possible top-to-bottom edges. We think of $R_{m,m}$ as having been chosen at random. Also we think of a "txt" as any pair consisting of a t-subset of top nodes and a t-subset of bottom nodes.

LEMMA

Suppose $t \geq \frac{2 \log m}{\log 2}$, and let $\bar{E}$ be the event that an $R_{m,m}$ has an empty induced subgraph on at least one txt.

Then $\bar{E}$ has probability $< \frac{1}{(t!)^2}$.

PROOF OF LEMMA

The number of possible $R_{m,m}$'s is 2^{m^2}. The number of txt's is $\binom{m}{t}^2$. Each possible edge empty txt is contained in $2^{m^2-t^2}$ of the $R_{m,m}$'s. Thus, the number of $R_{m,m}$'s which contain at least one empty txt is $< 2^{m^2-t^2}\binom{m}{t}^2$.

Therefore $\bar{E}$ has probability $< \dfrac{2^{m^2-t^2}\binom{m}{t}^2}{2^{m^2}} = \dfrac{\binom{m}{t}^2}{2^{t^2}}$.

With $m \leq 2^{\frac{t}{2}}$, we calculate as follows.

$$\frac{\binom{m}{t}^2}{2^{t^2}} \leq \frac{\binom{2^{t/2}}{t}^2}{2^{t^2}} \leq \frac{\left(\left(\frac{2^{t/2}}{t!}\right)^t\right)^2}{2^{t^2}} = \frac{1}{(t!)^2}.$$

THEOREM

Suppose $\frac{\log m}{\log 6} > 121$, and $t = \left\lceil \frac{2 \log m}{\log 2} \right\rceil$.

Then with probability $>1-\frac{1}{(t!)^2}$, we have

$$\frac{\log m}{\log 6} < \text{choice } \# \, R_{m,m} < \frac{3\log m}{\log 6} \quad .$$

PROOF

For the upper bound, we know from the discussion of $N(2,k)$ that if $2^{k-3} < m \le 2^{k-2}$, then choice $\#K_{m,m} \le k$. This tells us that choice $\#R_{m,m} \le$ choice $\#K_{m,m} < \frac{\log m}{\log 2} + 3 < \frac{3\log m}{\log 6}$.

To derive the lower bound, let $k = \left\lfloor \frac{\log m}{\log 6} \right\rfloor > 120$. Using the fact that $e^k > k^k/k!$, and a calculator if necessary, we obtain:

$$m \ge 6^k > 7k^2 2^k e^k > 7k^2\left(\frac{2^k k^k}{k!}\right) > 7k^2 \binom{2k-1}{k} > t\cdot k\cdot\binom{2k-1}{k} .$$

Harmlessly supposing $m = t\cdot k\cdot\binom{2k-1}{k}$ we next describe an assignment of letters the adversary can use to show that almost all $R_{m,m}$ have choice number $> k$. $\binom{2k-1}{k}$ is the number of k-subsets of letters from $\{1,2,\ldots,2k-1\}$. Each k-subset is put on $k\cdot t$ of the top nodes, and likewise on the bottom nodes. Now consider what must happen when a choice is attempted.

First we argue that on top there must be $\ge k$ letters, each chosen from $\ge t$ nodes. Because otherwise, if $\le k-1$ letters were chosen $\ge t$ times each, we could look at a k-subset of remaining letters. That k-subset was put on $k\cdot t$ top nodes, therefore one of the letters in it must have been chosen $\ge k\cdot t/k = t$ times.

Similarly there must be $\ge k$ letters, each chosen from $\ge t$ bottom nodes. But then, since only $2k-1$ letters were used, there must be one letter simultaneously chosen from t top nodes and t bottom nodes. The attempted choice fails if this txt has an edge. Now, according to the lemma, in almost all possible $R_{m,m}$'s every txt does have an edge. Thus the lemma tells us that choice $\#R_{m,m} > k$, with probability $> 1-\frac{1}{(t!)^2}$. In other words, we have proved the lower bound:

with probability $>1-\frac{1}{(t!)^2}$, we have choice $\#R_{m,m} > \frac{\log m}{\log 6}$.

THE RANDOM COMPLETE GRAPH - OPEN QUESTIONS

We do not know good bounds for the choice number of the random complete graph. Having fixed n nodes, let R_n denote any one of the graphs whose edges constitute a subset of the $\binom{n}{2}$ possible edges. We think of R_n as having been chosen at random, and look for bounds $L(n)$ and $U(n)$ for which we can prove that

$$L(n) < \text{choice } \#R_n < U(n) \quad ,$$

with probability $\to 1$ as n gets large.

From the known bounds for χR_n, we know that there exists a constant c such that $\frac{cn}{\log n} \leq L(n)$. On the upper side we merely know that $U(n) \leq \frac{n}{2}$.

Thus a specific open problem is to prove that with probability $\to 1$, $\frac{\text{choice } \#R_n}{n} \to 0$ as $n \to \infty$.

It would be even better to find good bounds for $K_{m*r} = K_{m,m,\ldots,m}$ where the number of nodes is $n=rm$, and m is about the size of log n.

We do know that choice $\#K_{3*r} > \frac{4}{3}r+c$.

The only one of this kind for which we know the exact value is K_{2*r}, which may be of interest because it is the only example we have whose proof uses the P. Hall theorem.

THEOREM

Choice $\#K_{2*r} = r$.

PROOF

Starting at $r=2$, we already know that choice $\#K_{2*2} = 2$. (it is the 4-cycle).

To induct, suppose $r > 2$, and suppose we know the theorem for all cases $< r$. Let the adversary put r letters on every node. If some letter is on both nodes of a nonadjacent pair, we can choose that letter from both nodes of that pair, and delete it from all other nodes. We can complete the choice by induction in this case.

Otherwise every nonadjacent pair has a disjoint pair of sets of letters. Any union of $\le r$ of the sets of letters on nodes will have $\ge r$ letters. Any union of $> r$ of the sets will have $\ge 2r$ letters, because it will include a disjoint pair. The conditions for the P. Hall theorem are satisfied, so there exists a system of distinct representatives. That is, the choice exists.

Here are some more specific numbers which are easily proved.

$K_{k-1,m}$ is k-choosable for all m, all $k \ge 2$.

$$K_{k,m} \quad \text{is} \begin{cases} \text{k-choosable for } m < k^k \\ \text{not k-choosable for } m \ge k^k. \end{cases}$$

PLANAR GRAPHS

Since every planar graph has a node of valence ≤ 5, it follows easily that every planar graph is 6-choosable. Perhaps some mathematicians, who are dissatisfied with the recent computer proof of the 4-color theorem, still sense that there are some things we ought to know, but do not yet know, about the structure of planar graphs. Here we offer two conjectures which may incidentally add interest to that exploration.

CONJECTURE

Every planar graph is 5-choosable.

CONJECTURE

There exists a planar graph which is not 4-choosable.

QUESTION

Does there exist a planar bipartite graph which is not 3-choosable?

Here is a graph which is planar, and 3-colorable, but not 3-choosable.

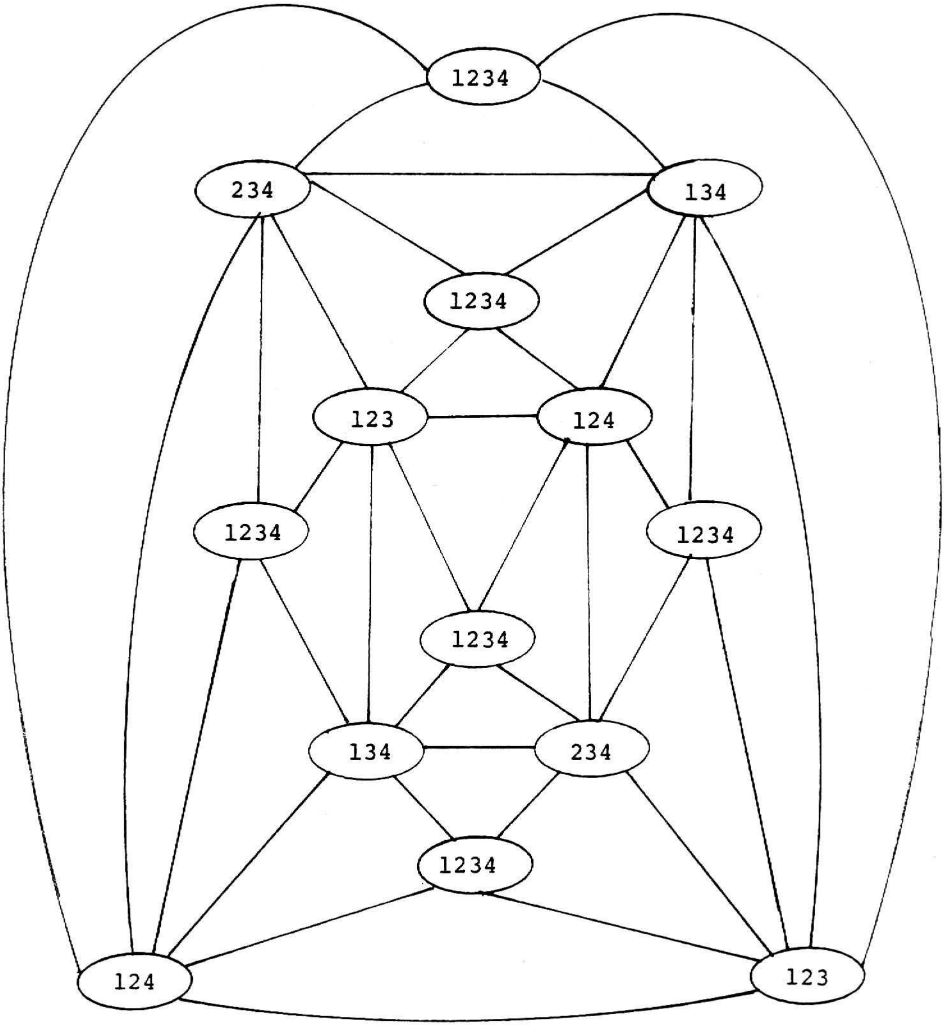
1234
234
134
1234
123
124
1234
1234
1234
134
234
1234
124
123

(a:b)-CHOOSABILITY

Suppose a is the number of distinct letters on each node, put there by the adversary, and we want to choose a b-subset from each node, keeping the chosen subsets disjoint whenever the nodes are adjacent. G will be (a:b)-choosable if such a choice can be made no matter what letters the adversary puts.

In terms of (a:b)-choosability we can say that if there does exist a planar bipartite graph which is not 3-choosable, it will have been a very close call in the following sense. If $a/b < 3$, then there exists a planar bipartite graph which is not (a:b)-choosable. In fact $K_{2,\binom{a}{b}^2}$ will be not (a:b)-choosable.

OPEN QUESTION

If G is (a:b)-choosable, does it follow that G is (am:bm)-choosable?

OPEN QUESTION

If G is (a:b)-choosable, and $\frac{c}{d} > \frac{a}{b}$, does it follow that G is (c:d)-choosable?

COMPOSITION LEMMA

Suppose H is obtained from G by adding edges. Let S be the subgraph consisting of those edges and their nodes.

If S is (d:a)-choosable, and G is (a:b)-choosable, then H is (d:b)-choosable.

PROOF

Let the adversary put d letters on each node. First make a choice consisting of an a-subset from each node, with disjoint a-subsets on S-adjacent nodes. This first choice can be made because S is (d:a)-choosable. Next make a choice consisting of a b-subset from each node's a-subset, with disjoint b-subsets on G-adjacent nodes. This second choice can be made because G is (a:b)-choosable. The resulting choice makes the b-subsets disjoint on adjacent nodes of H. Thus H is (d:b)-choosable.

COROLLARY

If H is not 2k-choosable, and G is obtained from H by erasing disjoint edges, then G is not k-choosable.

Here is just one more theorem - a direct consequence of the fact that for given k and g there exists a family F of k-sets with the following three properties.

1. F does not have property B
2. For any two distinct $X, Y \in F$, $|X \cap Y| \leq 1$.
3. The smallest cycle has length $> g$, in the graph which has nodeset $= F$, with an edge between nodes X and Y iff $|X \cap Y| = 1$.

For a proof of the above fact, please refer to P. Erdos and A. Hajnal, "On Chromatic Numbers of Graphs and Set Systems", Acta Math. Acad. Sci. Hungar., 17 (1966), pp. 61-99.

THEOREM

For given k and g, there exists a bipartite graph G such that the smallest cycle in G has length $> g$, and choice $\#G > k$.

PROOF

Let F be a family of 2k-sets with properties 1., 2., 3. above. Let H be the bipartite graph having the 2k-sets of F as top nodes, and likewise as bottom nodes, with an edge between a top node X and a bottom node Y iff $X \cap Y \neq \phi$.

First observe that H is not 2k-choosable, because F does not have property B. Any choice including one from each top node would use all the letters belonging to some 2k-set on the bottom.

Next obtain G by erasing those edges of H which connect two nodes having the same 2k-set. Thus G will inherit from F the property of having smallest cycle length $> g$.

The corollary to the composition lemma tells us that G is not k-choosable.

ACKNOWLEDGEMENT

It got started when we tried to solve Jeff Dinitz's problem. Jeffrey Dinitz is a mathematician at Ohio State University. At the 10^{th} S.E. Conference on Comb., Graph Theory, and Computing at Boca Raton in April 1979 he posed the following problem.

Given a m x m array of m-sets, is it always possible to choose one from each set, keeping the chosen elements distinct in every row, and distinct in every column?

To the best of our knowledge Jeff Dinitz problem remains unsolved for $m \geq 4$.