Generating All Invertible Matrices by Row Operations

Petr Gregor
Department of Theoretical Computer Science and Mathematical Logic,
Charles University, Prague, Czech Republic

Hung P. Hoang
Algorithm and Complexity Group, Faculty of Informatics, TU Wien, Austria

Arturo Merino
Institute of Engineering Sciences, Universidad de O'Higgins, Rancagua, Chile

Ondřej Mička
Department of Theoretical Computer Science and Mathematical Logic,
Charles University, Prague, Czech Republic

Abstract

We show that all invertible $n \times n$ matrices over any finite field $\mathbb{F}_q$ can be generated in a *Gray code* fashion. More specifically, there exists a listing such that (1) each matrix appears exactly once, and (2) two consecutive matrices differ by adding or subtracting one row from a previous or subsequent row, or by multiplying or dividing a row by the generator of the multiplicative group of $\mathbb{F}_q$. This even holds in the more general setting where the pairs of rows that can be added or subtracted are specified by an arbitrary transition tree that has to satisfy some mild constraints. Moreover, we can prescribe the first and the last matrix if $n \geq 3$, or $n = 2$ and $q > 2$. In other words, the corresponding flip graph on all invertible $n \times n$ matrices over $\mathbb{F}_q$ is Hamilton connected if it is not a cycle. This solves yet another special case of Lovász conjecture on Hamiltonicity of vertex-transitive graphs.

2012 ACM Subject Classification Mathematics of computing → Combinatorics; Mathematics of computing → Graph theory

Keywords and phrases Hamilton cycle, combinatorial Gray code, invertible matrices, finite field, general linear group, generation algorithms

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2024.35

Related Version Full proofs are provided in the preprint version of this extended abstract.
arXiv Version: `https://arxiv.org/abs/2405.01863` [8]

Funding This work was supported by Czech Science Foundation grant GA 22-15272S.
Hung P. Hoang: Austrian Science Foundation (FWF, project Y1329 START-Programm)
Arturo Merino: This work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 850979).

Acknowledgements This work was initiated at the 2nd Combinatorics, Algorithms, and Geometry workshop in Dresden, Germany in 2022. We would like to thank the organizers and participants for the inspiring atmosphere.

1 Introduction

Combinatorial generation is one of the most basic tasks we can perform on combinatorial objects and a key topic in Volume 4A of Knuth's seminal series *The Art of Computer Programming* [11]. In this task, we are given an implicit description of the objects and need to produce a listing of all objects fitting the description, with each object appearing exactly once. The goal is to develop an algorithm that can generate these objects at a fast rate.

35th International Symposium on Algorithms and Computation (ISAAC 2024).
Editors: Julián Mestre and Anthony Wirth; Article No. 35; pp. 35:1–35:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

If consecutive objects produced by a generation algorithm differ by large changes, the algorithm must spend a lot of time updating its data structures. Therefore, a natural first step towards creating an efficient generation algorithm is to ensure that consecutive objects differ by only a *small change*. Such a listing is known as a **(combinatorial) Gray code**; see Mütze's survey [13] for many Gray codes of various objects. In addition to combinatorial generation, Gray codes are also relevant in the field of combinatorial reconfiguration, which examines the relationships between combinatorial objects through their local changes; see, e.g., Nishimura's recent introduction on reconfiguration [14].

In this paper, we study Gray codes for invertible matrices over a finite field. A natural attempt for enumerating all invertible $n \times n$ matrices over a finite field $\mathbb{F}_q$, is to choose any nonzero first row and then selecting the following rows to be independent to the previous rows. However, this attempt is not efficient as it requires multiple checks for independence to generate even a single matrix. Furthermore, consecutive matrices in this listing may differ in multiple rows. Instead, we focus on generating matrices in a Gray code order, i.e., every matrix is obtained from the previous one by a single elementary row operation. We note that generating invertible matrices with specific properties has applications in cryptography, e.g., in McEliece cryptosystems [7].

1.1 Strong Lovász conjecture

All invertible $n \times n$ matrices over $\mathbb{F}_q$ with matrix multiplication form the **general linear group** $GL(n, q)$. Each elementary row operation can be represented by multiplying on the left by a matrix that corresponds to this row operation. Hence, we are interested in finding a Hamilton path in an (undirected) Cayley graph on $GL(n, q)$ generated by the allowed row operations, which is in turn an instance of Lovász conjecture [12] on the Hamiltonicity of vertex-transitive graphs.[1]

Stronger versions of Lovász conjecture have been considered in the literature. For example, Dupuis and Wagon [6] asked which non-bipartite vertex-transitive graphs are not Hamilton connected. A graph is **Hamilton connected** if there is a Hamilton path between any two vertices. Similarly, they asked which bipartite vertex-transitive graphs are not Hamilton laceable [6]. A bipartite graph is **Hamilton laceable** if there is a Hamilton path between any two vertices from different bipartite sets. Note that the bipartite sets must be of equal size, which is true for all vertex-transitive bipartite graphs except K_1.

▶ **Conjecture 1** (Strong Lovász conjecture). *For every finite connected vertex-transitive graph G it holds that G is Hamilton connected, or Hamilton laceable, or a cycle, or one of the five known counterexamples.*

The five known counterexamples are the dodecahedron graph, the Petersen graph, the Coxeter graph, and the graphs obtained from the latter two by replacing each vertex with a triangle. The dodecahedron graph is a non-bipartite vertex-transitive graph that has a Hamilton cycle, but it is not Hamilton connected [6]. The other four well-known counterexamples are non-bipartite vertex-transitive graphs that do not admit a Hamilton cycle. Note that except when $G \in \{K_1, K_2, C_3, C_4\}$ the cases in the conjecture are mutually exclusive.

There are many results in line with Conjecture 1. Particularly relevant to us is a result of Tchuente [18] showing that the Cayley graph of the symmetric group S_n, generated by any connected set of transpositions, is Hamilton laceable when $n \geq 4$. Another relevant example is

[1] A graph is **vertex-transitive** if its automorphism group acts transitively on the vertices.

Chen and Quimpo's Theorem [4] showing that all Abelian Cayley graphs satisfy Conjecture 1. Nevertheless, Conjecture 1 remains open even for Cayley graphs of the symmetric group with every generator an involution [16]. Note that none of the five counterexamples to Conjecture 1 is a Cayley graph, leading to Cayley graph variants of Conjecture 1 (e.g., [15]).

1.2 Row operations

Our aim when generating all invertible matrices by row operations is to restrict the allowed operations as much as possible. Note that for $q > 2$ we must allow row multiplications by some scalar to be able to generate all 1×1 matrices. Thus, we allow row multiplications by a fixed generator α of the multiplicative group of nonzero elements of $\mathbb{F}_q$. We will also allow row multiplication by α^{-1}; i.e., division by α, to have an inverse operation for an undirected version of the problem. Furthermore, we specify allowed row additions and subtractions by a directed **transition graph** T on the vertex set $[n]$, where $[n] := \{1, \ldots, n\}$. An edge $(i, j) \in E(T)$ specifies that we can add to or subtract from the j-th row the i-th row. Then, each allowed row operation above corresponds to the left multiplication by a corresponding matrix from a set $\mathrm{ops}(T)$, formally defined by (2).

Observe that to generate all invertible matrices by the allowed operations, the transition graph T *must* be strongly connected; see Lemma 3 below. For our main result we require the following stronger condition.

▶ **Definition 1** (Bypass transition graph)**.** A transition graph T on the vertex set $[n]$ is a **bypass transition graph** if either (i) $n = 1$, or (ii) $n \geq 2$ and

- there exist an edge (i, n) and an edge (n, j) for some $i, j \in [n-1]$, and
- the graph $T - n$ obtained by removing n from T is also a bypass transition graph.

In other words, a bypass transition graph is obtained from a single vertex 1 by repeatedly adding a directed path (a '*bypass*') from some vertex i to some vertex j via a new vertex n. An example of a transition graph with the above property is the one comprised by edges $(i, i+1)$ and $(i+1, i)$ for all $i \in [n-1]$; i.e., a bidirectional path. In the language of row operations, this corresponds to allowing row additions or subtractions between any two consecutive rows. It can be easily seen by induction that a bypass transition graph is strongly connected.

1.3 Our results

For any integer $n \geq 1$, a finite field $\mathbb{F}_q$, and an n-vertex transition graph T we define the (undirected) Cayley graph

$$G(n, q, T) := \mathrm{Cay}(GL(n, q), \mathrm{ops}(T)),$$

where the set $\mathrm{ops}(T)$ is given by (2). Our main result is as follows.

▶ **Theorem 1.** *Let $n \geq 2$ be an integer and q be a prime power such that $q \geq 3$ if $n = 2$. Let T be an n-vertex bypass transition graph. Then the graph $G(n, q, T)$ is Hamilton connected.*

Note that for $n = 1$ the transition graph T has no edges, so $G(1, q, T)$ for any $q \geq 3$ is simply a $(q-1)$-cycle and $G(1, 2, T) = K_1$. For $n = q = 2$, we have that $T = (\{1, 2\}, \{(1, 2), (2, 1)\})$ is the only bypass transition graph, so $G(n, q, T)$ is a 6-cycle, which is not Hamilton connected. Thus, we may restate our result as follows.

▶ **Corollary 2.** *Let $n \geq 1$ be an integer and q be a prime power, and let T be an n-vertex bypass transition graph. Then the graph $G(n, q, T)$ is Hamilton connected unless it is a cycle.*

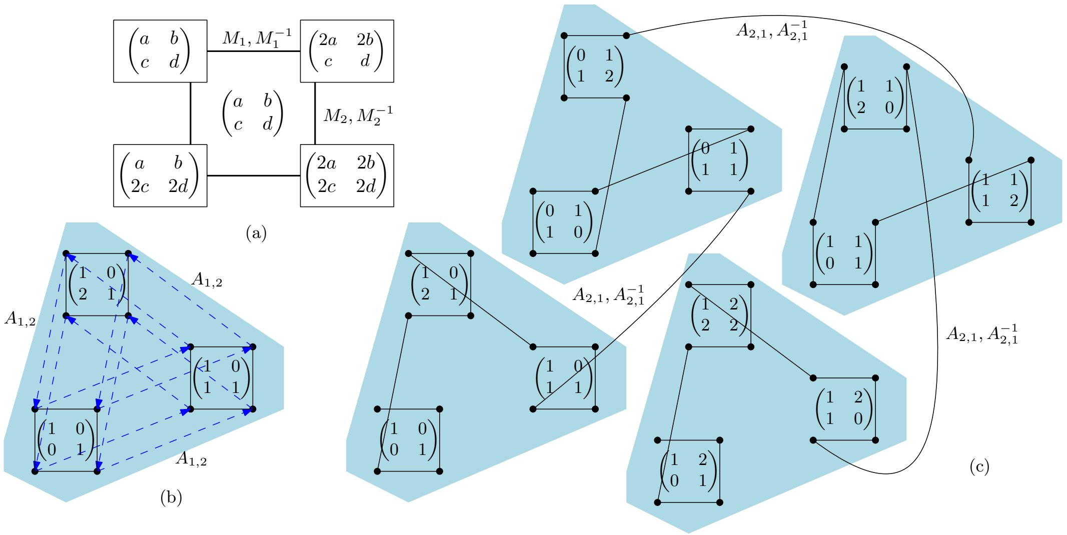
$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$
M_1, M_1^{-1}
$\begin{pmatrix} 2a & 2b \\ c & d \end{pmatrix}$
$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$
M_2, M_2^{-1}
$\begin{pmatrix} a & b \\ 2c & 2d \end{pmatrix}$
$\begin{pmatrix} 2a & 2b \\ 2c & 2d \end{pmatrix}$
(a)
$A_{1,2}$
$A_{1,2}$
$A_{1,2}$
(b)
$A_{2,1}, A_{2,1}^{-1}$
$A_{2,1}, A_{2,1}^{-1}$
$A_{2,1}, A_{2,1}^{-1}$
(c)

Figure 1 The part (c) illustrates a Hamilton path in the graph $G(2, 3, ([2], \{(1, 2), (2, 1)\}))$. Four vertices around a matrix Z are those obtained from Z by multiplying or dividing a row by α (which is 2 for $q = 3$); see the part (a). The part (b) shows the edges within a shaded component, where the black solid edges are row multiplications/divisions, while the (directed) dashed edges are additions from the first row to the second row. Note that the other directions of the latter edges indicate subtractions of the first row from the second row. Furthermore, while these shaded components exhibit a Cartesian product structure, the same does not hold for the whole graph.

This shows that the family of graphs $G(n, q, T)$ where T is a bypass transition graph is yet another example of a family of Cayley graphs satisfying Conjecture 1. A particularly interesting example is when T is a bidirectional path. See Figure 1 for an illustration for $n = 2$ and $q = 3$.

Moreover, we discuss how to turn the proof of Theorem 1 algorithmic in Section 7.

1.4 Related work

Permutations of $[n]$ can be represented as (invertible binary) permutation matrices forming a subgroup of $GL(n, 2)$. Thus, all the vast results on generating permutations such as in [17, 18] can be directly translated into the context of generating permutation matrices. In particular, there is a general permutation framework developed in [10] that allows us to generate many combinatorial classes by encoding them into permutations avoiding particular patterns. However, the row operations that we consider here do not preserve the subgroup of permutation matrices, so our results do not fall into this framework.

A related task to generation is random sampling. The construction of a random invertible $n \times n$ matrix over $\mathbb{F}_q$ is usually done by constructing a uniformly random matrix and checking whether it is non-singular. The success probability is lower-bounded by a constant independent of n but dependent on q (e.g., see [5] and the citations therein). Hence, there is only a constant factor overhead for random sampling of an invertible matrix over a finite field compared to that of a matrix over the same field. The latter task can be achieved, for example, by independently constructing each row (or column).

2 Preliminaries

The **(undirected) Cayley graph** of a group Γ with a generator set S is the graph $\mathrm{Cay}(\Gamma, S) := (\Gamma, \{\{x, sx\} : x \in \Gamma, s \in S\})$, assuming that S is closed under inverses and does not contain the neutral element. Note that we apply generators on the left as it is more natural for row operations on matrices.

The **general linear group** $GL(n, q)$ is the group of all invertible $n \times n$ matrices over the finite field $\mathbb{F}_q$ with matrix multiplication. Note that for $\mathbb{F}_q$ to be a field, q has to be a prime power. For example, $GL(1, 2)$ is the trivial group, $GL(2, 2) \simeq S_3$, and $GL(3, 2) \simeq PSL(2, 7)$ is also known as the group of automorphisms of the Fano plane. The number of elements in $GL(n, q)$ is $a_n := (q^n - 1)(q^n - q) \cdots (q^n - q^{n-1})$, which is obtained by counting choices for (nonzero) rows that are not spanned by the previous rows. It also satisfies the recurrence

$$a_n = (q^n - 1)q^{n-1}a_{n-1}, \tag{1}$$

for $n \geq 2$, and $a_1 = q - 1$ (i.e., the number of nonzero elements of $\mathbb{F}_q$).

By Gaussian elimination, the group $GL(n, q)$ can be generated by row additions and row multiplications by a scalar. As we consider the Cayley graph to be undirected, we also consider the inverse operations, which we call row subtractions and row divisions by a scalar. The formal definitions of these operations are as follows.

For $i \in [n] := \{1, \ldots, n\}$, let $r_i = r_i(A)$ denote the i-th row in A. For distinct $x, y \in [n]$, we denote by $A_{xy} = (a_{ij})$ the binary matrix with $a_{ij} = 1$ if and only if $i = j$, or ($i = y$ and $j = x$). Note that left multiplication by A_{xy} corresponds to adding the x-th row to the y-th row; i.e., the operation $r_x + r_y \to r_y$. Similarly, multiplication by A_{xy}^{-1} then corresponds to subtracting the x-th row to the y-th row; i.e., the operation $-r_x + r_y \to r_y$.

Let α be a generator of the multiplicative group of $\mathbb{F}_q$. For $x \in [n]$, we denote by $M_x = (a_{ij})$ the matrix with $a_{ij} = \alpha$ if $i = j = x$, $a_{ij} = 1$ if $i = j \neq x$, and $a_{ij} = 0$ otherwise. Left multiplication by M_x corresponds to multiplying the x-th row by α; i.e., the operation $\alpha r_x \to r_x$, and multiplication by M_x^{-1} corresponds to the inverse operation $\alpha^{-1} r_x \to r_x$ that we call **dividing** the x-th row by α. Note that for $q = 2$ the multiplicative group is trivial, that is, $M_x = I$, where I denotes the identity matrix.

A **transition graph** T is any directed graph on the vertex set $[n]$ with the edge set $E(T)$. For a transition graph T and a field $\mathbb{F}_q$ we define

$$\mathrm{ops}(T) := \{A_{ij}, A_{ij}^{-1} : (i, j) \in E(T)\} \cup \{M_i, M_i^{-1} : i \in [n]\}, \tag{2}$$

for $q > 2$, and $\mathrm{ops}(T) := \{A_{ij}, A_{ij}^{-1} : (i, j) \in E(T)\}$ for $q = 2$. In other words, $\mathrm{ops}(T)$ contains the row additions and subtractions induced by the edges of T, and all row multiplications and divisions by α if they are nontrivial. A directed graph is strongly connected if for any two vertices i, j, there is a directed path from i to j. A (strongly connected) component of a directed graph is a maximal induced subgraph that is strongly connected. We make the following observation, whose proof can be found in [8].

▶ **Lemma 3.** *For every transition graph T, the set $\mathrm{ops}(T)$ generates the group $GL(n, q)$ if and only if T is strongly connected.*

We denote by $\mathbb{F}_q^n$ the vector space of all n-tuples over the field $\mathbb{F}_q$. The span of $u_1, \ldots, u_k \in \mathbb{F}_q^n$ is denoted by $\langle u_1, \ldots, u_k \rangle$. Its orthogonal space $\langle u_1, \ldots, u_k \rangle^\perp$ is the kernel of the matrix with rows $u_1, \ldots, u_k$.

For $k \geq 3$ we denote by C_k a cycle on k vertices, and for $k \in \{1, 2\}$ we define C_k as the complete graph K_k. We also denote the path on k vertices by P_k for $k \geq 1$. The **Cartesian product** $G \,\square\, H$ of two graphs G and H is the graph with the vertex set $V(G) \times V(H)$ and the edge set $\{(u,v)(u',v) : uu' \in E(G), v \in V(H)\} \cup \{(u,v)(u,v') : u \in V(G), vv' \in E(H)\}$. For a graph G and a subset U of vertices, we denote by $G[U]$ the subgraph of G induced by U. Similarly, for a graph G and two subsets of vertices $U_1, U_2 \subseteq V$, we use $E[U_1, U_2]$ to denote the set of edges between U_1 and U_2, i.e., $E[U_1, U_2] = \{xy \in E : x \in U_1, y \in U_2\}$.

For an edge-colored graph, a trail in a graph is **alternating** if any two consecutive edges on the trail differ in color.

3 Joining lemma for Hamilton connectivity

In this section, we present a lemma that joins many Hamilton connected graphs into a larger one. This lemma seems quite versatile. Not only is it useful in our proof in the next section, but it also allows us to easily reprove several classical results on Hamilton connectivity, for example for the permutahedron [18].

▶ **Lemma 4** (Joining lemma). *Let G be a graph with the vertex set partitioned into $k \geq 2$ disjoint subsets $V_1, \ldots, V_k$ such that following conditions hold.*

(1) *$G[V_i]$ is Hamilton connected for every $i \in [k]$;*

(2) *Every vertex in every set V_i has a neighbor in some different set V_j;*

(3) *There are at least three pairwise disjoint edges between every two sets V_i, V_j.*[2]

Then G is Hamilton connected.

Proof. Let $x, y \in V$ be two vertices to be connected by a Hamilton path. First we consider the case when they are in different sets V_i. We can assume that $x \in V_1$ and $y \in V_k$, otherwise we rename the sets. We select vertices $x_i, y_i \in V_i$ for every $i \in [k]$ so that $x_1 = x$, $y_k = y$, $x_i \neq y_i$ for every $i \in [k]$, and y_i is a neighbor of x_{i+1} for every $i \in [k-1]$. Such vertices exist since there are at least three edges between V_i and V_{i+1} for every $i \in [k-1]$ by the condition (3). Then we concatenate Hamilton paths in $G[V_i]$ between x_i and y_i for each $i = 1, \ldots, k$ that exist by the condition (1) into a Hamilton xy-path in G. Note that in this case we did not use the condition (2).

In the second case x and y are in the same set V_i. We can assume that $x, y \in V_1$. Let P be a Hamilton path in $G[V_1]$ between x and y. If $k = 2$, let ab be an edge of P such that the neighbors a' and b' of a and b in V_2, respectively, are distinct. Such neighbors exist by the condition (2), and such an edge ab exists, because otherwise all vertices in V_1 are only adjacent to one vertex in V_2, a contradiction to the condition (3). By replacing the edge ab with the edge aa', a Hamilton path of $G[V_2]$ between a' and b', and the edge $b'b$ we obtain a Hamilton xy-path in G.

If $k > 2$, let ab be an edge of P such that a and b have neighbors a' and b', respectively, in different sets V_i for $i > 1$. Such an edge ab exists since every vertex of V_1 has a neighbor in some other set V_i by the condition (2), and they cannot be all from the same set, say V_2, for otherwise, the condition (3) for the sets V_1 and V_3 would not hold. By the same argument as in the first case, there exists a Hamilton path R between a' and b' in the subgraph $G[V_2 \cup \cdots \cup V_k]$. Finally, replacing the edge ab on P with the edge aa', the path R, and the edge $b'b$ yields a Hamilton xy-path in G. ◀

[2] We could weaken the condition (3) for $k \geq 4$ so that we only need two disjoint edges between all pairs of sets except for two disjoint pairs.

4 Proof of Theorem 1

We will prove the theorem by induction on n, and let $G_n := G(n, q, T)$. We say that an edge $\{X, A_{ij}X\}$ of G_n is **labeled** ij and an edge $\{X, M_iX\}$ of G_n is **labeled** i.

The proof of the base cases for $q = 2$ and $n = 3$, and for $q \geq 3$ and $n = 2$ is deferred to Section 5; see Lemmas 5 and 6. Here, we prove the inductive step, so we assume that $q = 2$ and $n \geq 4$, or $q \geq 3$ and $n \geq 3$, and that the statement holds for the graph $G(n-1, q, T-n)$. Our main tool is the joining lemma from the previous section (Lemma 4).

Proof of the inductive step. We view rows of an invertible $n \times n$ matrix A as an ordered basis $(r_1, \ldots, r_n)$ of the vector space $\mathbb{F}_q^n$. The first $n-1$ rows span a subspace of dimension $n-1$ which is orthogonal to some subspace of dimension 1. That is,

$$\langle r_1, \ldots, r_{n-1} \rangle = \langle u \rangle^\perp$$

for a nonzero $u \in \mathbb{F}_q^n$ that satisfies $Ru^T = \mathbf{0}$, where R is the $(n-1) \times n$ matrix whose rows are $r_1, \ldots, r_{n-1}$.

We denote by S_u the set of all the $(n-1) \times n$ matrices whose rows form a basis of $\langle u \rangle^\perp$. Observe that this operation gives a bijection between S_u and $GL(n-1, q)$: Remove the i-th column of every matrix in S_u, where i is the index of the first nonzero element of u (which exists, as u is not the zero vector). Furthermore, for every matrix X in S_u, we can add any vector as the last row to form an $n \times n$ invertible matrix, as long as this added vector is independent of the rows of X (i.e., any vector in $\mathbb{F}_q^n \setminus \langle u \rangle^\perp$).

Note that this tallies with the count in (1). Recall that a_{n-1} denotes the number of elements in $GL(n-1, q)$. There are $(q^n - 1)/(q-1)$ choices for the one-dimensional subspace $\langle u \rangle$. For each subspace (with a representative basis u), S_u has a_{n-1} elements, due to the aforementioned bijection. Lastly, for each matrix X in S_u, there are $q^{n-1}(q-1)$ possible last rows, which can be obtained by adding a linear combination of the rows of X to an initial last row and then multiplying the sum by a power of α. Together, we recover the recurrence statement (1).

Following the above analysis, we prove the inductive step in four smaller steps.

First, given a one dimensional subspace with a basis u and a vector v independent of the rows of any matrix in S_u, we denote by the tuple (S_u, v) the set of all matrices in $GL(n, q)$ formed by adding v as the last row to each of the matrices in S_u. Since the aforementioned bijection is an isomorphism of $G_n[(S_u, v)]$ and $G(n-1, q, T-n)$, by inductive hypothesis we conclude that $G_n[(S_u, v)]$ is Hamilton connected.

Before we proceed, we note that row multiplications, divisions, and row additions that do not involve the last row only transform a matrix into another matrix in the same set (S_u, v) for some u, v. Next, adding a row to the last row transforms a matrix in (S_u, v) into another matrix in (S_u, v') for $v \neq v'$. Lastly, adding the last row to another row transforms a matrix in (S_u, v) into another matrix in $(S_{u'}, v)$, where $\langle u \rangle \neq \langle u' \rangle$.

Second, we denote by $(S_u, \langle v \rangle)$ the set of all matrices in $GL(n, q)$ formed by adding any multiple of v as the last row to each of the matrices in S_u. Since multiplication by α generates all nonzero elements of $\mathbb{F}_q$, the edges of label n that multiply the last row form a cycle of length $q-1$. Hence, the graph $G_n[(S_u, \langle v \rangle)] \simeq G_n[(S_u, v)] \,\Box\, C_{q-1}$, which can be easily showed to be Hamilton connected (see [8]).

Third, given a one dimensional subspace with a basis u, we denote by $(S_u, *)$ the set of all matrices in $GL(n, q)$ whose first $n-1$ rows form a matrix in S_u. Here, we use Lemma 4 to join the subgraphs $G_n[(S_u, \langle v \rangle)]$ for all applicable v to prove the Hamilton connectivity of $G_n[(S_u, *)]$. The joining edges between these subgraphs have label in for $i \in [n-1]$ (such an i is guaranteed by the bypass property of T). In order to use the lemma, we show that all of its conditions hold. The condition (1) follows the second step above. The condition (2) is satisfied, because for every $u, v \in \mathbb{F}_q^n$ and a matrix X in $(S_u, v) \subseteq (S_u, \langle v \rangle)$, $A_{in}X$ is a neighbor of X in G_n and in $(S_u, \langle v + r_i(X) \rangle)$, a different set than $(S_u, \langle v \rangle)$. For the condition (3), given u and two distinct $v, v' \notin \langle u \rangle^\perp$ such that $\langle v \rangle \neq \langle v' \rangle$, we have that v is a linear combination of a basis of $\langle u \rangle^\perp$ and v', and consequently $v = x + av'$ for some nonzero $x \in \langle u \rangle^\perp$ and a nonzero $a \in \mathbb{F}_q$. As S_u contains all matrices whose rows form a basis in $\langle u \rangle^\perp$, there exist three matrices X_1, X_2, and X_3 in S_u such that their i-th row is $x = v - av'$. We can guarantee three matrices in S_u, because in the inductive step, $n \geq 3$ and $q \geq 3$, or $n \geq 4$ and $q = 2$, and hence, when we fix the $n - 2$ rows including the i-th row, there are at least three different choices for the remaining row. Then the edges $\{X_1, A_{in}X_1\}$, $\{X_2, A_{in}X_2\}$, and $\{X_3, A_{in}X_3\}$ are the three distinct edges as required by the condition (3). We can now apply Lemma 4 and conclude that $G_n[(S_u, *)]$ is Hamilton connected.

Last, we again apply Lemma 4 to join the different subgraphs $G_n[(S_u, *)]$ for all subspaces $\langle u \rangle$ to complete the inductive step. Here, the joining edges have the label nj for some $j \in [n-1]$, which exist because T is a bypass transition graph. The condition (1) of the lemma follows the previous step. The condition (2) is satisfied, because for any X in some $(S_u, *)$, $A_{nj}X$ is a neighbor of X in G_n and belongs to a different set $(S_{u'}, *)$. For the condition (3), given u, u' not in the same one-dimensional subspace, $\langle u, u' \rangle^\perp$ is a subspace of dimension $n - 2$.

If $n \geq 4$, or $n = 3$ and $q > 3$, there exist three distinct matrices B, B', and B'' whose rows form bases of this $(n-2)$-dimensional subspace. The remaining case $n = 3$ and $q = 3$ is considered separately below. Let $v_u \in \langle u \rangle^\perp \setminus \langle u' \rangle^\perp$ and $v_{u'} \in \langle u' \rangle^\perp \setminus \langle u \rangle^\perp$. Clearly, we have that $v_{u'} - v_u$ is independent of the rows of each matrix B, B', and B''. Let $\tilde{B}$, $\tilde{B}'$, and $\tilde{B}''$ be the $n \times n$ matrix obtained from B, B', and B'' respectively by inserting a new row v_u at the j-th position and $v_{u'} - v_u$ as the last row.

If $n = 3$ and $q = 3$ we have $\langle u \rangle^\perp \cap \langle u' \rangle^\perp = \langle w \rangle = \{0, w, 2w\}$ for some nonzero $w \in \mathbb{F}_3^3$, so there are only two distinct matrices whose rows form bases of this 1-dimensional subspace, in particular $B = (w)$ and $B' = (2w)$. In this case, we define $\tilde{B}$ and $\tilde{B}'$ as in the previous case, but for $\tilde{B}''$ we take the matrix obtained from B by inserting a new row $2v_u$ at the j-th position and $2v_{u'} - 2v_u$ as the last row.

In both cases, $\{\tilde{B}, A_{nj}\tilde{B}\}$, $\{\tilde{B}', A_{nj}\tilde{B}'\}$, and $\{\tilde{B}'', A_{nj}\tilde{B}''\}$ are the three edges as required by the condition (3). This completes all the conditions of Lemma 4 and completes the inductive step. ◀

5 Base cases for the induction

We verify the case when $n = 3$ and $q = 2$ by computer search in SageMath; see [8].

▶ **Lemma 5.** *For every bypass transition graph T, the graph $G(3, 2, T)$ is Hamilton connected.*

Since the only bypass transition graph T for $n = 2$ is the complete graph $T = ([2], \{(1,2), (2,1)\})$, the remaining base cases for $q \geq 3$ and $n = 2$ are captured in the following lemma.

▶ **Lemma 6.** *For a prime power $q \geq 3$, $G(2, q, ([2], \{(1,2), (2,1)\}))$ is Hamilton connected.*

To prove Lemma 6, we first consider the graph that arises by removing the edges that add the second row to the first row; i.e., for a prime power $q \in \mathbb{N}$, we define $G'(q) := G(2, q, ([2], \{(1,2)\}))$. The graph $G'(q)$ is disconnected, and thus, we consider the connected component in $G'(q)$ that contains the identity matrix and denote it by $H(q)$. We will usually write H and G' for $H(q)$ and $G'(q)$ whenever there is no risk of confusion.

It is easy to see that the vertices of H are of the form:

$$V(H) = \left\{ \begin{pmatrix} \alpha^i & 0 \\ \alpha^j a & \alpha^j \end{pmatrix} : i, j \in \{0, \dots, q-2\}, a \in \mathbb{F}_q \right\}.$$

Furthermore, the graph H has a simple structure when analyzing the components by fixing $a \in \mathbb{F}_q$, as described in the following. Let us define

$$V_a = \left\{ \begin{pmatrix} \alpha^i & 0 \\ \alpha^j a & \alpha^j \end{pmatrix} : i, j \in \{0, \dots, q-2\} \right\}$$

and let H_a be the graph induced by fixing a in H; i.e., $H_a := H[V_a]$. We have the following simple observations regarding H and its decomposition by fixing $a \in \mathbb{F}_q$.

(p1) *(H splits into copies of H_a.)* Removing the edges that add the first row to the second row in H disconnects H and splits it into the connected components $\{H_a : a \in \mathbb{F}_q\}$.

(p2) *(The graphs H_a have good Hamiltonicity properties.)* For every $a \in \mathbb{F}_q$, we get that H_a is a toroidal grid of dimensions $(q-1) \times (q-1)$ where each dimension of the grid is given by multiplication by α in the respective row; i.e., $H_a \cong C_{q-1} \square C_{q-1}$. In particular, H_a is isomorphic to H_b for every $a, b \in \mathbb{F}_q$.

(p3) *(The components H_a are well-connected.)* For every $i \in \{0, \dots, q-2\}$ and $a, b \in \mathbb{F}_q$ such that $a \neq b$, we have that

a. $\alpha^i \begin{pmatrix} a-b & 0 \\ a & 1 \end{pmatrix} \in V_a$ and $\alpha^i \begin{pmatrix} a-b & 0 \\ b & 1 \end{pmatrix} \in V_b$ are connected by an edge,

b. $\alpha^i \begin{pmatrix} b-a & 0 \\ a & 1 \end{pmatrix} \in V_a$ and $\alpha^i \begin{pmatrix} b-a & 0 \\ b & 1 \end{pmatrix} \in V_b$ are connected by an edge,

and no other edges between V_a and V_b exist. In particular, for every $a, b \in \mathbb{F}_q$ such that $a \neq b$ we have that $|E[V_a, V_b]| = 2(q-1)$ with all the edges being disjoint.

We exploit these properties as follows: We split H into the components H_a for $a \in \mathbb{F}_q$, then since the graphs H_a are either Hamilton laceable or connected and the components H_a are well-connected we can glue the corresponding Hamilton paths in each H_a to form a Hamilton path in H.

If q is even, for every $a \in \mathbb{F}_q$ the graphs H_a are Hamilton connected. This makes it easier to lift the Hamilton paths from H_a to a Hamilton path in H. However, when q is odd, the picture is much different. In particular, there are parity constraints given by the fact that for every $a \in \mathbb{F}_q$ the graph H_a is now bipartite. To make this formal, we partition $V(H)$ into two colors. We say that $x = \begin{pmatrix} \alpha^i & 0 \\ \alpha^j a & \alpha^j \end{pmatrix}$ is **blue** whenever $i+j$ is even, and it is **red** whenever $i+j$ is odd. We denote the **color of a vertex** $x \in V(H)$ by $\mathrm{col}(x)$. It is easy to show that if $x \in V_a$, $y \in V_b$ for some $a \neq b$ and there is an edge $xy \in E[V_a, V_b]$, then $\mathrm{col}(x) = \mathrm{col}(y)$. Thus, for every edge $xy \in E[V_a, V_b]$ we can define its **(edge) color** as $\mathrm{col}(xy) := \mathrm{col}(x) = \mathrm{col}(y)$.

If $q \equiv 3 \pmod 4$, a simple computation shows that whenever there is a red or blue edge between V_a and V_b for $a, b \in \mathbb{F}_q$ we also have an edge of the opposite color. Thus, the coloring does not impose any extra restrictions.

The problematic case occurs whenever $q \equiv 1 \pmod 4$. In this case, all the edges between V_a and V_b have the same color for $a, b \in \mathbb{F}_q$. Hence, it is natural to consider the graph where we contract every V_a for $a \in \mathbb{F}_q$. Thus, we obtain a new graph $\bar{K}_q$ with $\mathbb{F}_q$ as vertices, and for the edges $xy \in E[V_a, V_b]$ we put a new edge ab colored with $\mathrm{col}(xy)$. This graph is a complete graph on $\mathbb{F}_q$ where the coloring of the edges can be succinctly described as follows:

(*) For x and y in $\mathbb{F}_q$, the edge xy has color red (blue) if there exists an odd (even) $z \in \mathbb{Z}$, such that $x - y = \alpha^z$. (See Figure 2 for an example.)

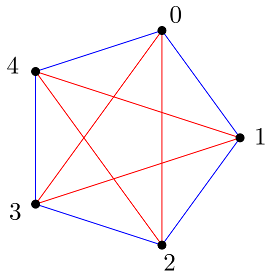

Figure 2 The graph $\bar{K}_5$ for $\alpha = 2$.

If we plan to have a Hamilton path of H that traverses each set V_a at a time for $a \in \mathbb{F}_q$, then for any $a, b \in \mathbb{F}_q$, there is at most one edge of the Hamilton path that crosses between V_a and V_b. Further, as each V_a has even size, this means that as we traverse this Hamilton path, any two consecutive such "crossing" edges have to differ in color. This translates to the requirement that we should have an alternating Hamilton path in $\bar{K}_q$. We show in the next lemma that this holds, even for Hamilton connectivity.

▶ **Lemma 7.** *Let q be a prime power, $q \equiv 1 \pmod 4$. For any two distinct vertices $a, b \in \mathbb{F}_q$ and a color c of either red or blue, there exists an alternating Hamilton ab-path of $\bar{K}_q$ such that a is incident to an edge of color c on the path.*

We defer the proof of Lemma 7 to Section 6. We can use Lemma 7 to prove Hamiltonicity properties of subgraphs of H. To this end, we have the following definition.

▶ **Definition 1.** *An induced subgraph H' of H is **structured** if and only if the following holds:*

1. *For every $a \in \mathbb{F}_q$ we have that $H'[V_a]$ is isomorphic to either $C_{q-1} \,\square\, C_{q-1}$ or $C_{q-1} \,\square\, P_\ell$ for $\ell \geq (q-1)/2$, and*
2. *For every distinct $a, b \in \mathbb{F}_q$, there is at least one edge between $H'[V_a]$ and $H'[V_b]$.*

▶ **Lemma 8.** *Let $q \geq 5$ be an odd integer and H' be a structured induced subgraph of $H(q)$. Let $x, y \in V(H(q))$ be two vertices of different colors such that $x \in V_a$, $y \in V_b$ with distinct $a, b \in \mathbb{F}_q$. Then there exists a Hamilton xy-path in H'.*

The reader may notice similarities between Lemma 8 and the joining lemma (Lemma 4). In particular, they may wonder why we require only *one* edge between components, instead of the three needed in the joining lemma. Recall that the need for three edges in the joining lemma was in the case where we want to have an xy-subpath that spans two consecutive components, but the edges that cross between these two components are incident to either x or y. However, this cannot happen for structured graphs, because the coloring conditions force these endpoints not to be used in crossing edges. The proof of this lemma is presented in [8].

Equipped with Lemmas 4 and 8, we can show the following lemma, whose full proof is also presented in [8].

▶ **Lemma 9.** *For every $q \geq 3$ the graph $H(q)$ is Hamilton connected.*

We are now ready to prove Lemma 6.

Proof of Lemma 6. Let us denote $S = \{(a, 1) : a \in \mathbb{F}_q\} \cup \{(1, 0)\} \subseteq \mathbb{F}_q^2$. For any nonzero vector $u \in S$, let us define $V^u \subseteq V(G)$ to be the set of all matrices in G with the first row in $\langle u \rangle$. Note that each V^u corresponds to a unique component of G', with $V^{(1,0)} = V(H)$.

We apply Lemma 4 with partitioning $\{V_u : u \in S\}$.

We begin by simplifying our arguments using symmetry. Note that any two components of G' are isomorphic via $f_A : x \mapsto xA$ for some $A \in GL(2, q)$. Moreover, for any A the mapping f_A is an automorphism of G that preserves operations on the edges; i.e., if an edge corresponds to the operation M_1, then it will be mapped to some edge that also corresponds to M_1. In particular, H is isomorphic to every other component.

For the condition (1) of Lemma 4, we know that H is Hamilton connected by Lemma 9, so by isomorphism the same holds for every $G[V^u]$. The condition (2) is satisfied simply by adding the second row to the first row. For the condition (3), we first observe that if there is an edge between $G[V^u]$ and $G[V^{u'}]$, there are actually at least $q - 1$ disjoint edges as we can multiply both matrices by α^i. By isomorphism, it is enough to show that there is an edge between H and any V^u distinct from $V^{(1,0)}$. Since $u = (a, 1)$ for some $a \in \mathbb{F}_q$, we can use the edges between $\begin{pmatrix} 1 & 0 \\ a-1 & 1 \end{pmatrix} \in V(H)$ and $\begin{pmatrix} a & 1 \\ a-1 & 1 \end{pmatrix} \in V^{(a,1)}$. ◀

6 Alternating path in two-edge-colored complete graph

In this section, we prove Lemma 7, which is needed in the proof of Lemma 8.

We start with a brief recap of the context needed for this lemma. Suppose q is a prime power and $q \equiv 1 \pmod 4$. We remind the reader that α is a generator of the multiplicative group of $\mathbb{F}_q$. Recall that $\bar{K}_q$ is the complete graph on the vertex set $\mathbb{F}_q$ with edges colored by the following scheme:

(*) For x and y in $\mathbb{F}_q$, the edge xy has color red (blue) if there exists an odd (even) $z \in \mathbb{Z}$ such that $x - y = \alpha^z$.

Our goal is to find an alternating Hamilton path between two prescribed vertices a and b with a prescribed color of the edge incident to a.

We begin by arguing that $\bar{K}_q$ is well-defined. Let 0 be the additive identity and 1 be the multiplicative identity of $\mathbb{F}_q$. By the definition of α, the nonzero elements of $\mathbb{F}_q$ are exactly $\alpha^0, \dots, \alpha^{q-2}$. Furthermore, $\alpha^i = \alpha^{q-1+i}$ for all integers $i \in \mathbb{Z}$. Since q is odd, we conclude that if $\alpha^z = \alpha^{z'}$ for some z, z', then z and z' have the same parity. Thus, for x and y in $\mathbb{F}_q$, there exists a unique $p \in \{0, 1\}$ such that if $x - y = \alpha^z$ then $z \equiv p \pmod 2$. Further, since $\alpha^{(q-1)/2} = -1$, if $x - y = \alpha^z$, then $y - x = \alpha^{z'}$ for $z' = z + (q - 1)/2$. As $q \equiv 1 \pmod 4$, z and z' have the same parity. Therefore, the color of each edge of $\bar{K}_q$ is well-defined.

The problem of finding an alternating cycle/path in a graph has a long history and a wide range of applications; see the survey by Bang-Jensen and Gutin [1]. However, the existing results on alternating Hamilton cycles/paths in two-edge-colored complete graphs (e.g., [2, 3]) cannot be readily applied in our setting. Furthermore, we also specify the color of the first edge of the path, which is not guaranteed by these results. Therefore, we provide a direct and constructive proof of an alternating Hamilton path in our special complete graph.

In the following proof, we use the observation that the operation of adding a constant to all vertex labels preserves edge colors since the difference between any two vertices remains the same.

▶ **Lemma 7.** *Let q be a prime power, $q \equiv 1 \pmod 4$. For any two distinct vertices $a, b \in \mathbb{F}_q$ and a color c of either red or blue, there exists an alternating Hamilton ab-path of $\bar{K}_q$ such that a is incident to an edge of color c on the path.*

Proof. For $i \in \{0, \dots, q-1\}$, define $v_i := \sum_{j=0}^{i} \alpha^j$. Note that $v_{q-2} = 0$, and $v_0 = v_{q-1} = 1$. It is easy to see that $(v_0, \dots, v_{q-2})$ forms an alternating cycle C in $\bar{K}_q$. The only missing vertex in C is $u := -(\alpha-1)^{-1}$. Indeed, if this vertex is on the cycle, then for some t, we must have $(\alpha^{t+1}-1)(\alpha-1)^{-1} = -(\alpha-1)^{-1}$, which implies $\alpha^{t+1} = 0$, a contradiction with the fact that α generates nonzero elements of $\mathbb{F}_q$.

For any $i \in \{0, \dots, q-2\}$ we have $v_i - u = (\alpha^{i+1}-1)(\alpha-1)^{-1} + (\alpha-1)^{-1} = \alpha^{i+1}(\alpha-1)^{-1}$. By a similar argument, we have that $v_{i+1} - u = \alpha^{i+2}(\alpha-1)^{-1} = \alpha(v_i - u)$. Thus, by the coloring scheme (*), uv_i and uv_{i+1} have different colors.

▷ Claim. For any vertex v in C, there exists an alternating Hamilton uv-path such that on the path, u is incident to an edge whose color is different from that of uv.

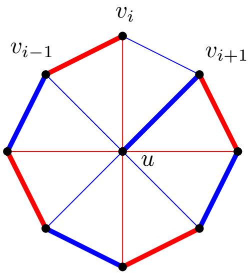

Figure 3 Illustration of the claim's proof. The outer cycle is the cycle C, and the bold edges indicate an alternating Hamilton path.

Proof. Suppose $v = v_i$ for some $i \in \{0, \dots, q-2\}$. By the argument above, uv_{i-1} and uv_{i+1} have the same color. Further, since C is an alternating cycle, $v_{i-1}v_i$ and v_iv_{i+1} have different colors. Hence, one of these two edges have the same color as uv_{i-1} and uv_{i+1}. Without loss of generality, suppose this edge is v_iv_{i+1}. Then we have $(u, v_{i+1}, v_{i+2} \dots, v_{q-2}, v_0, \dots, v_{i-1}, v_i)$ is the desired alternating Hamilton path. See Figure 3 for an illustration. ◁

Consider adding $b - u$ to all vertex labels. The missing vertex from the cycle C above is now b. By the claim above, we obtain an alternating Hamilton ba-path such that the incident edge to b has different color than that of ba. Since q is odd, this implies that the edge incident to a on this path has the same color as ba. Next, we add $a - u$ to all original vertex labels. The missing vertex from C is now a. Again by the claim above, we obtain another alternating Hamilton ab-path such that the incident edge to a has different color than that of ab.

Since the two alternating Hamilton paths above have different colors for the edge incident to a, the lemma follows. ◀

7 Algorithmization

The proof of Theorem 1 can be easily turned into an algorithm that computes a Hamilton path in $G(n, q, T)$ running in time polynomial in $|GL(n, q)|$. This can be obtained by a straightforward recursion based on the joining lemma. More specifcally, the main idea of the proof is to split the graph and proceed recursively. Close examination of the proof of the joining lemma and its applications along our proof shows that such a recursion can be computed in time polynomial in $|GL(n, q)|$; we omit the details.

However, the typical goal from a generation perspective is to have an algorithm that outputs objects one by one with a small delay and preprocessing time. Here, the **delay** is the worst-case time the algorithm takes between consecutively generated objects and the **preprocessing time** is the time before generating any objects. Thus, the natural objective from generation perspective for invertible matrices is an enumeration algorithm running in delay $\mathrm{poly}(n,q) := n^{\mathcal{O}(1)}q^{\mathcal{O}(1)}$ with $\mathrm{poly}(n,q)$ preprocessing. Note that such an algorithm immediately gives a solution to computing Hamilton paths in $G(n,q,T)$ in time $\mathrm{poly}(n,q)|GL(n,q)|$.

The naive implementation of our inductive proof uses a call stack that needs space exponential in n and takes exponential time in n to put the recursive calls in the stack. Despite that, it is still possible to obtain a polynomial delay algorithm by following the recursive structure of the main proof. As highlighted in the stack approach, we cannot store *all* the information given by the recursion. Instead, we only store information related to the *current path* in the recursion tree. More specifically, if we are at a vertex z, we trace back the ℓ recursive calls, each utilizing the joining lemma. The i-th call indicates a pair of a source x_i and a target y_i for which we traverse a Hamilton path. For every $i \in [\ell]$ we store x_i, y_i and a small amount of extra bits serving as a compressed history. It turns out that this is enough information to reconstruct the path of z in the recursion tree and decide how to proceed; more details are given in [8].

8 Open questions

We conclude with several remarks and open questions.

1. **Non-bypass transition graphs.** Does Theorem 1 hold for any strongly connected (and not necessarily bypass) transition graph, in particular for the directed n-cycle? We verified by computer that the result holds for the directed cycle if $q = 2$ and $n = 3$.
2. **Other generators.** Does Theorem 1 hold for other generators of the group $GL(n,q)$? For example, there is the generator $\{M_2A_{1,n}, P_{2\ldots n1}\}$ of size 2, where $P_{2\ldots n1}$ refers to the permutation matrix corresponding to the permutation $2\ldots n1$ [19]. This problem is similar to the sigma-tau problem for permutations solved by Sawada and Williams [17].
3. **Subgroups of $GL(n,q)$.** As an intermediate step, we show that Cayley graphs of certain subgroups of $GL(n,q)$ are Hamilton connected. Can we prove it for other subgroups that correspond to given restrictions of matrices?
4. **Symmetric Hamilton cycles.** Instead of Hamilton connectivity we may ask for Hamilton cycles that are preserved under a large cyclic subgroup of automorphisms. This problem was recently studied by Gregor, Merino, and Mütze [9] for several highly symmetric graphs. The graphs considered here are also highly symmetric.
5. **Matrices over rings.** Another natural extension is to explore if our results extend to invertible matrices in the ring setting. This is particularly interesting for cyclic rings; i.e., checking Hamiltonicity of Cayley graphs of invertible matrices in $\mathbb{Z}_k$ for $k \in \mathbb{N}$. Naturally, the methods will highly depend on the chosen generators for which there does not seem to be an obvious choice.
6. **Alternating Hamilton paths in 2-colored K_{2n+1}.** Despite our efforts and many existing results on properly colored Hamilton cycles in complete graphs (see a survey [1]), we did not find an answer to the following question. Is it true that the complete graph K_{2n+1} with 2-colored edges so that every vertex is incident with exactly n edges of each color contains an alternating Hamilton path between any two vertices?
7. **Efficient algorithms.** Is there a generating algorithm that achieves $\mathcal{O}(n)$ delay? Is there a simple greedy algorithm?

References

1 J. Bang-Jensen and G. Gutin. Alternating cycles and paths in edge-coloured multigraphs: a survey. *Discrete Math.*, 165/166:39–60, 1997. doi:10.1016/S0012-365X(96)00160-4.

2 J. Bang-Jensen, G. Gutin, and A. Yeo. Properly coloured Hamiltonian paths in edge-coloured complete graphs. *Discrete Appl. Math.*, 82(1-3):247–250, 1998. doi:10.1016/S0166-218X(97)00062-0.

3 M. Bánkfalvi and Zs. Bánkfalvi. Alternating Hamiltonian circuit in two-coloured complete graphs. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 11–18. Academic Press, New York-London, 1968.

4 C. C. Chen and N. F. Quimpo. On strongly Hamiltonian abelian group graphs. In *Combinatorial mathematics, VIII (Geelong, 1980)*, volume 884 of *Lecture Notes in Math.*, pages 23–34. Springer, Berlin-New York, 1981. doi:10.1007/BFb0091805.

5 C. Cooper. On the rank of random matrices. *Random Structures Algorithms*, 16(2):209–232, 2000. doi:10.1002/(SICI)1098-2418(200003)16:2<209::AID-RSA6>3.0.CO;2-1.

6 M. Dupuis and S. Wagon. Laceable knights. *Ars Math. Contemp.*, 9:115–124, 2015. doi:10.26493/1855-3974.420.3C5.

7 T. Fabšič, O. Grošek, K. Nemoga, and P. Zajac. On generating invertible circulant binary matrices with a prescribed number of ones. *Cryptogr. Commun.*, 10(1):159–175, 2018. doi:10.1007/s12095-017-0239-4.

8 P. Gregor, H. P. Hoang, A. Merino, and O. Mička. Generating all invertible matrices by row operations. *arXiv preprint*, 2024. Full preprint version of the present article. arXiv:2405.01863.

9 P. Gregor, A. Merino, and T. Mütze. The Hamilton Compression of Highly Symmetric Graphs. In S. Szeider, R. Ganian, and A. Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 54:1–54:14, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.MFCS.2022.54.

10 E. Hartung, H. P. Hoang, T. Mütze, and A. Williams. Combinatorial generation via permutation languages. I. Fundamentals. *Trans. Amer. Math. Soc.*, 375(4):2255–2291, 2022. doi:10.1090/tran/8199.

11 D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial algorithms. Part 1.* Addison-Wesley, Upper Saddle River, NJ, 2011.

12 L. Lovász. Problem 11. In *Combinatorial Structures and Their Applications (Proc. Calgary Internat. Conf., Calgary, AB, 1969)*. Gordon and Breach, New York, 1970.

13 T. Mütze. Combinatorial Gray codes—an updated survey. *Electron. J. Combin.*, Dynamic Surveys DS26:93 pp., 2023. doi:10.37236/11023.

14 Naomi Nishimura. Reasons to fall (more) in love with combinatorial reconfiguration. In *WALCOM: algorithms and computation*, volume 14549 of *Lecture Notes in Comput. Sci.*, pages 9–14. Springer, Singapore, 2024. doi:10.1007/978-981-97-0566-5_2.

15 David J. Rasmussen and Carla D. Savage. Hamilton-connected derangement graphs on S_n. *Discrete Math.*, 133(1-3):217–223, 1994. doi:10.1016/0012-365X(94)90028-0.

16 Frank Ruskey and Carla Savage. Hamilton cycles that extend transposition matchings in Cayley graphs of S_n. *SIAM J. Discrete Math.*, 6(1):152–166, 1993. doi:10.1137/0406012.

17 J. Sawada and A. Williams. Solving the sigma-tau problem. *ACM Trans. Algorithms*, 16(1):Art. 11, 17 pp., 2020. doi:10.1145/3359589.

18 M. Tchuente. Generation of permutations by graphical exchanges. *Ars Combin.*, 14:115–122, 1982.

19 W. C. Waterhouse. Two generators for the general linear groups over finite fields. *Linear and Multilinear Algebra*, 24(4):227–230, 1989. doi:10.1080/03081088908817916.