A Removal Lemma for Systems of Linear Equations over Finite Fields

Daniel Král'[*] Oriol Serra[†] Lluís Vena[‡]

Abstract

We prove a removal lemma for systems of linear equations over finite fields: let $X_1, \ldots, X_m$ be subsets of the finite field $\mathbb{F}_q$ and let A be a $(k \times m)$ matrix with coefficients in $\mathbb{F}_q$ and rank k; if the linear system $Ax = b$ has $o(q^{m-k})$ solutions with $x_i \in X_i$, then we can destroy all these solutions by deleting $o(q)$ elements from each X_i. This extends a result of Green [Geometric and Functional Analysis 15(2) (2005), 340–376] for a single linear equation in abelian groups to systems of linear equations. In particular, we also obtain an analogous result for systems of equations over integers, a result conjectured by Green. Our proof uses the colored version of the hypergraph Removal Lemma.

1 Introduction

In 2005, Green [4, Theorem 1.5] proved the so-called Removal Lemma for abelian groups. It roughly says that if a linear equation over an abelian group has not many solutions one can delete all the solutions by removing

[*]Institute for Theoretical Computer Science (ITI), Faculty of Mathematics and Physics, Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic. ITI is supported as project 1M0545 by the Czech Ministry of Education. E-mail: kral@kam.mff.cuni.cz.

[†]Departament de Matemàtica Aplicada IV, Universitat Politècnica de Catalunya. Supported by the Catalan Research Council under project 2005SGR0258. E-mail: oserra@ma4.upc.edu.

[‡]Departament de Matemàtica Aplicada IV, Universitat Politècnica de Catalunya. Supported by the Spanish Research Council under project MTM2005-08990-C01-C02. E-mail: lvena@ma4.upc.edu.

few elements. This Removal Lemma for groups has its roots in the well–known Triangle Removal Lemma of Ruzsa and Szemerédi [10] (see also [7] and [8] for generalizations and many applications of this important result in combinatorics) which roughly says that if a certain graph has not many triangles, then they are supported over not many edges.

In [5], the authors gave a purely combinatorial proof, by using the Removal Lemma for graphs, of the algebraic version of the Removal Lemma for linear equations. This allows for an extension of the result to non-abelian groups. In the same paper, the authors considered some extensions of the result to systems of equations in abelian and non–abelian groups which could be proved along the same lines. However to extend the result to general linear systems, the graph representation used in the mentioned paper presented serious limitations. Instead, the extensions to hypergraphs of the removal lemma, which have been recently proved by Nagle, Rödl, Schacht [9], Gowers [3] or Tao [13], seem to be the natural tool to achieve this goal.

Our main result is the following:

Theorem 1 (Removal Lemma for systems of equations). *Let $F = \mathbb{F}_q$ be the finite field of order q. Let $X_1, \dots, X_m$ be subsets of F, A a $(k \times m)$ matrix with coefficients in F whose rank is k and b a k-dimensional vector over F.*

If there are $o(q^{m-k})$ solutions of the system $Ax = b$ with $x_i \in X_i$, then there exist sets $X'_1, \dots, X'_m$ with $|X_i \setminus X'_i| = o(q)$ such that there is no solution to the system $Ax = b$ with $x_i \in X'_i$.

So, if a linear system has not many solutions, then these solutions are supported by not many elements. Equivalently, Theorem 1 asserts that for every $\varepsilon > 0$, k and m, there exists $\delta > 0$ such that if the number of solutions is at most δq^{m-k}, all the solutions can be destroyed by removing at most εq elements from each of the sets X_i. The value of δ depends only on ε, k and m, in particular, it is independent of q.

As an example of application of Theorem 1, we state the following Corollary.

Corollary 2. *If a subset $X \subseteq \mathbb{F}_q$, $q = p^n$, $p \geq k \geq 3$, has $o(q^2)$ arithmetic progressions of length k, then the set A has $o(q)$ elements.*

Theorem 1 implies a result conjectured by Green [4, Conjecture 9.4]:

Theorem 3. *Let k and m be integers with $k \leq m$ and let A be an $k \times m$ matrix of rank k. Let $X \subseteq [N]$, and suppose that there are $o\left(N^{m-k}\right)$ vectors in X^m for which $Ax = 0$. Then $X = B \cup C$, where B has no solution $Ax = 0$ with $x \in B^m$ and $|C| = o(N)$.*

To prove Theorem 3 one can apply Theorem 1 with $X_i = X$ and a convenient large enough prime $c(A) \cdot N \leq p \leq 2c(A) \cdot N$, where $c(A)$ is a constant that depends on the matrix A.

Independently of us, Conjecture 9.4 from [4] was proved by Shapira [11] whose method also yields a different proof of Theorem 1. Shapira's proof also reduces the problem to finding an appropriate representation of the system to a hypergraph and uses the colored version of the hypergraph Removal Lemma (Theorem 4) as our proof does. However, his proof involves $O(m^2)$-uniform hypergraphs and our proof involves $(k+1)$-uniform hypegraphs. At the high level, the two proofs are similar, but the particular ideas used to represent systems by hypergraphs differ.

When the system is reduced to one equation our construction coincides with the one in [5], thus it can be viewed as its natural generalization.

We note that Theorem 1 might also be derived from the main result in Szegedy [12]. There the author proves a Symmetry Removal Lemma and describes a framework to apply it to Cayley Hypergraphs. Theorem 1 would follow once the conditions of validity within this setting are properly verified.

Let us also mention that the conclusion of Theorem 1 can be proven substantially easier if we assume that every k columns of the matrix are linearly independent; we have reported on this result in [6]. Also, independently of us, it was obtained by Candela [2].

Our proof of Theorem 1 follows the main idea of the one presented in [5]. As we have already mentioned, we use the edge-colored version of the hypergraph Removal Lemma which follows from a more general result of Austin and Tao [1, Theorem 2.1].

Theorem 4 (Austin and Tao [1]). *Let H be an edge-colored $(k+1)$-uniform hypergraph with m vertices. Let $K = (V, E)$ be an edge colored $(k+1)$-uniform t-partite hypergraph with M vertices. If the number of copies of H in K (preserving the colors of the edges) is $o(M^m)$, then there is a set $E' \subseteq E$ of size $|E'| = o(M^{k+1})$ such that $K' = (V, E \setminus E')$, as an edge-colored hypergraph, is H-free.*

2 Main result and its proof

In this section, we present the proof of Theorem 1 except for an auxiliary lemma (Lemma 5) whose proof is given in the next section. We first explain the main steps of the proof.

2.1 Outline of the proof

We associate to the system $Ax = b$, where A has size $k \times m$, an edge-colored $(k+1)$-uniform hypergraph H with m edges and m vertices. We shall construct a large m-partite $(k+1)$-uniform hypergraph K on mq vertices. Its vertex set is composed of m copies, $F_1, F_2, \ldots, F_m$, of the field $F = \mathbb{F}_q$. The edges of K are defined in such a way that each solution of the system corresponds to q^k edge–disjoint copies of H with each edge representing an element of X_i. The bound on the number of solutions of our linear system translates to the fact that K contains $o(q^m)$ copies of H. By the Removal Lemma for hypergraphs, Theorem 4, we will find a set E' of edges with size $o(q^{k+1})$, such that, by removing E' from K we delete all copies of H. We then apply a pigeonhole argument to find $o(q)$ elements to be removed from each set X_i based on the $o(q^{k+1})$ edges given by the Removal Lemma; this argument uses the fact that the q^k copies corresponding to the same solution are edge–disjoint.

The key point above is the construction of the auxiliary hypergraph. Before we explain the details of this construction, we show that we can assume the matrix A to be of a certain special form.

2.2 Reductions of the system

First observe that, by the nature of the statement of Theorem 1, there is no loss of generality in assuming that the matrix A has full rank k. By permuting the columns and an appropriate choice of a basis of the vector space F^m, the matrix A can be assumed to be of the form $A = (I_k|B)$ where I_k is the identity matrix.

Our second observation is that it suffices to prove Theorem 1 for homogenous systems. Indeed, if A is written in the form $(I_k|B)$, the general statement follows by applying it to the system $Ax = 0$ once we replace the given k sets $X_1, \ldots, X_k$ by $X_1' = X_1 - b_1, \cdots, X_k' = X_k - b_k$, where $b = (b_1, \ldots, b_m)$.

We will make a further assumption on $A = (I_k|B)$: any two rows in B have rank 2. Suppose on the contrary that rows B_i and B_j are not linearly independent, say $B_i = \lambda B_j$. This implies that every solution of the system $Ax = 0$ satisfies $x_i = \lambda x_j$. Therefore we can replace X_i by $X'_i = X_i \cap (\lambda \cdot X_j)$, delete the j-th equation together with the j-th variable and apply our theorem in the resulting setting: the obtained system contains one less equation and one less variable.

We may also assume that any row B_i has, at least, two non-zero entries. Otherwise the i'th equation would read $x_i + b_{i,j}x_j = 0$ for some $j \in [k+1, \ldots, m]$. As in the preceding paragraph, we can replace the set X_j by $X'_j = X_j \cap \left(-b_{i,j}^{-1} \cdot X_i\right)$ and consider the system obtained by eliminating the i-th equation and the i-th variable.

Consequently, we can assume that A and b satisfy the following:

(i) $b = 0$.

(ii) The matrix A has the form $A = (I_k|B)$ where I_k is the identity matrix.

(iii) Every two rows of B are linearly independent.

(iv) Each row of B has at least two non-zero entries.

Notice that the condition (iv) implies that $m \geq k+2$.

2.3 Construction of the hypergraph K

For the construction of the hypergraph K with the properties given in Subsection 2.1, we shall use an auxiliary matrix associated to the matrix A which is described in Lemma 5. Before we state the lemma, let us introduce some additional notation. If M is a matrix, the i-th row of a matrix M is denoted by M_i and M^j denotes its j-th column. The support of a vector $x \in F^n$, denoted by $s(x)$, is the set of coordinates with a nonzero entry.

Lemma 5. *Let $A = (I_k|B)$ be a $(k \times m)$-matrix with coefficients in $\mathbb{F}_q$. There are an $(m \times m)$ matrix C and m pairwise distinct $(k+1)$-subsets $S_1, \ldots, S_m \subseteq [1, m]$ with the following properties:*

1. *$AC = 0$*

2. *$\mathrm{rank}(C) = m - k$ (maximal subject to the first condition).*

3. *For every i, $s(C_i) \subseteq S_i$ and $i \in s(C_i)$.*

4. For every i, there exists a subset $S'_i \subseteq S_i$ with $|S'_i| = k$ such that the set of columns $\{C^j,\ j \in [1,\ldots,m] \setminus S'_i\}$ has rank $m-k$.

The proof of Lemma 5 is postponed to Section 3.

We are ready to define a suitable hypergraph representation of the linear system along the lines described in Subsection 2.1. Let $Ax = 0$ be a linear system, where A is a $k \times m$ matrix with entries in F satisfying the properties (i)–(iv) at the end of Subsection 2.2. Let C be the matrix associated to A and $S_1, \ldots, S_m$ be the $(k+1)$–subsets of $[1,m]$ satisfying the properties stated in Lemma 5.

First, the hypergraph H is the $(k+1)$-uniform edge-colored hypergraph with vertex set $\{1,2,\ldots,m\}$ and with edges $S_1, S_2, \ldots, S_m$, where the edge S_i is colored i.

The hypergraph K is the $(k+1)$-uniform m-partite hypergraph with the vertex set $F \times [1,m]$. For every $x \in X_i$, K contains an edge $\{(a_j, j), j \in S_i, a_j \in F\}$ if and only if

$$\sum_{j \in S_i} C_{ij} a_j = x,$$

and this edge is colored by i and labelled by x. Since the support $s(C_i)$ is nonempty and $|S_i| = k+1$, K contains precisely q^k edges colored by i and labelled by x for each $x \in X_i$.

2.4 Proof of Theorem 1

We now show that the hypergraphs K and H have the properties given in the outline of the proof.

Claim 1. *If H' is a copy of H in K, then $x = (x_1, \ldots, x_m)$ is a solution of the system, where x_i is the label of the edge colored by i in H'.*

Proof. The copy H' has an edge of each color and is supported over m vertices. By Lemma 5 (3) we have $i \in S_i$ for each i which implies $\cup_{i=1}^m S_i = [1,m]$. Hence the vertex set of H' is of the form $\{(a_1,1),(a_2,2),\ldots,(a_m,m)\}$. By the construction of K, it holds that $Ca = x$ where $a = (a_1, a_2, \ldots, a_m)$. Hence, $0 = ACa = Ax$ and x is a solution of the system. $\square$

Claim 2. *For any solution $x = (x_1, \ldots, x_m)$ of the system $Ax = 0$ with $x_i \in X_i$, there are precisely q^k edge–disjoint copies of the edge–colored hypergraph H in the hypergraph K.*

Proof. Fix a solution $x = (x_1, \dots, x_m)$ of $Ax = 0$ with $x_i \in X_i$, $1 \le i \le m$. First, we show that there is a copy of H in K in which the edge colored i has label x_i, $1 \le i \le m$.

Since the matrix C has rank $m - k$ and satisfies $AC = 0$, the columns in C spans the solution space in F^m and thus there is a vector $u = (u_1, \dots, u_m)$ with $x = Cu$. In particular,

$$x_i = (C_i, u) = \sum_{j=1}^{m} C_{ij} u_j = \sum_{j \in S_i} C_{ij} u_j,$$

where the second equality follows from Lemma 5 (3). Therefore, for every i, the set $\{(u_j, j), j \in S_i\}$ is an edge of K colored i and labeled x_i. It follows that the edges $\{(u_j, j), j \in S_i\}$, $i = 1, \dots m$, span a copy of H in K. Since the kernel of C is k-dimensional, there are q^k vectors u satisfying $x = Cu$, and each of them corresponds to a copy of H in K. We next verify that these q^k copies are edge–disjoint.

Let $e = \{(a_j, j), j \in S_i\}$ be an edge of K colored by i and labeled $x_i \in X_i$. We show that all the q^k copies of H in K contain different edges colored by i and labelled x_i. By Lemma 5 (4), there is a subset $S'_i \subseteq S_i$ of size k such that $\{C^j, j \notin S'_i\}$ is a set of $m - k$ linearly independent solutions of the system $Ax = 0$. Hence, we may find a vector $u = (u_1, \dots, u_m)$ with $x = Cu$ such that $u_j = a_j$ for each $j \in S'_i$. With this choice, we must also have $u_j = a_j$ for each $j \in S_i$ and the copy of H associated to this u contains the edge e. Hence, for each edge colored i and labeled x_i there is a copy of H associated to x in K which contains this edge.

Since there are q^k such edges and the same number of copies of H associated to the solution x, no two copies can share the same edge colored i and labelled x_i. By applying the same argument to each of the colors $1, \dots, m$, we conclude that the q^k copies of H associated to the solution x are edge–disjoint. □

We now proceed with the proof of Theorem 1.

Proof of Theorem 1. Since the number of solutions is $o(q^{m-k})$, by Claims 1 and 2, K contains $o(q^m)$ copies of H. By the Removal Lemma for colored hypergraphs (Theorem 4), there is a set E' of edges of K with size $o(q^{k+1})$ such that, by deleting the edges in E' from K, the resulting hypergraph is H-free.

The sets X'_i are constructed as follows: if E' contains at least q^k/m edges colored with i and labelled with x_i, remove x_i from X_i. In this way, the total number of elements removed from all the sets X_i together is at most $m \cdot o(q) = o(q)$. Hence, $|X_i \setminus X_i'| = o(q)$ as desired. Assume that there is still a solution $x = (x_1, x_2, \ldots, x_m)$ with $x_i \in X_i'$. Consider the q^k edge–disjoint copies of H in K corresponding to x. Since each of these q^k copies contains at least one edge from the set E' and the copies are edge–disjoint, E' contains at least q^k/m edges with the same color i and the same label x_i for some i. However, such x_i should have been removed from X_i. $\square$

3 Proof of Lemma 5

In this section, we give an effective construction of the matrix C with the properties stated in Lemma 5.

We first define a sequence $\mathcal{B}_1, \ldots, \mathcal{B}_m$ of basis of the column space of A which is formed by columns of A. For a base $\mathcal{B}_i$, $T_i = \{j \in [1, m] : A^j \in \mathcal{B}_i\}$ denotes the set of indexes of the columns of A contained in $\mathcal{B}_i$.

We set $T_k = [1, k]$, i.e., $\mathcal{B}_k = \{A^1, \ldots, A^k\}$. Since no row of the submatrix B is the zero vector, we may assume, up to reordering the columns from B, that $A_{1,k+1} \neq 0$.

Suppose that $\mathcal{B}_i$ has been defined for some $k \leq i < m$. Express each vector A^j, $i+1 \leq j \leq m$ in the basis $\mathcal{B}_i$, i.e., $A^j = \sum_{r \in T_i} \lambda_{r,j} A^r$. Let $g(i+1)$ be the smallest r such that $\lambda_{r,j} \neq 0$ for some $j \in [i+1, m]$. Without loss of generality we may assume that the last $m - i$ columns of A are ordered in such a way that $j = i+1$, that is, $\lambda_{g(i+1),i+1} \neq 0$. The base $\mathcal{B}_{i+1}$ is then obtained from $\mathcal{B}_i$ by replacing the column $A^{g(i+1)}$ with A^{i+1}, i.e., T_{i+1} if builded from T_i by deleting $g(i+1)$ and adding $i+1$:

$$T_{i+1} = (T_i \setminus \{g(i+1)\}) \cup \{i+1\}.$$

In particular, if $i = k$, $T_{k+1} = [2, k+1]$ since $A_{1,k+1} \neq 0$.

We have now defined the bases $\mathcal{B}_k, \ldots, \mathcal{B}_m$. Observe that $1 = g(k+1) < \cdots < g(m)$. Moreover, $T_i \subseteq [1, i]$ for $k \leq i \leq m$. We further set $\mathcal{B}_0 = \mathcal{B}_m$ for our convenience.

Suppose that $\mathcal{B}_i$ is defined for some $0 \leq i < k$. We proceed to define $\mathcal{B}_{i+1}$ in a similar manner. Let $g(i+1)$ be the smallest index in $T_i \setminus [1, i]$ such that the corresponding coefficient of the vector A^{i+1} expressed in the base $\mathcal{B}_i$ is non-zero. Note that $g(i+1)$ is well-defined since the vectors $A^1, \ldots, A^{i+1}$ are linearly independent. The base $\mathcal{B}_{i+1}$ is obtained from $\mathcal{B}_i$ by replacing $A^{g(i+1)}$ with A^{i+1}. In particular, for $1 \leq i \leq k$, it always holds

$\{1,\ldots,i\} \subseteq T_i$. Moreover, the base $\mathcal{B}_k$ defined in this way coincides with our original choice of it.

An intuitive way of understanding the basis $\mathcal{B}_i$ and sets T_i is as follows: the set T_i is the lexicographically maximal decreasing sequence of k indices from $[i-m+1,i]$ (indices are taken modulo m) such that the columns with these indices are linearly independent. Let us prove this claim formally:

Claim 3. *The set T_i viewed as a subset of $[i-m+1,i]$ and decreasingly ordered is the lexicographically maximal subset of $[i-m+1,i]$ such that the columns A^j, $j \in T_i$, generate the column space of A (indices are taken modulo m).*

Proof. We prove the claim by induction for $i = k,\ldots,m+k-1$. The claim holds if $i = k$ as $T_k = \{1,\ldots,k\}$ and this is the lexicographically maximal subset of $[k-m+1,k]$. Let $T'_i = \{t'_1,\ldots,t'_k\}$ be the lexicographically maximal subset of $[i-m+1,i]$ such that $A^{t'_1},\ldots,A^{t'_k}$ are linearly independent and let $T_i = \{t_1,\ldots,t_k\}$. For simplicity, assume $t_1 > t_2 > \ldots > t_k$ and $t'_1 > t'_2 > \ldots > t'_k$.

Let j be the first index such that $t_j \neq t'_j$. Clearly, $t'_j > t_j$. Since $t_1 = t'_1 = i$, it holds $j > 1$. Since $A^{t'_2},\ldots,A^{t'_j}$ are linearly independent, $\{t'_2,\ldots,t'_{j-1}\} = \{t_2,\ldots,t_{j-1}\} \subseteq T_{i-1}$ and T_{i-1} is the lexicographically maximal k-subset of $[i-m,i-1]$ corresponding to linearly independent columns of A, T_{i-1} must contain an element r that is larger or equal to t'_j. Clearly, $g(i) = r$ (otherwise, $r \in T_i$).

If $r = t'_j$, then the vectors $A^{t'_1}, A^{t'_2},\ldots,A^{t'_j}$ are not linearly independent (as $g(i) = r$ so that A^r is in the span of $A^{t_1},\ldots,A^{t_{j-1}}$). If $r > t'_j$, then the sets of vectors A^{t_1}, $A^{t_2},\ldots,A^{t_{j-1}}$ and $A^{t_2},\ldots,A^{t_{j-1}}$, A^r span the same linear space which implies that the vectors $A^{t_2},\ldots,A^{t_{j-1}}$, A^r, $A^{t'_j}$ are linearly independent. Consequently, $T_{i-1} = \{t_2,\ldots,t_k\} \cup \{r\}$ is not lexicographically maximal independent k-subset of $[i-m,i-1]$ (recall that $t'_2 = t_2$, $\ldots$, $t'_{j-1} = t_{j-1}$ and $t'_j > t_j$). ☐

The next claims will be needed to check that the matrix C which we define has the properties given in Lemma 5.

Claim 4. *The function $g : [1,m] \to [1,m]$ is bijective.*

Proof. In the described construction, i can only be inserted to T_j in order to build T_{j+1} if $j+1 = i$. If $g(r) = g(s) = i$ for a pair of distinct r and s (which involves deleting i twice in the process), then an element i would be deleted twice from T_j but inserted only once which is impossible because we start and end with the same base. $\square$

Claim 5. *For every $i = 1, 2, \ldots, m$, the set T_{i-1} does not contain i. Moreover, for every $i = 2, \ldots, k+1$ the set T_{i-2} does not contain i.*

Proof. The claim holds by construction if $i \geq k+1$. Hence, we have to focus on $i \in [1, k]$. Since the columns A^j, $j \in [1, m] \setminus \{i\}$ span the column space of A (by assumption (iv) on A), $i \notin T_{i-1}$ by Claim 3.

To prove the second part of the statement, observe that, by assumption (iii) on A applied to rows $i-1, i$ with $2 \leq i \leq k$, the columns A^j, $j \in [1, m] \setminus \{i-1, i\}$ span the column space of A. Again by Claim 3, we also have $i \notin T_{i-2}$.

In the extremal case $i = k+1$ we also use assumption (iv) on A to ensure that there is $j > k+1$ such that $A^1, \ldots, A^{k-1}, A^j$ is a base. Again by Claim 3, $k+1 \notin T_{k-1}$. $\square$

We can now define the matrix C. The j-th column of C has its support in $T_{j-1} \cup \{j\}$. For $i \in T_{j-1}$, the entry C_{ij} is the coefficient of A^i in the expression of A^j in the base $\mathcal{B}_{j-1}$:

$$A^j = \sum_{i \in T_{j-1}} C_{ij} A^i,$$

and $C_{jj} = -1$ (recall that, by Claim 5, we have $j \notin T_{j-1}$.)

Clearly, each column of C belongs to the space of solutions of the system $Ax = 0$, so that Lemma 5 (1) holds.

The submatrix of C formed by the last $m-k$ columns and the last $m-k$ rows is an upper triangular matrix with nonzero entries on the diagonal which implies that the rank of C is $m-k$. This proves Lemma 5 (2)

We next define the family $\{S_1, \ldots, S_m\}$ of $(k+1)$-subsets of $[1, m]$. By the definition of the function g and of the matrix C, the nonzero elements in the j-th column of C are in the rows $[1, j] \cup [g(j), m]$ if $j \in [1, k]$ and in the rows $[g(j), j]$ if $j \in [k+1, m]$. Let $R \subseteq [1, m] \times [1, m]$ be the set of subscripts defining this area, i.e, $(i, j) \in R$ if and only if either $j \in [1, k]$ and $i \in [1, j] \cup [g(j), m]$ or $j \in [k+1, m]$ and $i \in [g(j), j]$ (see Figure 3 for a typical portrait of R.)

When reading off this area in the matrix by rows we get the subsets S_i, namely,

$$S_i = \begin{cases} g^{-1}([1, i]) \cup [i, k], & i \in [1, k] \\ g^{-1}(T_i) \cup \{i\}, & i \in [k+1, m]. \end{cases}$$

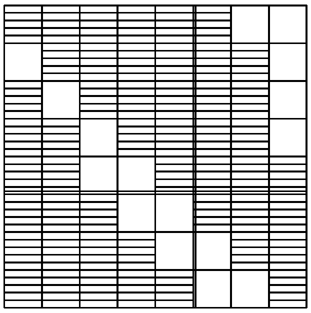

Figure 1: An example of the area R in matrix C which corresponds to the permutation $g(1,2,3,4,5,6,7,8) = (3,4,6,7,8,1,2,5)$.

By the definition of g, the support of the row C_i is contained in S_i for every $i \in [1,m]$ and none of the rows is zero (the entry in the main diagonal is -1).

Let us observe that $|S_i| = k+1$. It follows from the definition of g that $g(i) \notin T_i$. For $i \in [1,k]$, the sets $g^{-1}([1,i])$ and $[i,k]$ are disjoint as we now argue: $g^{-1}([1,i])$ can be divided into two parts $g^{-1}([1,i]) \cap [k+1,m]$ and $g^{-1}([1,i]) \cap [1,k]$. We show that $g^{-1}([1,i]) \cap [1,k]$ does not intersect $[i,k]$ as it is clear that the other part of $g^{-1}([1,i])$ and $[i,k]$ are disjoint. Claim 5 implies that $j \notin T_{j-1}$, for all j, which means that all the elements in $[1,i] \cap T_0$ have been replaced with elements in $[1,i-1]$ before building T_{i-1}. This shows that $g^{-1}([1,i])$ does not intersect $[i,k]$. Since g is a bijection (Claim 4) the sets S_i have cardinality $k+1$ for each $i \in [1,m]$.

Let us now show that the sets S_i are pairwise distinct. Recall that the region R contains in a column $j \in [1,k]$ the rows $[1,j] \cup [g(j),m]$. It follows from the second part of Claim 5 that $j \notin T_{j-2}$ for $j = 2,...,k+1$, which implies $g(j-1) > j$. Hence S_j does not contain $j-1$ but it does contain j. On the other hand, the column $j \in [k+1,m]$ contains in the region R the rows $[g(j),j]$, so again S_j contains j but does not contain $j-1$.

Let $j < j'$. If $j' \leq k$ then $\{j'-1, j'\} \subseteq [j,k] \subseteq S_j$, which implies $S_j \neq S_{j'}$. If $j' > k$ then, either $j' \notin S_j$ or, as g is increasing in $[k+1,m]$, $\{j'-1, j'\} \subseteq S_j$, which again implies $S_j \neq S_{j'}$.

In order to prove the last part of Lemma 5, we show that the columns $\{C^j,\ j \notin S_i\}$ form a set of $m-k-1$ linearly independent vectors. Together with Lemma 5 (3) this fact implies Lemma 5 (4) and completes the proof of the Lemma.

Let $C' = \{C^j : j \notin S_i\}$ be the submatrix of C formed by the columns with indices not in S_i. We divide this matrix into four parts: the upper left $UL = \{C_{rs} : r < i, s \in [1,i] \setminus S_i\}$ formed by the first $i-1$ rows of C and the columns with index at most i, the upper right $UR = \{C_{rs} : r < i, s \in [i+1,m] \setminus S_i\}$ formed by the same rows and the remaining columns, the lower right $LR = \{C_{rs} : r \geq i, s \in [1,i] \setminus S_i\}$ formed by the last $m-i+1$ rows and the columns with index at most i and the lower left $LR = \{C_{rs} : r \geq i, s \in [i+1,m] \setminus S_i\}$ with the remaining entries.

By our construction of the matrix C, UR is an all-zero matrix, while, as discussed in the proof of Lemma 5 (2), the columns C^j with $j \in [i+1,m] \setminus S_i$ are linearly independent because the columns C^j, $j \in [k+1,m]$, are linearly independent. On the other hand, again by the construction of C, UL is an upper triangular matrix (maybe with the steps higher than one). It follows that the columns of C' are linearly independent. The proof of Lemma 5 is now finished.

Acknowledgments

We would like to thank Balázs Szegedy, Vojta Rödl and Mathias Schacht for helpful discussions and comments.

References

[1] T. Austin and T. Tao, On the testability and repair of hereditary hypergraph properties, arxiv0801.2179v1.

[2] P. Candela, On systems of linear equations and uniform hypergrpahs, manuscript, 2008.

[3] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. Submitted, preprint avaliable at: `http://arxiv.org/abs/0710.3032v1`

[4] B. Green. A Szemerédi-type regularity lemma in abelian groups, with applications. *Geometric and Functional Analysis* 15(2) (2005), 340–376.

[5] D. Král', O. Serra, L. Vena. A combinatorial proof of the Removal Lemma for groups. *J. Combin. Theory Ser. A* (2008) (to appear).

[6] D. Král', O. Serra, L. Vena. A removal lemma for linear systems over finite fields. Proc. VI Jornadas Matemàtica Discreta y Algorítmica, Ediciones y Publicaciones de la UdL, 2008, 417–424.

[7] J. Komlós, M. Simonovits, Szemerédi's regularity lemma and its applications in graph theory. Combinatorics, Paul Erdős is eighty, Vol.2 (Keszthely, 1993), 295-352, Bolyai Soc. Math. Stud., 2, János bolyai Math Soc., Budapest, 1996.

[8] J. Komlós, A. Shokoufandeh, M. Simonovits, E. Szemerédi, The regularity lemma and its applications in graph theory, *Theoretical aspects of computer science* (Tehran, 2000), 84–112, Lecture Notes in Comput. Sci., **2292**, Springer, Berlin, (2002).

[9] B. Nagle, V. Rödl, M. Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures Algorithms* 28 (2006), no. 2, 113–179.

[10] I.Z. Ruzsa, E. Szemerédi. Triple systems with no six points carrying three triangles. *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely*, 1976), Vol. II, pp. 939–945, Colloq. Math. Soc. János Bolyai, 18, North-Holland, Amsterdam-New York, 1978.

[11] A. Shapira. A proof of Green's conjecture regarding the removal properties of sets of linear equations. Submitted, available as arXiv:0807.4901v2 [math.CO].

[12] B. Szegedy, The Symmetry Preserving Removal Lemma. manuscript, preprint available as arXiv:0809.2626.

[13] T. Tao. A variant of the hypergraph removal lemma. *J. Combin. Theory Ser. A*, **113** (2006), 1257–1280.