

HAMILTONICITY OF SCHRIJVER GRAPHS AND STABLE KNESER GRAPHS

TORSTEN MÜTZE AND NAMRATA

ABSTRACT. For integers $k \geq 1$ and $n \geq 2k + 1$, the Schrijver graph $S(n, k)$ has as vertices all k -element subsets of $[n] := \{1, 2, \dots, n\}$ that contain no two cyclically adjacent elements, and an edge between any two disjoint sets. More generally, for integers $k \geq 1$, $s \geq 2$, and $n \geq sk + 1$, the s -stable Kneser graph $S(n, k, s)$ has as vertices all k -element subsets of $[n]$ in which any two elements are in cyclical distance at least s . We prove that all the graphs $S(n, k, s)$, in particular Schrijver graphs $S(n, k) = S(n, k, 2)$, admit a Hamilton cycle that can be computed in time $\mathcal{O}(n)$ per generated vertex.

1. INTRODUCTION

For integers $k \geq 1$ and $n \geq 2k + 1$, the *Kneser graph* $K(n, k)$ has as vertices all k -element subsets of $[n] := \{1, 2, \dots, n\}$, and an edge between any two sets A and B that are disjoint, i.e., $A \cap B = \emptyset$. Kneser graphs were popularized by Lovász [Lov78] in his celebrated proof of Kneser's conjecture, where he proved that the chromatic number of $K(n, k)$ equals $n - 2(k - 1)$, using topological arguments. Schrijver [Sch78] considered an induced subgraph of $K(n, k)$, obtained by considering as vertices only the k -element sets that contain no two cyclically adjacent numbers, i.e., for any vertex $A \subseteq [n]$ and any two elements $i, j \in A$ with $i < j$, we require that $j - i \geq 2$ and $j - i \neq n - 1$. This graph is referred as *Schrijver graph* $S(n, k)$, and Schrijver showed that it is vertex-critical, i.e., it has the same chromatic number as the Kneser graph $K(n, k)$, but removing any of its vertices will decrease the chromatic number. Both Kneser and Schrijver graphs are generalized by *s-stable Kneser graphs* $S(n, k, s)$. Those are defined for integers $k \geq 1$, $s \geq 1$, and $n \geq \max\{2, s\}k + 1$, as the induced subgraph of $K(n, k)$ obtained by considering as vertices only the k -element sets that contain no two numbers in cyclical distance less than s , i.e., for any vertex $A \subseteq [n]$ and any two elements $i, j \in A$ with $i < j$, we require that $j - i \geq s$ and $j - i \notin \{n - 1, n - 2, \dots, n - s + 1\}$. Note that $S(n, k, 1)$ is precisely the Kneser graph $K(n, k)$, whereas $S(n, k, 2)$ is precisely the Schrijver graph $S(n, k)$. Meunier [Meu11, Conj. 2] conjectured that the chromatic number of $S(n, k, s)$ equals $n - s(k - 1)$ when $s \geq 2$, and some progress in this direction has been made [Jon12, Che15].

In this paper, we consider the problem of finding Hamilton cycles in Schrijver graphs, and more generally, in s -stable Kneser graphs. For Kneser graphs $K(n, k)$, this problem was recently settled affirmatively by Merino, Mütze and Namrata [MMN23]. Specifically, we showed that all Kneser graphs $K(n, k)$, $n \geq 2k + 1$, admit a Hamilton cycle, unless $(n, k) = (5, 2)$, which is the exceptional Petersen graph. This result followed a long line of work (see, e.g. [CF02, Che03, Joh11, MNW21]), and it resolved a problem that was open for 50 years.

(Torsten Mütze) DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF WARWICK, UNITED KINGDOM & DEPARTMENT OF THEORETICAL COMPUTER SCIENCE AND MATHEMATICAL LOGIC, CHARLES UNIVERSITY, PRAGUE, CZECH REPUBLIC

(Namrata) DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF WARWICK, UNITED KINGDOM
E-mail addresses: `torsten.mutze@warwick.ac.uk`, `namrata@warwick.ac.uk`.

This work was supported by Czech Science Foundation grant GA 22-15272S. Both authors participated in the workshop 'Combinatorics, Algorithms and Geometry' in March 2024, which was funded by German Science Foundation grant 522790373.

In this work, we apply the methods developed in [MMN23] to Schrijver graphs and s -stable Kneser graphs, which yields the following new results.

Theorem 1. *For any $k \geq 1$, $s \geq 2$ and $n \geq sk + 1$, the s -stable Schrijver graph $S(n, k, s)$ has a Hamilton cycle. In particular, the Schrijver graph $S(n, k) = S(n, k, 2)$ has a Hamilton cycle.*

We also provide an efficient algorithm for computing these cycles.

Theorem 2. *Let $k \geq 1$, $s \geq 2$ and $n \geq sk + 1$. There is an algorithm for computing a Hamilton cycle in $S(n, k, s)$ that takes time $\mathcal{O}(n)$ to compute the next vertex on the cycle.*

The initialization time and memory required by our algorithm is also $\mathcal{O}(n)$. The assumption $s \geq 2$ in Theorem 2 is important. In fact, we do not know any algorithm with running time polynomial in n and k for computing a Hamilton cycle in $S(n, k, 1) = K(n, k)$, even though by the aforementioned results [MMN23], $K(n, k)$ is known to have a Hamilton cycle (unless $(n, k) = (5, 2)$).

We implemented the algorithm given by Theorem 2 in C++, and made it available for experimentation and download on the Combinatorial Object Server website [cos].

After submitting this manuscript, we learnt that Theorem 1 was proved independently by Ledezma and Pastine [LP24].

1.1. Outline of this paper. In Sections 2 and 3 we present the proofs of Theorem 1 and 2, respectively, for the special case of Schrijver graphs $S(n, k) = S(n, k, 2)$. With minimal adaptations, these proofs can be adjusted to also work for the general case of s -stable Kneser graphs for any $s \geq 2$. We explain those adaptations in Section 4.

2. PROOF OF THEOREM 1 FOR SCHRIJVER GRAPHS $S(n, k)$

As mentioned before, we apply the methods developed in [MMN23] to construct a Hamilton cycle in the Schrijver graph $S(n, k)$. The construction proceeds in two steps: In the first step we construct a *cycle factor* in the graph, i.e., a collection of disjoint cycles that together visit all vertices. In the second step we connect the cycles of the factor to a single Hamilton cycle.

2.1. Cycle factor construction and gliders.

2.1.1. Preliminaries. We interpret vertices of the Schrijver graph $S(n, k)$ as bitstrings by considering the corresponding characteristic vectors. For every k -element subset A of $[n]$, this is a bitstring of length n with exactly k many 1s at the positions corresponding to the elements of A . For example, the vertex $A = \{2, 7, 9\}$ of $S(9, 3)$ is represented by the bitstring $x = 010000101$. Throughout this paper, we consider the indices into these bitstrings modulo n , i.e., $x_{n+i} = x_i$. By definition of Schrijver graphs, no two 1s in x are next to each other (considered cyclically), or in other words, every 1-bit is surrounded by 0-bits. We write $X_{n,k}$ for the set of all bitstrings that encode vertices of $S(n, k)$. A straightforward counting argument shows that $|X_{n,k}| = \binom{n-k+1}{k} - \binom{n-k-1}{k-2} = \frac{n}{n-k} \binom{n-k}{k}$ (see [Kap43]). For any bitstring x and any integer $r \geq 0$, we write x^r for the r -fold concatenation of x with itself.

For two strings x and y over any alphabet, we write $x \prec y$ if x is lexicographically strictly smaller than y , and for a set of strings X we write $\text{lexmin } X$ for the lexicographically least string in X , i.e., for the unique string $x \in X$ that satisfies $x \prec y$ for all $y \in X \setminus \{x\}$.

2.1.2. Cycle factor construction. For a bitstring $x \in X_{n,k}$, we write $\sigma(x)$ for the string obtained from x by a cyclic right shift by 1 position. We define the *necklace of x* by

$$\langle x \rangle := \{\sigma^i(x) \mid i \geq 0\},$$

i.e., the necklace $\langle x \rangle$ is the set of all cyclic shifts of x .

As $x \in X_{n,k}$ has no two 1s next to each other, $(x, \sigma(x))$ is an edge in the Schrijver graph. Consequently, for a bitstring x and $r := |\langle x \rangle|$, the sequence

$$C(x) := (x, \sigma(x), \sigma^2(x), \dots, \sigma^{r-1}(x)) \quad (1a)$$

defines a cycle in $S(n, k)$; see the top part of Figure 1. We thus obtain a cycle factor in $S(n, k)$ by

$$\mathcal{C}_{n,k} := \{C(x) \mid x \in X_{n,k}\}. \quad (1b)$$

2.1.3. Matched and unmatched bits. The operation $x \mapsto \sigma(x)$ can also be described as follows: To create the string $\sigma(x)$ from x , every 1-bit at some position i in x is transposed with the 0-bit immediately to the right of it, i.e., at position $i + 1$. Given $x \in X_{n,k}$, we refer to the bits at positions $\mu(x) := \{i \mid x_i = 1 \text{ or } x_{i-1} = 1\} \subseteq [n]$ as *matched*, and to the bits at positions $\bar{\mu}(x) := [n] \setminus \mu(x) = \{i \mid x_i = 0 \text{ and } x_{i-1} = 0\}$ as *unmatched*. The motivation for this terminology will become clear momentarily. We denote unmatched bits in x as $-$. For example, we write $x = 010001001010101 = 010--10-1010101$ and we have $\mu(x) = \{1, 2, 3, 6, 7, 9, 10, 11, 12, 13, 14, 15\}$ and $\bar{\mu}(x) = \{4, 5, 8\}$. Note that $|\mu(x)| = 2k$ and therefore $|\bar{\mu}(x)| = n - 2k$, so x has exactly $n - 2k$ many $-$ s. On the right hand side of Figure 1, matched bits are colored, and unmatched bits are white.

We refer to any maximal (cyclic) substring of matched bits in $x \in X_{n,k}$ as a *block* and to any maximal (cyclic) substring of unmatched bits as a *gap*. The *length* of a block or gap is the number of bits belonging to the block or gap, respectively. Clearly, the length of a block is always even. For example, the bitstring x from before has two blocks of lengths 2 and 10, and two gaps of lengths 2 and 1.

2.1.4. Gliders. We consider every substring 10 in x as an entity called a *glider*. Specifically, the set of all gliders in x is the set

$$\Gamma(x) := \{(i, i + 1) \mid x_i = 1 \text{ and } x_{i+1} = 0\}. \quad (2)$$

For example, for $x = 010--10-1010101$ from before we have $\Gamma(x) = \{(2, 3), (6, 7), (9, 10), (11, 12), (13, 14), (15, 1)\}$.

The motivation for the word ‘glider’ becomes clear when comparing the gliders in x with those in $\sigma(x)$, and when interpreting one application of σ (i.e., moving one step along the cycle $C(x)$) as one unit of time moving forward. Indeed, if $(i, i + 1) \in \Gamma(x)$, then we have $(i + 1, i + 2) \in \Gamma(\sigma(x))$, i.e., a glider that occupies positions i and $i + 1$ in x moves to positions $i + 1$ and $i + 2$ after one time step; see Figure 1. We can thus view these pairs of matched bits as entities that change position over time, namely they move one position to the right with each time step, i.e., their ‘speed’ is 1.

2.1.5. Connection to parenthesis matching in Kneser graphs. The aforementioned mapping $x \mapsto \sigma(x)$ can be generalized to yield a more general mapping $x \mapsto f(x)$ that recovers precisely the construction of a cycle factor in $K(n, k)$ used in [MMN23], which agrees with the one before on the Schrijver subgraph $S(n, k) \subseteq K(n, k)$, as follows: To create the string $f(x)$ from x , where x now is an arbitrary bitstring of length n with k many 1s, every 1-bit at some position i in x is transposed with the 0-bit that terminates the shortest possible substring at positions $i + 1, i + 2, \dots$ that has strictly more 0s than 1s; see the bottom part of Figure 1. An alternative and equivalent definition of f is to interpret the 1s in x as opening brackets and the 0s as closing brackets, to match pairs of opening and closing brackets in the natural way (leaving $n - 2k$ closing brackets unmatched), and to transpose every bit with the corresponding matched partner to obtain $f(x)$. This idea, pioneered by Greene and Kleitman [GK76] in another context, is often referred to as *parenthesis matching*, which explains our choice of terminology for matched and unmatched bits before. In fact, the paper [MMN23] describes a more general partition of groups of matched bits into gliders that works in general Kneser graphs, and then different gliders may move at

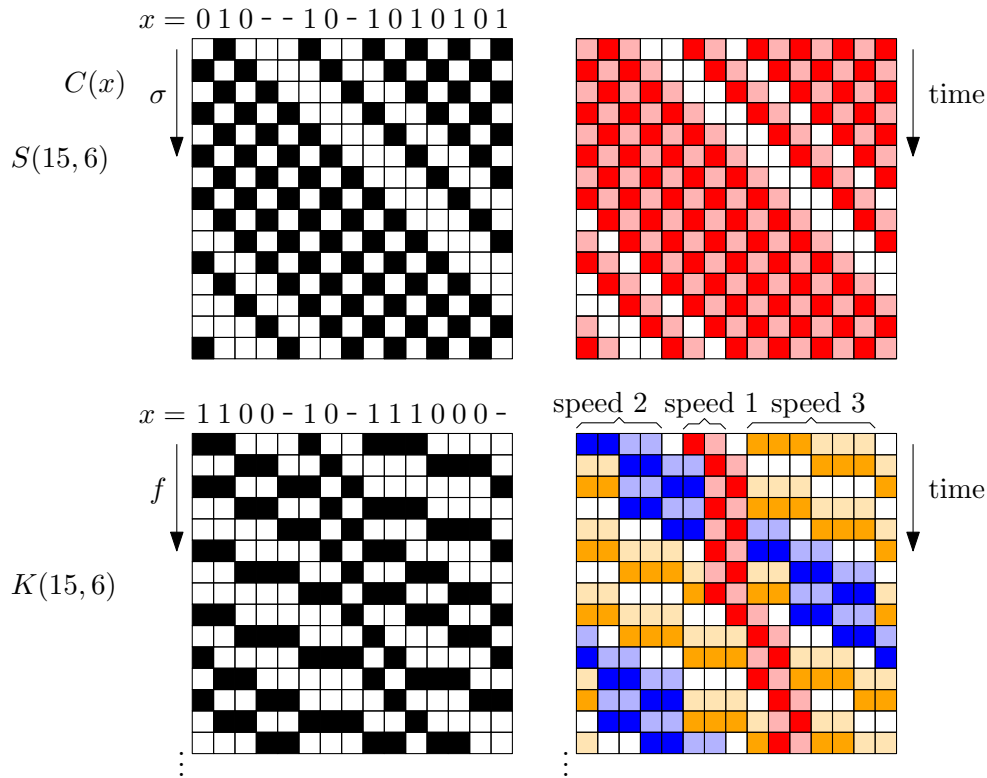


FIGURE 1. A cycle in the Schrijver graph $S(15,6)$ constructed from the mapping σ (top), and a cycle in the Kneser graph $K(15,6)$ constructed from the more general mapping f . The cycle in $S(15,6)$ is shown completely, whereas only the first 15 vertices of the cycle in $K(15,6)$ are displayed. On the left, each vertex is represented by a bitstring, with 1-bits colored black and 0-bits colored white. The vertices are printed from top to bottom in the order of the cycle. The right-hand side shows the interpretation of certain groups of bits as gliders, and their movement over time. Matched bits belonging to the same glider are given the same color, with the opaque filling given to 1-bits, and the transparent filling given to 0-bits.

different speeds; see the bottom part of Figure 1. For our discussion of Schrijver graphs this more general definition of gliders is not needed.

2.2. Gluing the cycles together. We now show how to connect the cycles in the factor $\mathcal{C}_{n,k}$ defined in (1) together to a single Hamilton cycle of $S(n,k)$. Each joining step proceeds by gluing together two cycles from the factor to one cycle via a 4-cycle.

2.2.1. Connectors. Given two bitstrings $x, y \in X_{n,k}$, we refer to the pair (x, y) as a *connector*, if x and $\sigma^{-1}(y)$ agree in all but three positions $i, i+1, i+2$ such that $x_{i,i+1,i+2} = 10-$ and $\sigma^{-1}(y)_{i,i+1,i+2} = -10$, i.e., we have

$$\begin{bmatrix} x \\ \sigma^{-1}(y) \end{bmatrix} = \begin{bmatrix} \cdots 10 \cdots \\ \cdots -10 \cdots \end{bmatrix}, \quad (3)$$

where the vertically aligned \cdots stand for substrings in which x and $\sigma^{-1}(y)$ agree. In other words, x and $\sigma^{-1}(y)$ have the same gliders, except that the rightmost glider in a block of x is ‘pushed’ to the right by one position, i.e., it is transposed with the unmatched bit to its right.

We write $\mathcal{X}_{n,k}$ for the set of all connectors. Note that a bitstring $z \in \mathcal{X}_{n,k}$ may be part of several different connectors. Specifically, if there are p gaps in z , then z is contained in exactly p connectors of the form (z, y) and p connectors of the form (x, z) .

Lemma 3. *For any connector $(x, y) \in \mathcal{X}_{n,k}$, the sequence $C_4(x, y) := (x, \sigma(x), \sigma(y), y)$ is a 4-cycle in the Schrijver graph $S(n, k)$.*

Proof. We already know that $(x, \sigma(x))$ and $(y, \sigma(y))$ are edges in $S(n, k)$. Therefore, it suffices to prove that (x, y) is an edge in $S(n, k)$, as this will imply that $(\sigma(x), \sigma(y))$ is also an edge. Let $i, i+1, i+2$ be such that $x_{i,i+1,i+2} = 10-$ and $\sigma^{-1}(y)_{i,i+1,i+2} = -10$. From (3) we see that y agrees with $\sigma(x)$ in all positions but $i+1, i+2, i+3$, specifically that $\sigma(x)_{i+1,i+2,i+3} = 10-$ and $y_{i+1,i+2,i+3} = -10$. As $\sigma(x)$ and x have no 1s at the same position, and $x_{i+2} = -$ (this is the position where y has a 1), we obtain that y and x have no 1s at the same position, i.e., (x, y) is indeed an edge in $S(n, k)$. \square

Observe that if $(x, y) \in \mathcal{X}_{n,k}$ and $C(x)$ and $C(y)$ are two distinct cycles in the factor $\mathcal{C}_{n,k}$ defined in (1), then the symmetric difference of the edge sets $(C(x) \cup C(y)) \Delta C_4(x, y)$ is a single cycle in $S(n, k)$ on the same vertex set as $C(x) \cup C(y)$, i.e., the 4-cycle ‘glues’ two cycles from the factor together to a single cycle.

We define

$$Y_{n,k} := \{x \in X_{n,k} \mid x_1 = - \text{ and } x_n = 0\} = \{x \in X_{n,k} \mid 1 \in \bar{\mu}(x) \text{ and } n \in \mu(x)\}. \quad (4)$$

These are vertices of $S(n, k)$ in which the first bit is unmatched and the last bit is matched (specifically, a matched 0-bit). Any $x \in Y_{n,k}$ has the form

$$x = -\gamma_1(10)^{\beta_1/2} -\gamma_2(10)^{\beta_2/2} \dots -\gamma_r(10)^{\beta_r/2} \quad (5a)$$

for integers $r \geq 1$, $\gamma_1, \dots, \gamma_r > 0$, and even $\beta_1, \dots, \beta_r > 0$. In words, γ_i is the length of the i th gap in x , and β_i is the length of the i th block in x . We then define

$$L(x) := (-\beta_r, \gamma_r, -\beta_{r-1}, \gamma_{r-1}, \dots, -\beta_2, \gamma_2, -\beta_1, \gamma_1), \quad (5b)$$

i.e., $L(x)$ records the lengths of blocks and gaps in x from *right to left*, and the lengths of blocks are given a negative sign. For example, for the bitstring $x = ----10101010---10-1010--101010 \in Y_{30,10}$ we have $L(x) = (-6, 2, -4, 1, -2, 3, -8, 4)$.

For any $x \in X_{n,k}$ we define $L(\langle x \rangle)$ as the lexicographically smallest string $L(x')$ over all $x' \in \langle x \rangle \cap Y_{n,k}$, i.e.,

$$L(\langle x \rangle) := \text{lexmin} \{L(x') \mid x' \in \langle x \rangle \cap Y_{n,k}\}. \quad (5c)$$

In words, we consider all cyclic shifts x' of x that belong to the set $Y_{n,k}$, and take the lexicographically least string $L(x')$ among them. Clearly, cyclic shifts of $x \in Y_{n,k}$ yield different cyclic shifts of $L(x)$, all starting with a negative number. In the example from before, we obtain $L(\langle x \rangle) = (-8, 4, -6, 2, -4, 1, -2, 3)$. Informally, for a given $x \in X_{n,k}$, the definition $L(\langle x \rangle)$ locates the (cyclically) longest possible block in x , preceded by the shortest possible gap, preceded by the longest possible block, preceded by the shortest possible gap, etc., which motivates the negative signs in (5b) and the right-to-left reading.

We will also need the following variant of the definition (5b) from before. For any $x \in X_{n,k}$ and any position $p \in [n]$ with $x_{p,p+1} = 0-$, i.e., the p th bit of x is matched and the $(p+1)$ st bit is unmatched, we have $\sigma^{-p}(x) \in Y_{n,k}$ and may thus define

$$L(x, p) := L(\sigma^{-p}(x)). \quad (6)$$

In words, this is the sequence of lengths of blocks and gaps in x starting from position p read in the right-to-left direction (with negative signs for block lengths).

2.2.2. Auxiliary graph. Our strategy is to join the cycles of the factor $\mathcal{C}_{n,k}$ by repeatedly gluing pairs of them together via connectors, as described before. For any set of connectors $\mathcal{U} \subseteq \mathcal{X}_{n,k}$ we define a graph $\mathcal{H}_{n,k}[\mathcal{U}]$ as follows: The nodes of $\mathcal{H}_{n,k}[\mathcal{U}]$ are the cycles of the factor $\mathcal{C}_{n,k}$. Furthermore, for any connector $(x, y) \in \mathcal{U}$ for which $C(x)$ and $C(y)$ are distinct cycles, we add an edge that connects $C(x)$ and $C(y)$ to the graph $\mathcal{H}_{n,k}[\mathcal{U}]$. To obtain a Hamilton cycle in $S(n, k)$, we require two things: The graph $\mathcal{H}_{n,k}[\mathcal{U}]$ must be connected. Furthermore, we require that any




FIGURE 2. Illustration of the proof of Lemma 4.

two of the 4-cycles $C_4(x, y)$ and $C_4(x', y')$ with $(x, y), (x', y') \in \mathcal{U}$ used for the joining are edge-disjoint. To ensure the second property, it is enough to guarantee that $C_4(x, y)$ and $C_4(x', y')$ do not have an edge in common on one of the cycles of the factor $\mathcal{C}_{n,k}$. This condition is equivalent to $\{x, y\} \cap \{x', y'\} = \emptyset$, and if it holds then we call the connectors (x, y) and (x', y') *disjoint*. We therefore require a set of pairwise disjoint connectors $\mathcal{U} \subseteq \mathcal{X}_{n,k}$ such that $\mathcal{H}_{n,k}[\mathcal{U}]$ is a connected graph. Note the following subtlety: There may be two connectors $(x, y), (x', y') \in \mathcal{U}$ with $x' = \sigma(x)$ and $y' = \sigma(y)$, which by our definition are disjoint, even though the 4-cycles $C_4(x, y)$ and $C_4(x', y')$ share the edge $(x', y') = (\sigma(x), \sigma(y))$. However, in the joining process, only one of the two 4-cycles will be used, as they connect the same pair of cycles $C(x) = C(x')$ and $C(y) = C(y')$.

A set of connectors with the aforementioned properties can be defined explicitly for any fixed integer $p \in [n]$ by

$$\mathcal{U}_p := \{(x, y) \in \mathcal{X}_{n,k} \mid x_{p,p+1,p+2} = 10- \text{ and } \sigma^{-1}(y)_{p,p+1,p+2} = -10\}; \quad (7)$$

recall (3). In words, these are connectors in which the ‘pushing’ of a glider occurs at position p .

Lemma 4. *For any $k \geq 1$, $n \geq 2k + 1$, and $p \in [n]$, let $\mathcal{U}_p \subseteq \mathcal{X}_{n,k}$ be the set of connectors defined in (7). The connectors in \mathcal{U}_p are pairwise disjoint, and $\mathcal{H}_{n,k}[\mathcal{U}_p]$ is a connected graph.*

Proof. Consider two connectors $(x, y), (x', y') \in \mathcal{U}_p$ with $x \neq x'$ which implies that $y \neq y'$. Clearly, we have $y \neq x$ and $y' \neq x'$. From (7) we also see that $x_{p+2} = -$ and $\sigma^{-1}(y)_{p+1} = y'_{p+2} = 1$, implying that $x \neq y'$, and symmetrically that $x'_{p+2} = -$ and $\sigma^{-1}(y)_{p+1} = y_{p+2} = 1$, implying that $x' \neq y$. This proves that $\{x, y\} \cap \{x', y'\} = \emptyset$, as desired.

It remains to prove that $\mathcal{H}_{n,k}[\mathcal{U}_p]$ is connected. We say that a vertex $x \in X_{n,k}$ is *connectable*, if $x_{p,p+1,p+2} = 10-$. By (7), this means that there is a unique corresponding connector $(x, y) \in \mathcal{U}_p$. For any connectable vertex x with at least two blocks, we let β be the length of the block ending at position $p + 1$ in x , γ be the length of the gap to the left of it, γ' be the length of the gap to the right of it, and we let β' be the length of the block to the right of this gap; see Figure 2 (a).

We define

$$P(x) := L(x, p + 1 + \gamma' + \beta') \quad (8)$$

with $L(x, p)$ as defined in (6). In words, $P(x)$ is obtained by recording the lengths of blocks and gaps in x from right to left, starting at the first block strictly to the right of position $p + 1$ in x , where block lengths receive an additional negative sign. We thus have $P(x) = (-\beta', \gamma', -\beta, \gamma, \dots)$. We also define the connectable vertex

$$\tau(y) := \begin{cases} \sigma^{-2}(y) & \text{if } \gamma' > 1, \\ \sigma(y) & \text{if } \gamma' = 1 \text{ and } \beta > 2, \\ \sigma^{\gamma+1}(y) & \text{if } \gamma' = 1 \text{ and } \beta = 2; \end{cases} \quad (9)$$

see Figure 2. It satisfies

$$P(\tau(y)) = \begin{cases} (-\beta', \gamma' - 1, \dots) & \text{if } \gamma' > 1, \\ (-(\beta' + 2), \dots) & \text{if } \gamma' = 1, \end{cases} \quad (10)$$

implying that $P(\tau(y)) \prec P(x)$, i.e., $P(\tau(y))$ is lexicographically smaller than $P(x)$. This shows that we can move in the graph $\mathcal{H}_{n,k}[\mathcal{U}_p]$ along lexicographically decreasing connectors, to the lexicographically least connectable vertex, which lies in the cycle $C(-n-2k(10)^k)$, i.e., it has only a single block and a single gap (all gliders are adjacent). This proves that $\mathcal{H}_{n,k}[\mathcal{U}_p]$ is indeed connected, which completes the proof. \square

We are now ready to present the proof of Theorem 1 for Schrijver graphs $S(n, k)$.

Proof of Theorem 1 for Schrijver graphs $S(n, k) = S(n, k, 2)$. By Lemma 4, the set $\mathcal{U}_p \subseteq X_{n,k}$ of connectors defined in (7) for any fixed value of $p \in [n]$ has the property that the connectors in \mathcal{U}_p are pairwise disjoint, and $\mathcal{H}_{n,k}[\mathcal{U}_p]$ is a connected graph. Let \mathcal{T} be a minimal subset of \mathcal{U}_p such that $\mathcal{H}_{n,k}[\mathcal{T}]$ is connected, i.e., this graph is a spanning tree. As the connectors in \mathcal{U}_p are pairwise disjoint, no two of the 4-cycles $C_4(x, y)$, $(x, y) \in \mathcal{T}$, defined in Lemma 3 have an edge in common with one of the cycles of the factor $\mathcal{C}_{n,k}$ in $S(n, k)$ defined in (1). Furthermore, by the minimality of \mathcal{T} , no two of these 4-cycles have an edge in common that does not lie on one of the cycles of the factor (otherwise, \mathcal{T} would contain two edges connecting the same pair of nodes), so these 4-cycles are pairwise edge-disjoint. Consequently, as $\mathcal{H}_{n,k}[\mathcal{T}]$ is connected, the symmetric difference of the edge set of the cycle factor $\mathcal{C}_{n,k}$ with the 4-cycles $C_4(x, y)$, $(x, y) \in \mathcal{T}$, is a Hamilton cycle in $S(n, k)$. \square

3. PROOF OF THEOREM 2 FOR SCHRIJVER GRAPHS $S(n, k)$

To turn the proof presented in the previous section into an efficient algorithm, we first define an explicit spanning tree in the graph $\mathcal{H}_{n,k}[\mathcal{U}_p]$. To define such a spanning tree, we select exactly one connectable vertex x on each cycle of the factor $\mathcal{C}_{n,k}$, namely the one with lexicographically least value of $P(x)$, i.e.,

$$\mathcal{T}_p := \{(x, y) \in \mathcal{U}_p \mid x \text{ has at least 2 blocks and } P(x) = L(\langle x \rangle)\}, \quad (11)$$

where $P(x)$ and $L(\langle x \rangle)$ are defined in (8) and (5), respectively; see Figure 3.

Lemma 5. *For any $k \geq 1$, $n \geq 2k + 1$ and $p \in [n]$, let $\mathcal{T}_p \subseteq \mathcal{U}_p \subseteq \mathcal{X}_{n,k}$ be the set of connectors defined in (11). Then $\mathcal{H}_{n,k}[\mathcal{T}_p]$ is a spanning tree of $\mathcal{H}_{n,k}[\mathcal{U}_p]$.*

Proof. We consider the edges $(C(x), C(y))$ of $\mathcal{H}_{n,k}[\mathcal{T}_p]$ for connectors $(x, y) \in \mathcal{T}_p$ as being *oriented* from $C(x)$ to $C(y)$.

We first show that $\mathcal{H}_{n,k}[\mathcal{T}_p]$ is connected. We have argued before that for $(x, y) \in \mathcal{T}_p$ we have $P(\tau(y)) \prec P(x)$ (recall (10)). Let z be the lexicographically least connectable vertex on the cycle $C(\tau(y))$, i.e., we have $P(z) = P(\tau(y))$ or $P(z) \prec P(\tau(y))$. Combining these inequalities shows that $P(z) \prec P(x)$. This shows that we can move in the graph $\mathcal{H}_{n,k}[\mathcal{T}_p]$ along lexicographically decreasing connectors to the lexicographically least connectable vertex, which




FIGURE 3. Illustration of the tree \mathcal{T}_p for $n = 17$ and $k = 7$. Each node of $\mathcal{H}_{n,k}[\mathcal{T}_p]$, i.e., each cycle $C(x) \in \mathcal{C}_{n,k}$ is represented by the string $L(\langle x \rangle)$, corresponding to the lexicographically least connectable vertex on that cycle. The edges are oriented to point towards the lexicographically smaller string. Horizontal and vertical edges correspond to the first and second case in (10), respectively.

lies in the cycle $C(-^{n-2k}(10)^k)$ (corresponding to the bottom-right vertex in Figure 3). This proves that $\mathcal{H}_{n,k}[\mathcal{T}_p]$ is indeed connected.

It remains to argue that $\mathcal{H}_{n,k}[\mathcal{T}_p]$ is acyclic. Firstly, there can be no cycles in which the orientations of all edges agree, because of the lexicographically decreasing connectors. Secondly, if there was a cycle in which not all edges are oriented the same, then such a cycle would contain a node of out-degree 2, but in our graph $\mathcal{H}_{n,k}[\mathcal{T}_p]$, every node has out-degree exactly 1 by definition, except the node $C(-^{n-2k}(10)^k)$ which has out-degree 0 (as all vertices have only a single block). This proves that $\mathcal{H}_{n,k}[\mathcal{T}_p]$ is indeed acyclic. \square

Proof of Theorem 2 for Schrijver graphs $S(n, k) = S(n, k, 2)$. Our algorithm computes the Hamilton cycle that corresponds to the spanning tree $\mathcal{H}_{n,k}[\mathcal{T}_p]$ guaranteed by Lemma 5. It maintains the current vertex $x \in X_{n,k}$, and a variable $d \in \{-1, +1\}$ that controls the current direction of movement along the cycle. Specifically, the next vertex along the cycle $C(x)$ is obtained by setting

$$x \leftarrow \sigma^d(x) = \begin{cases} \sigma(x) & \text{if } d = +1, \\ \sigma^{-1}(x) & \text{if } d = -1, \end{cases} \quad (12)$$

i.e., in the first case we perform a cyclic right shift (move forward along the cycle) and in the second case the inverse operation, namely a cyclic left shift (move backward along the cycle). Clearly, these operations take linear time $\mathcal{O}(n)$.

Furthermore, in each step, the algorithm has to check whether the current vertex x equals one of the four vertices on a cycle $C_4(x', y') = (x', \sigma(x'), \sigma(y'), y')$ for one of the connectors $(x', y') \in \mathcal{T}_p$ (recall Lemma 3), and whether the next step according to (12) would move along an edge of this cycle. This is detected by the following conditions, followed by the corresponding alternative assignments instead of (12); see Figure 4:

- (i) if $x = x'$ and $d = +1$, then we set $x \leftarrow y'$, $d \leftarrow -1$;
- (ii) if $x = y'$ and $d = +1$, then we set $x \leftarrow x'$, $d \leftarrow -1$;
- (iii) if $x = \sigma(x')$ and $d = -1$, then we set $x \leftarrow \sigma(y')$, $d \leftarrow +1$;
- (iv) if $x = \sigma(y')$ and $d = -1$, then we set $x \leftarrow \sigma(x')$, $d \leftarrow +1$.




FIGURE 4. Illustration of the instructions (i)–(iv).




FIGURE 5. Hamilton cycles in different s -stable Kneser graphs computed by the algorithm from Section 3.

To decide whether $(x', y') \in \mathcal{T}_p$, we need to check whether $P(x') = M(\langle x' \rangle)$ (recall (11)), i.e., we need to compute the lexicographically smallest rotation of a string of length $\leq n$, which can be done using Booth's $\mathcal{O}(n)$ algorithm [Boo80]. \square

For more details, see the C++ implementation of our algorithm [cos]. Figure 5 shows several examples of Hamilton cycles computed by our algorithm.

4. GENERALIZING THE PROOFS TO s -STABLE KNESER GRAPHS $S(n, k, s)$

The proofs of Theorems 1 and 2 presented in the previous sections can be generalized straightforwardly from the case $s = 2$ (Schrijver graphs) to the more general case $s \geq 2$ (s -stable Kneser graphs), as follows.

We write $X_{n,k,s}$ for the vertices of $S(n, k, s)$ in their bitstring representation. It was shown in [MS08] that $|X_{n,k,s}| = \frac{n}{n-sk} \binom{n-sk}{k}$. The construction of the cycle factor $\mathcal{C}_{n,k}$ given in Section 2.1.2 remains unchanged, but we slightly generalize the definition of matched bits described in Section 2.1.3. Each 1-bit in $x \in X_{n,k,s}$ is matched to the next $s - 1$ many 0s to the right of it, so sk bits in total are matched (k of them 1s, and $(s - 1)k$ of them 0s) and the remaining $n - sk$ bits are unmatched. Consequently a glider in x has the form 10^{s-1} and occupies s positions $i, i + 1, \dots, i + s - 1$ (cf. (2)). Furthermore, a connector has the form

$$\begin{bmatrix} x \\ \sigma^{-1}(y) \end{bmatrix} = \begin{bmatrix} \dots 1 0^{s-1} \dots \\ \dots - 1 0^{s-1} \dots \end{bmatrix}$$

(cf. (3)). The decomposition (5a) generalizes accordingly to

$$x = -\gamma_1(10^{s-1})^{\beta_1/s} - \gamma_2(10^{s-1})^{\beta_2/s} \dots - \gamma_r(10^{s-1})^{\beta_r/s},$$

and the remaining definitions in Section 2.2.1 remain unchanged. The definition (7) generalizes to

$$\mathcal{U}_p := \{(x, y) \in \mathcal{X}_{n,k} \mid x_{p,p+1,\dots,p+s} = 10^{s-1} \text{ and } \sigma^{-1}(y)_{p,p+1,\dots,p+s} = -10^{s-1}\},$$

and then the proof of Lemma 4 goes through with minimal adjustments, specifically by generalizing (8) and (9) to

$$P(x) := L(x, p + s - 1 + \gamma' + \beta')$$

and

$$\tau(y) := \begin{cases} \sigma^{-2}(y) & \text{if } \gamma' > 1, \\ \sigma^{s-1}(y) & \text{if } \gamma' = 1 \text{ and } \beta > s, \\ \sigma^{\gamma'+s-1}(y) & \text{if } \gamma' = 1 \text{ and } \beta = s, \end{cases}$$

respectively. The remainder of the proofs remains unchanged.

REFERENCES

- [Boo80] K. S. Booth. Lexicographically least circular substrings. *Inform. Process. Lett.*, 10(4-5):240–242, 1980.
- [CF02] Y. Chen and Z. Füredi. Hamiltonian Kneser graphs. *Combinatorica*, 22(1):147–149, 2002.
- [Che03] Y. Chen. Triangle-free Hamiltonian Kneser graphs. *J. Combin. Theory Ser. B*, 89(1):1–16, 2003.
- [Che15] P.-A. Chen. On the multichromatic number of s -stable Kneser graphs. *J. Graph Theory*, 79(3):233–248, 2015.
- [cos] The Combinatorial Object Server: <http://www.combos.org/schrijver>.
- [GK76] C. Greene and D. J. Kleitman. Strong versions of Sperner’s theorem. *J. Combin. Theory Ser. A*, 20(1):80–88, 1976.
- [Joh11] J. R. Johnson. An inductive construction for Hamilton cycles in Kneser graphs. *Electron. J. Combin.*, 18(1):Paper 189, 12 pp., 2011.
- [Jon12] J. Jonsson. On the chromatic number of generalized stable Kneser graphs. <https://people.kth.se/~jakobj/doc/submitted/stablekneser.pdf>, 2012.
- [Kap43] I. Kaplansky. Solution of the “problème des ménages.”. *Bull. Amer. Math. Soc.*, 49:784–785, 1943.
- [Lov78] L. Lovász. Kneser’s conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, 25(3):319–324, 1978.
- [LP24] A. V. Ledezma and A. G. Pastine. s -stable Kneser graph are Hamiltonian. *arXiv:2401.15198*, 2024.
- [Meu11] F. Meunier. The chromatic number of almost stable Kneser hypergraphs. *J. Combin. Theory Ser. A*, 118(6):1820–1828, 2011.
- [MMN23] A. Merino, T. Mütze, and Namrata. Kneser graphs are Hamiltonian. In *STOC’23—Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 963–970. ACM, New York, 2023.
- [MNW21] T. Mütze, J. Nummenpalo, and B. Walczak. Sparse Kneser graphs are Hamiltonian. *J. Lond. Math. Soc. (2)*, 103(4):1253–1275, 2021.

- [MS08] T. Mansour and Y. Sun. On the number of combinations without certain separations. *European J. Combin.*, 29(5):1200–1206, 2008.
- [Sch78] A. Schrijver. Vertex-critical subgraphs of Kneser graphs. *Nieuw Arch. Wisk. (3)*, 26(3):454–461, 1978.