Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 15, pp. 210-214, July 2006

ELA
http://math.technion.ac.il/iic/ela

ON MINIMAL RANK OVER FINITE FIELDS*

GUOLI DING† AND ANDREĬ KOTLOV‡

Abstract. Let F be a field. Given a simple graph G on n vertices, its *minimal rank* (with respect to F) is the minimum rank of a symmetric $n \times n$ F-valued matrix whose off-diagonal zeroes are the same as in the adjacency matrix of G. If F is finite, then for every k, it is shown that the set of graphs of minimal rank at most k is characterized by finitely many forbidden induced subgraphs, each on at most $(\frac{|F|^k}{2}+1)^2$ vertices. These findings also hold in a more general context.

Key words. Minimal rank, Forbidden induced subgraph, Critical graph.

AMS subject classifications. 05C50, 05C75, 15A03, 15A33.

1. Introduction. Let F be a field. Given a simple graph G on n vertices, we say that a symmetric $n \times n$ F-valued matrix *represents* G if its off-diagonal zeroes are the same as in the adjacency matrix of G. The *min rank* [w.r.t. F] of G is the minimum rank [over F] of the matrices representing G. The reader is referred to [1] for motivation and importance of min rank.

For a fixed integer k, denote $\mathfrak{G}_k$ the class of all graphs of min rank at most k. It is obvious that $\mathfrak{G}_k$ is closed under vertex deletion. Call a graph *k-critical* if it is minimal (w.r.t. vertex deletion) of rank larger than k. Clearly, $\mathfrak{G}_k$ is characterized by the k-critical graphs as forbidden induced subgraphs.

It is a celebrated result of Robertson and Seymour [3] that any class of graphs closed under taking *minors* is characterized by a *finite* set of forbidden minors. At the 2005 Oberwolfach graph theory workshop, Hein van der Holst asked if there are only finitely many k-critical graphs for any k and F. (Barrett, van der Holst, and Loewy [2, 1] had recently confirmed this for $k \leq 2$ and any F.) When F is finite, we answer this question affirmatively, by providing an upper bound on the size of a k-critical graph. In fact, our arguments hold in a more general context. In the next section, we define k-critical graphs w.r.t. an arbitrary collection $\mathfrak{M}_k$ of matrices of rank at most k and show that the number of such graphs is finite as long as there is $c \in \mathbb{N}$ such that each matrix in $\mathfrak{M}_k$ uses at most c distinct elements from its field. If, in addition, $\mathfrak{M}_k$ is closed under row-and-column-duplication, the number of k-critical graphs can be bounded in terms of c and k. We derive such bounds in Section 3.

2. Definitions and Main Result. For a fixed $k \in \mathbb{N} \cup \{0\}$, let $\mathfrak{S}_k$ denote the set of all pairs (M, F) where F ranges over all fields and M is an F-valued symmetric matrix with $\mathrm{rk}_F(M) \leq k$. Fix some non-empty $\mathfrak{M}_k \subseteq \mathfrak{S}_k$. Denote by $\mathfrak{G}_k$ the set of graphs represented by matrices in $\mathfrak{M}_k$. As before, call a graph G *k-critical* if

*Received by the editors 03 November 2005. Accepted for publication 12 July 2006. Handling Editor: Raphael Loewy.

†Mathematics Department, Louisiana State University, Baton Rouge, LA 70803, USA. Partially supported by NSA grant H98230-05-1-0081, NSF grant ITR-0326387, and AFOSR grant F49620-03-1-0238-0239-0241 (ding@math.lsu.edu).

‡Ridgefield, CT 06877, USA (andrei@kotlov.com).

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 15, pp. 210-214, July 2006

$G \notin \mathfrak{G}_k$ while every proper induced subgraph of G is in $\mathfrak{G}_k$. It is easy to see that $\mathfrak{G}_k$ is characterized by [k-critical] graphs as forbidden induced subgraphs iff it is closed under vertex deletion.

For a matrix A, denote by $c(A)$ the number of distinct entries in A, and for any non-empty $\mathfrak{A} \subseteq \mathfrak{S}_k$, set $c(\mathfrak{A}) := \sup\{c(A) : (A, F) \in \mathfrak{A}\}$.

OBSERVATION 2.1. *A matrix A of rank r has at most $c(A)^r$ distinct rows.*

Proof. W.l.o.g. assume that the first r columns of A are linearly independent. Then two rows of A are identical iff they agree on the first r coordinates. □

Given disjoint graphs G and H, by *replacing a vertex v of G with H* we mean, as is common, the disjoint union of $G - v$ and H plus the edges $\{uw : uv \in G, w \in H\}$. If a matrix M represents G and the last $t \geq 2$ rows of M are identical then G can be obtained from the graph $G - \{n - t + 2, \ldots, n\}$ by replacing vertex $n - t + 1$ with either the clique K_t (if $M_{nn} \neq 0$) or the independent set $\overline{K}_t$ (if $M_{nn} = 0$). For $m \in \mathbb{N}$, call a graph *m-sprawling* if it can be obtained from a graph on at most m vertices by replacing each vertex with either a clique or an independent set.

COROLLARY 2.2. *Let $c := c(\mathfrak{M}_k) < \infty$. Then every graph in $\mathfrak{G}_k$ is c^k-sprawling. Consequently, every k-critical graph is $(c^k + 1)$-sprawling.* □

LEMMA 2.3. *Given $m \in \mathbb{N}$, an infinite sequence of distinct m-sprawling graphs has an infinite subsequence ascending w.r.t. induced-subgraph inclusion.*[1]

Proof. Fix an infinite sequence, $\boldsymbol{s}$, of [distinct] m-sprawling graphs. As there are finitely many graphs on at most m vertices, $\boldsymbol{s}$ has an infinite subsequence, $\boldsymbol{s}'$, of graphs which can be sprawled from the same graph, G, on some $n \leq m$ vertices. Further, as there are finitely many (namely, 2^n) choices of whether to use a clique or an independent set at each vertex of G, $\boldsymbol{s}'$ contains an infinite subsequence $\boldsymbol{s}''$ of graphs H_i ($i \in \mathbb{N}$) for which such choices coincide. Now, each H_i can be described by a string of n natural numbers $(a_{i1}, ..., a_{in})$ where a_{ij} is the number of vertices by which vertex j of G was replaced in obtaining H_i. Notice that an infinite sequence of natural numbers contains a monotone non-decreasing infinite subsequence. By sequentially applying this argument n times, we find that $\boldsymbol{s}''$ has an infinite subsequence, $\{H_{i_t} : t \in \mathbb{N}\}$, such that, for each $j = 1, ..., n$, the sequence $\{a_{i_t j}\}$ is monotone non-decreasing. But then, the sequence $\{H_{i_t}\}$ itself is monotone non-decreasing under induced-subgraph inclusion. □

COROLLARY 2.4. *If $c := c(\mathfrak{M}_k) < \infty$, then there are finitely many k-critical graphs.*

Proof. An infinite sequence of distinct k-critical (and thus, $(c^k + 1)$-sprawling) graphs would contain an ascending [infinite] subsequence, contrary to the definition of k-criticality. □

[1] A partially ordered set whose every infinite sequence of distinct elements has an increasing subsequence is called *well-partially-ordered*, cf. e.g. [4]. In this terminology, the lemma becomes: *the set of m-sprawling graphs is well-partially-ordered w.r.t. induced-subgraph inclusion.*

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 15, pp. 210-214, July 2006

We summarize our findings as follows.

THEOREM 2.5. *If $c(\mathfrak{M}_k) < \infty$ and $\mathfrak{G}_k$ is closed under vertex deletion, then $\mathfrak{G}_k$ is characterized by finitely many k-critical graphs as forbidden induced subgraphs.*

3. Upper Bound. In the previous section, we saw that once $c(\mathfrak{M}_k) < \infty$, the number of k-critical graphs is finite. However, in the absence of other requirements, this number need not be bounded by any function of c and k. Indeed, fix some $N \in \mathbb{N}$ and let $\mathfrak{M}_1$ consist of all the rank-at-most-1 symmetric 0-1 matrices of size less than $N \times N$. Then the N graphs $\{K_a \dot\cup \overline{K}_{N-a} : a = 1, ..., N\}$ (the disjoint unions of the a-clique and $N - a$ independent vertices) are 1-critical.

In this section, we bound the number of k-critical graphs in terms of c and k under the following additional constraint on $\mathfrak{M}_k$:

(∗) *If $(M, F) \in \mathfrak{M}_k$ then also $(M', F) \in \mathfrak{M}_k$ where M' has two identical rows i and j and M can be obtained by deleting row j and column j in M'.*

THEOREM 3.1. *If (∗) holds and $c := c(\mathfrak{M}_k) < \infty$, then $|G| \leq (\frac{c^k}{2} + 1)^2$ for a k-critical G.*

We precede the proof of Theorem 3.1 by the following discussion. Two vertices of a graph G are *twins* if they have the same neighbors in the rest of G (the two vertices themselves can be either adjacent or independent). It is easy to see that a twin relationship is transitive: if u and v are twins, and v and w are twins, then u and w are also twins and, moreover, the triple $\{u, v, w\}$ spans either a clique or an independent set. (We will call such a triple of pairwise-twin vertices a *triplet*.) Thus, the twin relationship induces a partition, $T(G)$, of the vertex set of G into *twin classes*.

Fix a vertex $v \in G$ and consider a partition $T' := T(G) - v$ of the vertex set of $G - v$. It is easy to see that: (a) T' is a refinement of $T(G - v)$; (b) more specifically, every set in $T(G - v)$ is a [disjoint] union of one or two sets in T'; (c) in particular, if G has no twins then $G - v$ has no triplet; (d) if v has exactly one twin, v', in G then T' and $T(G - v)$ coincide except perhaps on the twin class of v' in $G - v$; (e) if v has at least two twins in G then $T' = T(G - v)$. In addition, (f) if G has no triplet then v can be chosen so that $G - v$ has no triplet either. Indeed, if G is twin-free then (f) is trivially true by (c). Thus, assume G is not, and let v and v' be twins in G. Then, by (d), the only triplet $G - v$ may have is of the form $\{v', u, u'\}$, where u and u' are twins in G. Further, the subgraph induced by v, v', u, u' has either no edge except vv' or all edges except vv', and v, v', u, u' have the same neighbors in the rest of G. But then, $T(G) - u = T(G - u)$, as required.

Proof of Theorem 3.1: Clearly, we may assume that $\mathrm{rk}_F(M) = k$ for some $(M, F) \in \mathfrak{M}_k$. Suppose first that $k \geq 2$. Then necessarily $c \geq 2$. If G has no triplet then, by (f) and Observation 2.1, $n - 1 \leq 2c^k \leq (1 + \frac{c^k}{4})c^k = (\frac{c^k}{2} + 1)^2 - 1$, as required. Hence, we can assume w.l.o.g. that the last $t \geq 3$ vertices of G form its largest twin class, C. By criticality of G, there is $(M, F) \in \mathfrak{M}_k$ such that M represents $G - n$. Consider the diagonal entry M_{ii} of M for some $i > n - t$. If C is an independent set in G and $M_{ii} = 0$ or if C is a clique and $M_{ii} \neq 0$ then the symmetric $n \times n$ matrix M'

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 15, pp. 210-214, July 2006

ELA
http://math.technion.ac.il/iic/ela

whose rows i and n are identical and whose upper-left $(n-1)\times(n-1)$ corner is M, represents G and, by $(*)$, $(M', F) \in \mathfrak{M}_k$. This is a contradiction. Denoting by D the lower-right $(t-1)\times(t-1)$ corner of M, we thus have:

(†) *either D has no off-diagonal zero and its diagonal is all-zero, or D is all-zero except its diagonal has no zero.*

In particular, no two among the last $t-1$ rows of M are identical. Hence, by Observation 2.1, among the first $n-t$ rows of M, at most $c^k-(t-1)$ are distinct, whence at least $\frac{n-t}{c^k-(t-1)}$ are identical. The latter is a lower bound on the size of a twin class in $G-n$ and thus, by (e), also on t. We have $\frac{n-t}{c^k-(t-1)} \le t$ whence $n \le t(c^k+2-t) \le \left(\frac{c^k+2}{2}\right)^2$, as required.

It remains to consider the case $k \le 1$. If $k=0$ then $G = K_2$, and if $k=1$ and $c=1$ then $G=\overline{K}_2$, both in accord with what claimed. Thus assume that $k=1$ and $c \ge 2$ but, contrary to the claim, $|G| \ge 5$. Notice that G contains neither the path P_3 nor $K_2\dot{\cup}K_2$ as an induced subgraph as neither is in $\mathfrak{G}_1$. Hence, the non-isolated vertices of G, if any, form a clique. But then, in the above notation, $t \ge 3$ and, by (†), the lower-right 2×2 corner of M is of F-rank 2. This is a contradiction. □

We can improve the bound of Theorem 3.1 by a more diligent accounting. For the sake of brevity, we only establish the asymptotic version of such an improvement.

Theorem 3.2. *Under the assumptions of Theorem 3.1, $|G| = O(2^k + (c-1)^{2k})$.*

Proof. We utilize the notation of the proof of Theorem 3.1. As remarked in that proof, $n = O(c^k)$ if $t \le 2$, in accord with our claim. Thus, assume $t \ge 3$.

Set $r := \mathrm{rk}_F(D)$ and let, w.l.o.g., the last r columns of D be linearly independent. As in Observation 2.1, we can bound the number of distinct rows in M by considering its last r columns. Namely, among the first $n-t$ rows of M, the distinct ones come from at most c^{k-r} distinct rows whose last r components are all-zero and at most $c^{k-r}(c-1)^r$ distinct rows whose last r components have no zero. Proceeding as in the proof of Theorem 3.1, we find that $n \le t(c^{k-r}((c-1)^r+1)+1)$. Further, the number $t-1$ of [distinct] rows of D is, again by (†), at most $(c-1)^r + r$, whence

$$(\ddagger) \qquad n \le ((c-1)^r + r + 1)(c^{k-r}((c-1)^r+1)+1).$$

If $c \ge 3$ then (‡) implies $n = O(c^{k-r}(c-1)^{2r})$ which, in turn, implies $n = O((c-1)^{2k})$, and if $c=2$, (‡) becomes $n \le (r+2)(2^{k-r+1}+1) = O(2^k)$—all in accord with the claim. □

4. Concluding Remarks. The following improvements can be made to the argument for Theorem 3.2 should one desire to derive an exact (rather than asymptotic) bound there. Below, we assume that $(*)$ holds and G is a k-critical graph.

1. If a twin-class C of G is an independent set then, by (†), $|C| \le k+1$. Moreover, $|C| = k+1$ may hold only if the non-neighbors of C form an independent set themselves: this is because the direct sum $A \oplus B$ of matrices A and B satisfies $\mathrm{rk}(A\oplus B) = \mathrm{rk}(A) + \mathrm{rk}(B)$. In particular, G has at most $k-1$ isolated vertices unless $G = \overline{K}_k \dot{\cup} K_m$ for some $m \in \mathbb{N}$.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 15, pp. 210-214, July 2006

ELA
http://math.technion.ac.il/iic/ela

2. Similarly, if a twin-class C of G is a clique and $c = 2$ then, in the notation of the proof of Theorem 3.2, $r \leq |C| - 1 \leq r + (c-1)^r = r + 1$ and it is easy to see that $|C| - 1 = r + 1$ (equivalently, that the first row of D is a linear combination of the rest) iff $\mathrm{char}(F)$ divides r; further, $r \leq k$ and the equality may hold only if the non-neighbors of C form an independent set.
3. More generally, let $f(r, c)$ denote an upper bound on the size of an F-valued symmetric matrix D described by (†) such that $c(D) \leq c$ and $\mathrm{rk}_F(D) = r$. In the proof of Theorem 3.2, we implicitly set $f(r, c) := (c-1)^r + r$ and concluded that $n = O(f(r, c)c^{k-r}(c-1)^r)$. Thus, a tighter f, should one exist, would lead to a tighter bound in Theorem 3.2. (In the previous remark then, we showed how to improve $f(r, 2)$ in particular cases.)
4. Finally, in the notation of the proof of Theorem 3.2, the upper bound on the number of distinct rows among the first $n - t$ rows of M can be reduced by $t - 1 - r$, the number of those distinct rows with no zero (or no non-zero) in the last r positions which are already used among the last $t - 1$ rows of M.

REFERENCES

[1] Wayne Barrett, Hein van der Holst, and Raphael Loewy. Graphs whose minimal rank is two. *Electronic Journal of Linear Algebra*, 11:258–280, 2004.
[2] Wayne Barrett, Hein van der Holst, and Raphael Loewy. Graphs whose minimal rank is two: the finite fields case. *Electronic Journal of Linear Algebra*, 14:32–42, 2005.
[3] Neil Robertson and P. D. Seymour. Graph Minors XX: Wagner's conjecture. *Journal of Combinatorial Theory* Series B, 92:325–357, 2004.
[4] Wikipedia, `http://en.wikipedia.org/wiki/Well_partial_order`.