AN EXACT CHARACTERIZATION OF SATURATION FOR PERMUTATION MATRICES

Benjamin Aram Berendsohn*[1]

[1]*Freie Universität Berlin, Institut für Informatik, Berlin, Germany*
benjamin.berendsohn@fu-berlin.de

Submitted: Nov 10, 2021; Accepted: Jan 19, 2023; Published: Mar 15, 2023

Abstract. A 0-1 matrix M *contains* a 0-1 matrix *pattern* P if we can obtain P from M by deleting rows and/or columns and turning arbitrary 1-entries into 0s. The saturation function $\mathrm{sat}(P, n)$ for a 0-1 matrix pattern P indicates the minimum number of 1s in an $n \times n$ 0-1 matrix that does not contain P, but changing any 0-entry into a 1-entry creates an occurrence of P. Fulek and Keszegh recently showed that each pattern has a saturation function either in $\mathcal{O}(1)$ or in $\Theta(n)$. We fully classify the saturation functions of *permutation matrices*.

Keywords. Forbidden submatrices, saturation

Mathematics Subject Classifications. 05D99

1. Introduction

In this paper, all matrices are 0-1 matrices. For cleaner presentation, we write matrices with dots ($\bullet$) instead of 1s and spaces instead of 0s, for example:

$$\begin{pmatrix} 0&1&0 \\ 0&0&1 \\ 1&0&0 \end{pmatrix} = \begin{pmatrix} &\bullet& \\ &&\bullet \\ \bullet&& \end{pmatrix}$$

In line with this notation, we call a row or column *empty* if it only contains 0s. Furthermore, we refer to changing an entry from 0 to 1 as *adding* a 1-entry, and to the reverse as *removing* a 1-entry.

A *pattern* is a matrix that is not all-zero. A matrix M *contains* a pattern P if we can obtain P from M by deleting rows and/or columns, and removing arbitrary 1-entries. If M does not contain P, we say M *avoids* P. Matrix pattern avoidance can be seen as a generalization of two other well-known areas in extremal combinatorics. Pattern avoidance in permutations (see, e.g., Vatter's survey [Vat14]) corresponds to the case where both M and P are permutation

*Supported by DFG grant KO 6140/1-1.

matrices; and forbidden subgraphs in bipartite graphs correspond to avoiding a pattern P and all other patterns obtained from P by permutation of rows and/or columns.[1] There are also close connections to the extremal theory of ordered graphs [PT06] and posets [GNPV22].

A classical question in extremal graph theory is to determine the maximum number of edges in an n-vertex graph avoiding a fixed pattern graph H. The corresponding problem in forbidden submatrix theory is determining the maximum *weight* (number of 1s) of an $m \times n$ matrix avoiding the pattern P, denoted by $\mathrm{ex}(P, m, n)$. We call $\mathrm{ex}(P, n) = \mathrm{ex}(P, n, n)$ the *extremal function* of the pattern P. The study of the extremal function originates in its applications to (computational) geometry [Mit87, Für90, BG91]. A systematic study initiated by Füredi and Hajnal [FH92] has produced numerous results (e.g., [Kla00, Kla01, MT04, Tar05, Kesz09, Ful09, Gen09, Pet11a, Pet11b]), and further applications in the analysis of algorithms have been discovered [Pet10, CGK+15].

A natural counterpart to the extremal problem is the *saturation problem*. A matrix M is *saturating* for a pattern P, or P-*saturating* if it avoids P and is maximal in this respect, i.e., adding a 1-entry anywhere creates an occurrence of P. Clearly, $\mathrm{ex}(P, m, n)$ can also be defined as the maximum weight of an $m \times n$ matrix that is P-saturating. The function $\mathrm{sat}(P, m, n)$ indicates the *minimum* weight of an $m \times n$ matrix that is P-saturating. We focus on square matrices and the *saturation function* $\mathrm{sat}(P, n) = \mathrm{sat}(P, n, n)$.

The saturation problem for matrix patterns was first considered by Brualdi and Cao [BC21] as a counterpart of saturation problems in graph theory.[2] Fulek and Keszegh [FK21] started a systematic study. They proved that, perhaps surprisingly, every pattern P satisfies $\mathrm{sat}(P, n) \in \mathcal{O}(1)$ or $\mathrm{sat}(P, n) \in \Theta(n)$, where the hidden constants depend on P. This is in stark contrast to the extremal problem, where a wide range of different orders of magnitude is attained by various patterns (from linear and quasi-linear [Kesz09, Pet11a], to nearly quadratic [ARSz99]). Fulek and Keszegh also present large classes of patterns with linear saturation functions. For our purposes, their most important result is that every *decomposable* pattern has linear saturation function. We call a pattern P decomposable if it has the form

$$\begin{pmatrix} A & \mathbf{0} \\ \mathbf{0} & B \end{pmatrix} \text{ or } \begin{pmatrix} \mathbf{0} & A \\ B & \mathbf{0} \end{pmatrix}$$

for two matrices $A, B \neq \mathbf{0}$, where $\mathbf{0}$ denotes an all-0 matrix of the appropriate size. Otherwise, we call P *indecomposable*. Patterns of the first form $\left(\begin{smallmatrix} A & \mathbf{0} \\ \mathbf{0} & B \end{smallmatrix}\right)$ are called *sum decomposable*, and patterns not of that form are called *sum indecomposable*.[3]

Fulek and Keszegh also found a single non-trivial pattern with bounded saturation function (Q, pictured in Figure 1.1), and conjectured that there are many more. Geneson [Gen21] recently confirmed this by proving that almost all *permutation matrices* have bounded saturation function. A permutation matrix is a matrix with exactly one 1-entry in each row and each column. A different class of matrices with bounded saturation function, containing both permutation matrices and non-permutation matrices were found recently by the author [Ber20].[4]

[1] For this, we interpret the M and P as adjacency matrices of bipartite graphs.

[2] We refer to [FK21] for references to graph saturation results.

[3] These terms are derived from the theory of permutation patterns (see, e.g., Vatter [Vat14]). We are not aware of a standard term for this property in the context of 0-1 matrices.

[4] These results have been incorporated into this paper in Sections 1.1 and 2.

$$Q = \begin{pmatrix} & \bullet & & & \\ & & & & \bullet \\ & & \bullet & & \\ \bullet & & & & \\ & & & \bullet & \end{pmatrix}$$

Figure 1.1: The matrix with saturation function $\mathcal{O}(1)$ found by Fulek and Keszegh [FK21].

In this paper, we show that, in fact, *all* indecomposable permutation matrices have bounded saturation function. This completes the characterization of permutation matrices in terms of their saturation function.

Theorem 1.1. *A permutation matrix has linear saturation function if and only if it is decomposable.*

A simple generalization of the technique that Fulek and Keszegh used to prove that $\operatorname{sat}(Q, n)$ is bounded implies the following: To prove Theorem 1.1, it is sufficient to find a *vertical witness* for every indecomposable permutation matrix P, where we define a vertical witness for P to be a matrix M (of arbitrary size) that avoids P, has an empty row, and adding a 1-entry in that empty row creates an occurrence of P in M.

We therefore construct vertical witnesses for all permutation matrices. Our constructions are based on the fact that every indecomposable permutation matrix contains a *spanning oscillation* (defined in Section 1.2).

We also generalize a partial result to a class that contains non-permutation patterns:

Theorem 1.2. *Let P be a pattern that contains four 1-entries x_1, x_2, x_3, x_4 such that for each $i \in [4]$, there are no other 1-entries in the same row or column as x_i, and x_i is in the first or last row or column, and x_1, x_2, x_3, x_4 form one of the two patterns*

$$\begin{pmatrix} & \bullet & & \\ & & & \bullet \\ \bullet & & & \\ & & \bullet & \end{pmatrix}, \begin{pmatrix} & & \bullet & \\ \bullet & & & \\ & & & \bullet \\ & \bullet & & \end{pmatrix}.$$

Then $\operatorname{sat}(P, n) \in \mathcal{O}(1)$.

In Section 1.1, we define and discuss (vertical) witnesses, and in Section 1.2, we define spanning oscillations. In Section 1.3, we present the structure of the proof of Theorem 1.1. In Section 1.4 we introduce an alternative characterization of pattern containment that simplifies our proofs.

Section 1.5 gives an introduction to the witness-construction techniques used in the following chapters. In Sections 2 to 4, we construct vertical witnesses for all permutation matrices, based on different types of spanning oscillations, which proves Theorem 1.1. We also prove Theorem 1.2 in Section 2.

We now introduce conventions and notations used throughout the paper. Some more definitions needed for Sections 2 to 4 will be introduced in Section 1.4.

We identify 1-entries in an $m \times n$ matrix M as their positions $(i, j) \in [m] \times [n]$, where i is the row of the 1-entry (from top to bottom), and j is its column (from left to right). $E(M)$ denotes the set of 1-entries in M. For two 1-entries $x = (i, j) \in E(M)$ and $x' = (i', j') \in E(M)$, we write $x <_{\mathrm{v}} x'$ if $i < i'$ and $x <_{\mathrm{h}} x'$ if $j < j'$. Define $x \leqslant_{\mathrm{v}} x'$ and $x \leqslant_{\mathrm{h}} x'$ analogously. We also say x is *above* x' if $x <_{\mathrm{v}} x'$, and use *below*, *to the right*, and *to the left* similarly.

In a permutation matrix P, we denote the leftmost (rightmost, topmost, bottommost) 1-entry of P by ℓ_P (r_P, t_P, b_P). Note that if P is an indecomposable $k \times k$ permutation matrix with $k > 1$, then these four 1-entries are pairwise distinct.

Let M be an arbitrary matrix. Denote by $\mathrm{rot}(M)$ the matrix obtained by rotating M 90 degrees clockwise, denote by $\mathrm{rev}(M)$ the matrix obtained by reversing all rows of M, and denote by $\mathrm{trans}(M)$ the transpose of M, i.e., the matrix obtained by swapping the roles of rows and columns.[5]

If A is a $k \times m$ matrix and B is a $k \times n$ matrix, then the *horizontal concatenation* (A, B) is the $k \times (m+n)$ matrix M where $E(M) = E(A) \cup \{(i, j+m) \mid (i,j) \in E(B)\}$. Intuitively, M is obtained by placing A to the left of B. The horizontal concatenation $(A_1, A_2, \dots)$ of a sequence of matrices with the same height is defined accordingly.

1.1. Witnesses

Let P be a matrix pattern without empty rows or columns. An *explicit witness*[6] for P is a matrix M that is P-saturating and contains at least one empty row and at least one empty column.

Lemma 1.3 ([FK21]). *For each pattern P without empty rows and columns, we have $\mathrm{sat}(P, n) \in \mathcal{O}(1)$ if and only if P has an explicit witness.*

Proof. Suppose $\mathrm{sat}(P, n) \leqslant c_P$ for all $n \in \mathbb{N}$. Then there exists a $(c_P + 1) \times (c_P + 1)$ P-saturating matrix M with at most c_P 1-entries. Clearly, M has an empty row and an empty column, so M is an explicit witness for P.

Now suppose that P has an $m_0 \times n_0$ explicit witness M of weight w. We can replace the empty row (column) in M by an arbitrary number of empty rows (columns), and the resulting (arbitrarily large) matrix will still be P-saturating. As such, $\mathrm{sat}(P, m, n) \leqslant w$ for all $m \geqslant m_0$ and $n \geqslant n_0$. Note that it is critical here that P has no empty rows or columns. Otherwise, inserting empty rows or columns into M might create an occurrence of P. □

We call a row (column) of a matrix M *P-expandable* if the row (column) is empty and adding a single 1-entry anywhere in that row (column) creates a new occurrence of P in M. An explicit witness for P is thus a saturating matrix with at least one P-expandable row and at least one P-expandable column. We define a *witness* for P (used implicitly by Fulek and Keszegh) as a matrix that avoids P and has at least one P-expandable row and at least one P-expandable column. Clearly, an explicit witness is a witness. The following lemma shows that finding a witness is sufficient to show that $\mathrm{sat}(P, n) \in \mathcal{O}(1)$.

Lemma 1.4. *If a pattern P without empty rows or columns has an $m_0 \times n_0$ witness, then P has an $m_0 \times n_0$ explicit witness.*

[5] We do not use the common superscript $^{\mathrm{T}}$, as it will later be used with the meaning "top".

[6] An explicit witness is what Fulek and Keszegh [FK21] simply call a *witness*.

Proof. Let M be an $m_0 \times n_0$ witness for P. If M is P-saturating, then we are done. Otherwise, there must be a 0-entry (i, j) in M that can be changed to 1 without creating an occurrence P. Choose one such 0-entry and turn it into 1. Note that (i, j) cannot be contained in an expandable row or column of M, so the resulting matrix is still a witness. Thus, we obtain an explicit witness after repeating this step at most $m_0 \cdot n_0$ times. □

1.1.1 Vertical and horizontal witnesses

Fulek and Keszegh also considered the asymptotic behavior of the functions $\operatorname{sat}(P, m_0, n)$ and $\operatorname{sat}(P, m, n_0)$, where m_0 and n_0 are fixed. The dichotomy of $\operatorname{sat}(P, n)$ also holds in this setting:

Theorem 1.5 ([FK21, Parts of Theorem 1.3])**.** *For every pattern P, and constants m_0, n_0,*

(i) either $\operatorname{sat}(P, m_0, n) \in \mathcal{O}(1)$ *or* $\operatorname{sat}(P, m_0, n) \in \Theta(n)$*;*

(ii) either $\operatorname{sat}(P, m, n_0) \in \mathcal{O}(1)$ *or* $\operatorname{sat}(P, m, n_0) \in \Theta(m)$.

We can adapt the notion of witnesses in order to classify $\operatorname{sat}(P, m_0, n)$ and $\operatorname{sat}(P, m, n_0)$. Let P be a matrix pattern without empty rows or columns. A *horizontal (vertical) witness* for P is a matrix M that avoids P and contains an expandable column (row).[7] Clearly, P has a horizontal witness with m_0 rows if and only if $\operatorname{sat}(P, m_0, n)$ is bounded; and P has a vertical witness with n_0 columns if and only if $\operatorname{sat}(P, m, n_0)$ is bounded. Further note that M is a witness for P if and only if M is both a horizontal witness and a vertical witness.

We now prove that we can essentially restrict our attention to the classification of the functions $\operatorname{sat}(P, m_0, n)$ and $\operatorname{sat}(P, m, n_0)$. The following two lemmas are a generalization of the technique used by Fulek and Keszegh to prove that $\operatorname{sat}(Q, n) \in \mathcal{O}(1)$ for the pattern Q depicted in Figure 1.1.

Lemma 1.6. *Let P be a matrix pattern without empty rows or columns, and only one 1-entry in the last row (column). Let W be a horizontal (vertical) witness for P. Then, appending an empty row (column) to W again yields a horizontal (vertical) witness.*

Proof. We prove the lemma for horizontal witnesses with a row appended. The other case follows by symmetry. Let W be an $m_0 \times n_0$ horizontal witness for P, where the j-th column of W is expandable. Let W' be the matrix obtained by appending a row to W. Clearly, W' still does not contain P. Moreover, adding an entry in W' at (i, j) for any $i \neq n_0 + 1$ creates a new occurrence of P. It remains to show that adding an entry at $(n_0 + 1, j)$ creates an occurrence of P.

We know that adding an entry at (n_0, j) in W' creates an occurrence of P, say at positions $I \subseteq [n_0]^2$. Since P has only one entry in the last row, all positions $(i', j') \in I \setminus \{(n_0, j)\}$ satisfy $i' < n_0 + 1$. Thus, adding a 1-entry at $(n_0 + 1, j)$ instead of (n_0, j) creates an occurrence of P at positions $I \setminus \{(n_0, j)\} \cup \{(n_0 + 1, j)\}$. Thus, W' is a horizontal witness. □

[7] A horizontal witness can be expanded horizontally, a vertical witness can be expanded vertically.

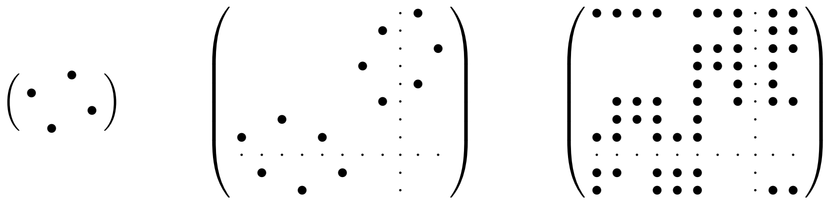

Figure 1.2: A pattern (left), a witness (middle) and an explicit witness (right) for the pattern. The small dots indicate the expandable row/column.

Lemma 1.7. *Let P be a indecomposable pattern without empty rows or columns, and with only one 1-entry in the last row and one 1-entry in the last column. Then* $\operatorname{sat}(P,n) \in \mathcal{O}(1)$ *if and only if there exist constants* m_0, n_0 *such that* $\operatorname{sat}(P,m_0,n) \in \mathcal{O}(1)$ *and* $\operatorname{sat}(P,m,n_0) \in \mathcal{O}(1)$.

Proof. Suppose that $\operatorname{sat}(P,n) \in \mathcal{O}(1)$. Then P has an $m_0 \times n_0$ witness M by Lemma 1.3, and thus $\operatorname{sat}(P,m_0,n)$ is at most the weight of M, for every $n \geqslant n_0$. Similarly, $\operatorname{sat}(P,m,n_0) \in \mathcal{O}(1)$.

Now suppose that $\operatorname{sat}(P,m_0,n) \in \mathcal{O}(1)$ and $\operatorname{sat}(P,m,n_0) \in \mathcal{O}(1)$. Then, for some m_1, n_1, there exists an $m_0 \times n_1$ horizontal witness W_{H} and an $m_1 \times n_0$ vertical witness W_{V}. Consider the following $(m_0+m_1) \times (n_0+n_1)$ matrix, where $\mathbf{0}_{m \times n}$ denotes the all-0 $m \times n$ matrix:

$$W = \begin{pmatrix} \mathbf{0}_{m_0 \times n_0} & W_{\mathrm{H}} \\ W_{\mathrm{V}} & \mathbf{0}_{m_1 \times n_1} \end{pmatrix}$$

We first show that W does not contain P. Suppose it does. Since P is contained neither in W_{H} nor in W_{V}, an occurrence of P in W must contain 1-entries in both the bottom left and top right quadrant. But then P is decomposable, a contradiction.

By Lemma 1.6, $W'_{\mathrm{V}} = (W_{\mathrm{V}}, \mathbf{0}_{m_1 \times n_1})$ is a vertical witness, and $W'_{\mathrm{H}} = \binom{W_{\mathrm{H}}}{\mathbf{0}_{m_1 \times n_1}}$ is a horizontal witness. The expandable row in W'_{V} and the expandable column in W'_{H} are both also present in W. This implies that W is a witness for P, so $\operatorname{sat}(P,n) \in \mathcal{O}(1)$. □

Figure 1.2 shows an example of a witness, constructed with Lemma 1.7, using vertical and horizontal witnesses presented later in Section 2, and an explicit witness constructed using Lemma 1.4.

Observe that the transformations rev, rot, and trans all preserve witnesses. However, the latter two change vertical witnesses to horizontal witnesses, and vice versa. Formally:

Observation 1.8. *Let P be a matrix with a vertical witness W. Then* $\operatorname{rev}(W)$ *is a vertical witness of* $\operatorname{rev}(P)$*,* $\operatorname{rot}(W)$ *is a horizontal witness of* $\operatorname{rot}(P)$*, and* $\operatorname{trans}(W)$ *is a horizontal witness of* $\operatorname{trans}(P)$. □

Recall that our goal is to show that every indecomposable permutation matrix has a witness. Since indecomposable permutation matrices are closed under transposition, Lemma 1.7 and Observation 1.8 imply that it suffices to find a *vertical* witness for each indecomposable permutation matrix. The same is true for every class of matrices satisfying the conditions of Lemma 1.7 that is closed under transposition or 90-degree clockwise rotation. This is useful to prove Theorem 1.2.

Lemma 1.9. *Let $\mathcal{P}$ be a class of indecomposable patterns without empty rows or columns, and with only one 1-entry in the last row and one 1-entry in the last column. If $\mathcal{P}$ is closed under transposition or 90-degree clockwise rotation and each pattern in $\mathcal{P}$ has a vertical witness, then* $\operatorname{sat}(P,n) \in \mathcal{O}(1)$ *for each* $P \in \mathcal{P}$.

Proof. Suppose that $\mathcal{P}$ is closed under transposition and each $P \in \mathcal{P}$ has a vertical witness. By Lemma 1.7, it suffices to show that each pattern in $\mathcal{P}$ also has a horizontal witness. Let $P \in \mathcal{P}$. Then $\operatorname{trans}(P) \in \mathcal{P}$ has a vertical witness W. By Observation 1.8, $\operatorname{trans}(W)$ is a horizontal witness for $\operatorname{trans}(\operatorname{trans}(P)) = P$.

The case that $\mathcal{P}$ is closed under 90-degree rotation can be handled analogously. □

1.2. Spanning oscillations

We now introduce *spanning oscillations*, a class of substructures that characterizes indecomposable permutation matrices.

For a permutation matrix P, the *permutation graph* G_P of the underlying permutation can be defined as follows: The vertex set is $E(P)$, and two 1-entries $x, y \in E(P)$ have an edge between them if x is below and to the left of y (or vice versa).

An *oscillation* in a permutation matrix of P is a sequence $X = (x_1, x_2, \dots, x_m)$ of distinct 1-entries in P such that X forms an induced path in G_P, i.e., there is an edge between x_i and x_{i+1} for each $i \in [m-1]$, and no other edges between 1-entries in X. Oscillations have been studied before in several contexts [Pra73, BRV08, Vat11]. Vatter showed that a permutation matrix P is sum indecomposable if and only if it has an oscillation that starts with ℓ_P and ends with r_P [Vat11, Propositions 1.4, 1.7]. Our characterization of indecomposable permutations is very similar. Call an oscillation $X = (x_1, x_2, \dots, x_m)$ *spanning* if $\{x_1, x_2\} = \{\ell_P, t_P\}$ and $\{x_{m-1}, x_m\} = \{b_P, r_P\}$.

Lemma 1.10. *Let P be a sum indecomposable permutation matrix such that t_P is to the left of b_P or ℓ_P is above r_P. Then P has a spanning oscillation.*

Proof. We write ℓ, t, b, r for ℓ_P, t_P, b_P, r_P. By symmetry, we can assume that t is to the left of b (otherwise, replace P by $\operatorname{trans}(P)$, noting that $G_P = G_{\operatorname{trans}(P)}$). Recall that ℓ, t, b, r are pairwise distinct, as P is indecomposable and not 1×1.

Since P is sum indecomposable, by the result of Vatter mentioned above, it has an oscillation $X' = (x'_1, x'_2, \dots, x'_m)$ with $x'_1 = \ell, x'_m = r$. Suppose first that t occurs in X'. Since G_P has an edge between ℓ and t, and X is an *induced* path in G_P, this means that $x'_2 = t$. Otherwise, note that t is connected in G_P to precisely those 1-entries that are to the left of t. Let i be maximal such that x'_i is to the left of t. If $i = 1$, then $(t, \ell, x'_2, \dots, x'_m)$ is an induced path in G_P. Otherwise, $\ell, t, x'_i, \dots, x'_m$ is an induced path in G_P. In either case, we have an oscillation $X'' = (x''_1, x''_2, \dots, x''_m)$ that starts with $\{\ell, t\}$ and ends with r.

It remains to make sure that b is among the last two 1-entries in the oscillation. If b occurs in X'', then $x''_{m-1} = b$, since X'' is an induced path. Otherwise, let j be minimal such that x_j is to the right of b. If $j = m$, then $X = (x''_1, x''_2, \dots, x''_{m-1}, r, b)$ is an induced path in G_P. Otherwise, $X = (x''_1, x''_2, \dots, x''_j, b, r)$ is an induced path in G_P. Since ℓ, t are both to the left of b, we have $j > 2$, so X is a spanning oscillation. □

We obtain the following characterization of indecomposable permutation matrices.

Corollary 1.11. *A permutation matrix P is indecomposable if and only if P or $\mathrm{rev}(P)$ has a spanning oscillation or P is the 1×1 permutation matrix.*

Proof. First, assume P is indecomposable. If t_P is to the left of b_P, then Lemma 1.10 implies that P has a spanning oscillation. If t_P is to the right of b_P, then Lemma 1.10 implies that $\mathrm{rev}(P)$ has a spanning oscillation. If $t_P = b_P$, then P is 1×1.

Second, assume P has a spanning oscillation. Then P is sum indecomposable. Suppose P is decomposable, then P has the form $\left(\begin{smallmatrix} \mathbf{0} & B \\ A & \mathbf{0} \end{smallmatrix}\right)$, so t is to the right of b and ℓ is below r. But then ℓ, b, t, r form the complete bipartite graph $K_{2,2}$ in G_P, implying that P has no spanning oscillation, a contradiction. A symmetric argument shows that P is indecomposable if $\mathrm{rev}(P)$ has a spanning oscillation. □

Spanning oscillations have a very rigid structure, which we now describe more explicitly, in terms of relative positions of 1-entries. Let P be a permutation matrix and $X = (x_1, x_2, \dots, x_m)$ be a spanning oscillation of P. For $2 \leqslant i \leqslant m-1$, call x_i an *upper* 1-entry if x_i is above and to the right of x_{i-1} and x_{i+1}, and call x_i a *lower* 1-entry if x_i is below and to the left of x_{i-1} and x_{i+1}. Since G_P contains the edges $\{x_{i-1}, x_i\}$ and $\{x_i, x_{i+1}\}$, but not the edge $\{x_{i-1}, x_{i+1}\}$, every 1-entry (except x_1, x_m) is either upper or lower. Clearly, upper and lower 1-entries alternate, i.e., x_i is upper if and only if x_{i+1} is lower, for $2 \leqslant i < m-1$. It is convenient to also call ℓ_P, b_P lower 1-entries and t_P, r_P upper 1-entries. We then have:

Observation 1.12. *Let P be a permutation matrix and $X = (x_1, x_2, \dots, x_m)$ be a spanning oscillation of P. If $x_1 = \ell_P$, then all x_i with odd i are lower 1-entries, and all x_i with even i are upper 1-entries. If $x_1 = t_P$, then all x_i with odd i are upper 1-entries, and all x_i with even i are lower 1-entries.* □

It is easy to see that, if $x_1 = \ell_P$, then x_3, x_4 must be below and to the right of x_1. By induction, and by considering symmetric cases, we can prove:

Observation 1.13. *Let P be a permutation matrix and $X = (x_1, x_2, \dots, x_m)$ be a spanning oscillation of P. Then x_i is above and to the left of x_j for each $i \in [m-2]$ and $i+2 \leqslant j \leqslant m$.*

This leaves us with only two possible forms of spanning oscillations for each length m, see Figure 1.3. Observe that spanning oscillations are preserved by transposition and 180-degree rotation, in the following sense.

Let P be a $k \times k$ permutation matrix and $X = (i_1, j_1), (i_2, j_2), \dots, (i_\ell, j_\ell)$ a spanning oscillation of P. Define

$$\begin{aligned} \mathrm{trans}(X) &= (j_1, i_1), (j_2, i_2), \dots, (j_\ell, i_\ell), \text{ and} \\ \mathrm{rot}^2(X) &= (k - i_\ell, k - j_\ell), (k - i_{\ell-1}, k - j_{\ell-1}), \dots, (k - i_1, k - j_1). \end{aligned}$$

It is easy to see that $\mathrm{trans}(X)$ is a spanning oscillation of $\mathrm{trans}(P)$ and $\mathrm{rot}^2(X)$ is a spanning oscillation of $\mathrm{rot}^2(P)$.

A spanning oscillation $X = (x_1, x_2, \dots, x_m)$ is *tall* if the following two properties are satisfied for each $2 \leqslant i \leqslant m-2$ where x_i is an upper 1-entry.

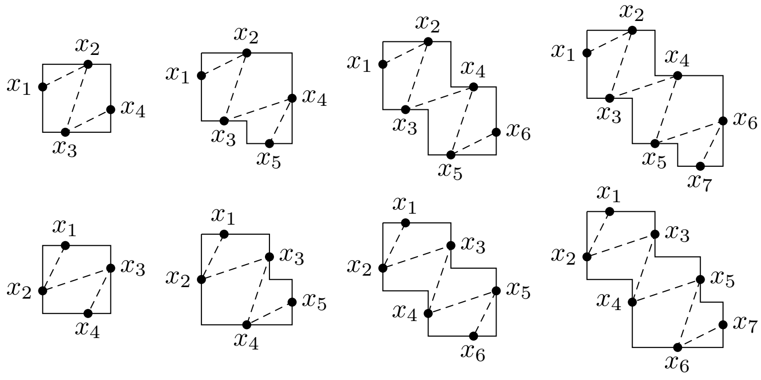

Figure 1.3: The spanning oscillations of length m, for $m = 4, 5, 6, 7$. The dashed line segments indicate the edges of the permutation graph. The borders indicate the possible positions for other 1-entries if the spanning oscillation is tall (top row) or wide (bottom row).

(i) P has no 1-entry that is below x_{i+1} and to the left of x_i.

(ii) P has no 1-entry that is above x_i and to the right of x_{i+1}.

A spanning oscillation X is *wide* if $\mathrm{trans}(X)$ is tall. We now show that we can always assume that a minimum-length spanning oscillation is tall (or wide).

Lemma 1.14. *Let P be a permutation matrix and $X = (x_1, x_2, \ldots, x_m)$ be a spanning oscillation of P of minimum length m. Then P has a tall spanning oscillation of length m that starts with x_1, x_2 and ends with x_{m-1}, x_m.*

Proof. Suppose X is not tall, so it violates (i) or (ii) at some index i with $2 \leqslant i \leqslant m-2$. We now show how to construct a spanning oscillation X' of length m that starts with x_1, x_2, ends with x_{m-1}, x_m, and violates (i) or (ii) less often than X. Repeating this, we eventually obtain a tall spanning oscillation.

Suppose first that X violates (i) at index i. Then x_i is an upper 1-entry, and there is a $y \in E(P)$ such that y is below x_{i+1} and to the left of x_i. Assume y is the bottommost such 1-entry. Note $y \notin \{\ell_P, b_P\}$, and that x_{i+2} is above x_{i+1} by Observation 1.12.

Let j be minimal such that x_j is to the right of y. Since $\ell_P <_{\mathrm{h}} y <_{\mathrm{h}} x_i$, we have $2 \leqslant j \leqslant i$. Let k be maximal such that x_k is above y. Since $x_{i+2} <_{\mathrm{v}} y <_{\mathrm{v}} b_P$, we have $i+2 \leqslant k \leqslant m-1$.

By Observation 1.13, x_j is above and to the left of x_k, meaning that both x_j and x_k are above and to the right of y. Thus, the sequence $X' = (x_1, x_2, \ldots, x_j, y, x_k, x_{k+1}, \ldots, x_m)$ is a path. We now show that X' is an *induced* path. Let $j' < j$. By definition of j, we know that $x_{j'}$ is to the left of y. By Observation 1.13, $x_{j'}$ is above x_{i+1}, implying that $x_{j'}$ is above y. Thus, G_P has no edge between $x_{j'}$ and y. Similarly, we can prove that there is no edge between y and $x_{k'}$ for each $k' > k$.

Since $2 \leqslant j$ and $k \leqslant m-1$, we know that X' starts with x_1, x_2 and ends with x_{m-1}, x_m, implying that X' is a spanning oscillation.

By assumption, P has no spanning oscillation shorter than P, so X must have length m, implying that $j = i$ and $k = i+2$. Further, X' does not violate (i) at index i, since, by choice of y, there are no 1-entries below y and to the left of $x_j = x_i$. Thus, X' has strictly less overall violations of (i) or (ii) than X.

The second case, where X violates (ii), can be proven symmetrically. $\square$

Clearly, the statement of Lemma 1.14 is also true when replacing "tall" with "wide", using the same proof on $\mathrm{trans}(P)$.

1.3. Structure of the main proof

We divide the proof of Theorem 1.1 into three cases, proven in Sections 2 to 4. In Section 2, we handle the special case of length-4 spanning oscillations:

Lemma 1.15. *Each permutation matrix with a spanning oscillation of length 4 has a vertical witness.*

In Section 3, we prove:

Lemma 1.16. *Each permutation matrix P with a wide spanning oscillation of length $m \geqslant 5$ that starts with t_P has a vertical witness.*

The final and most involved case is treated in Section 4:

Lemma 1.17. *Each permutation matrix P with a tall spanning oscillation of even length $m \geqslant 6$ that starts with ℓ_P has a vertical witness.*

It is not immediately obvious that Lemmas 1.15 to 1.17 cover all indecomposable permutation matrices. We now show that this is the case.

Corollary 1.18. *Every indecomposable permutation matrix has a vertical witness.*

Proof. Let P be an indecomposable permutation matrix. If P is 1×1, any all-zero matrix is a witness of P. Otherwise, one of P and $\mathrm{rev}(P)$ has a spanning oscillation X by Corollary 1.11. By Observation 1.8, it suffices to find a vertical witness for either P or $\mathrm{rev}(P)$, so without loss of generality, assume that X is a spanning oscillation of P, and that X has minimum length m. If $m = 4$, we can apply Lemma 1.15. If $m \geqslant 5$ and X starts with t_P, then Lemma 1.14 implies that P also has a wide spanning oscillation of size m that starts with t_P, so we can apply Lemma 1.16.

Now assume $m \geqslant 5$ and X starts with ℓ_P. If m is even, we can apply Lemma 1.17, since by Lemma 1.14 we can assume that X is tall. Otherwise, if m is odd, Observation 1.12 implies that X ends with b_P. This means that the spanning oscillation $\mathrm{rot}^2(X)$ of $\mathrm{rot}^2(P)$ starts with $t_{\mathrm{rot}^2(P)}$, so we can apply Lemma 1.16 to obtain a witness W' of $\mathrm{rot}^2(P)$. Observation 1.8 implies that $\mathrm{rot}^2(W')$ is a witness of P. $\square$

1.4. Embeddings

In the following sections, we use an alternative definition of pattern containment based on sets of 1-entries. Let P be a pattern and M be a matrix. We say a function $\phi \colon E(P) \to E(M)$ is an *embedding* of P into M if for $x, y \in E(P)$ we have $x <_{\mathrm{h}} y \Leftrightarrow \phi(x) <_{\mathrm{h}} \phi(y)$ and $x <_{\mathrm{v}} y \Leftrightarrow \phi(x) <_{\mathrm{v}} \phi(y)$.

Note that if we allow empty rows or columns in P, then $E(P)$ does not determine P, since appending an empty row or column to P does not change $E(P)$. This means that the existence of an embedding of P into M does not necessarily imply that P is contained in M. However, we only consider patterns without empty rows or columns in this paper, and in that case, equivalence holds.

Lemma 1.19. *Let P, M be matrices, and let P have no empty rows or columns. Then P is contained in M if and only if there is an embedding of P into M.*

A proof of Lemma 1.19 is provided in Appendix A. We now introduce some notations used in the following sections.

Let $x = (i, j)$, $y = (i', j')$ be two 1-entries. The *horizontal distance* between x and y is $\mathrm{d}^{\mathrm{h}}(x, y) = |i - i'|$, and the *vertical distance* between x and y is $\mathrm{d}^{\mathrm{v}}(x, y) = |j - j'|$. The *width* $\mathrm{width}(A)$ (resp. *height* $\mathrm{height}(A)$) of a set $A \subseteq E(M)$ is the maximum horizontal (resp. vertical) distance between two 1-entries in A.

Let ϕ be an embedding of P into M. We say a row (column) is *hit* by ϕ if $\phi(x)$ is in that row (column) for some 1-entry $x \in E(P)$. We define variants of the above notions that only "count" rows and columns of M that are hit by ϕ. This will be useful when we have partial information about ϕ, or when we know that certain rows/columns are empty and thus cannot be hit by ϕ. Let $\mathrm{d}^{\mathrm{v}}_{\phi}((i, j), (i', j'))$ be the number of rows i'' such that i is hit by ϕ and $i < i'' \leqslant i'$. Similarly, let $\mathrm{d}^{\mathrm{h}}_{\phi}((i, j), (i', j'))$ be the number of columns j'' such that j is hit by ϕ and $j < j'' \leqslant j'$. For $A \subseteq E(M)$, let $\mathrm{width}_{\phi}(A) = \max_{x,y \in A} \mathrm{d}^{\mathrm{h}}_{\phi}(x, y)$, and $\mathrm{height}_{\phi}(A) = \max_{x,y \in A} \mathrm{d}^{\mathrm{v}}_{\phi}(x, y)$.

Observation 1.20. *Let ϕ be an embedding of P into M, let $x, y \in E(P)$, and let $\phi(x), \phi(y) \in A \subseteq E(M)$. Then*

$$\mathrm{d}^{\mathrm{h}}(x, y) = \mathrm{d}^{\mathrm{h}}_{\phi}(\phi(x), \phi(y)) \leqslant \mathrm{d}^{\mathrm{h}}(\phi(x), \phi(y)) \leqslant \mathrm{width}(A); \text{ and}$$
$$\mathrm{d}^{\mathrm{v}}(x, y) = \mathrm{d}^{\mathrm{v}}_{\phi}(\phi(x), \phi(y)) \leqslant \mathrm{d}^{\mathrm{v}}(\phi(x), \phi(y)) \leqslant \mathrm{height}(A). \qquad \square$$

1.5. Constructing witnesses

In this subsection, we describe intuitively how our various constructions work. Recall that a vertical witness for a matrix P avoids P and has a P-expandable row. As we will see in the following, there are many different ways of constructing a matrix with a P-expandable row; ensuring that it also avoids P is the hard part.

Suppose we have a permutation matrix P. Place a copy of P in the middle of a large matrix, and remove the leftmost 1-entry ℓ of the copy (Figure 1.4 *a, b*). Then, adding a 1-entry at the position of ℓ or in the same row to the left of ℓ will complete an occurrence of P. Thus, we have constructed the "left part" of an expandable row. We can similarly use a copy of P missing the

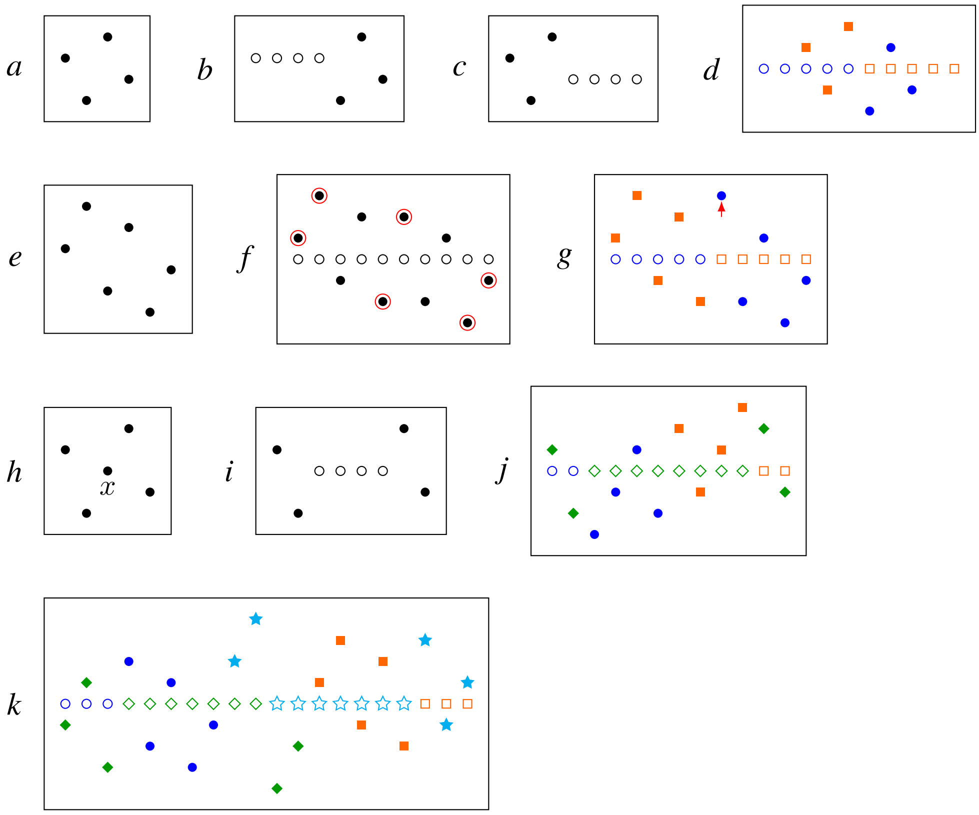

Figure 1.4: Constructing witnesses as described in Section 1.5. Empty dots indicate positions that complete the pattern, i.e., (partial) expandable rows. Colors/shapes of dots (1-entries) in *d, g, j*, and *k* indicate different parts of the construction. Red circles in *f* indicate the occurrence of the pattern.

rightmost 1-entry r to get the "right part" of an expandable row (Figure 1.4 *c*). If we now place the copy without ℓ (call it L) to the right of the copy without r (call it R), we can create a whole expandable row (Figure 1.4 *d*).

It turns out that the resulting matrix avoids P for a large class of patterns, in particular for patterns with a spanning oscillation of length 4. We prove this in Section 2.

The smallest example where the construction does contain P is shown in Figure 1.4 *e, f*. However, observe that changing the vertical position of a 1-entry preserves the expandable row as long as the vertical order within both L and R is maintained. Thus, we can try to vertically "stretch" L and/or R to make the matrix avoid P. In the given example, moving the top entry of L up one row suffices (Figure 1.4 *g*).

When stretching does not help, another option is to construct larger matrices in the following way. Consider some 1-entry x of P that is not the leftmost or rightmost 1-entry. Place a copy of P into a large matrix M, remove x, and then move all 1-entries to the left of x further to the left. This creates a "middle part" of an expandable row (Figure 1.4 *h, i*). Call that modified matrix P_x. Arranging L, R, and P_x as in Figure 1.4 *j* completes an expandable row. In Section 3 we use this construction, except that P_x is additionally stretched in a certain way.

More generally, we can obtain an expandable row by interleaving copies of L, R and $P_{x_1}, P_{x_2}, \ldots$ for several 1-entries $x_1, x_2, \ldots$ of P. Figure 1.4 *k* shows another witness for the matrix in Figure 1.4 *e* constructed in this way. Observe that each partial copy of P takes care of a certain part of the expandable row. This idea forms the basis for the construction in Section 4, where we use a number of different partial copies P_{x_i} depending on the pattern P.

2. Spanning oscillations of length 4

In this section, we show Theorem 1.2, which immediately implies Lemma 1.15.

Theorem 1.2. *Let P be a pattern that contains four 1-entries x_1, x_2, x_3, x_4 such that for each $i \in [4]$, there are no other 1-entries in the same row or column as x_i, and x_i is in the first or last row or column, and x_1, x_2, x_3, x_4 form one of the two patterns*

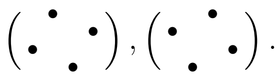

Then $\operatorname{sat}(P, n) \in \mathcal{O}(1)$.

Let $\mathcal{P}$ denote the class of patterns defined in Theorem 1.2. Note that $\mathcal{P}$ is closed under transposition. Thus, by Lemma 1.9, it is sufficient to prove that each $P \in \mathcal{P}$ has a vertical witness.

Let $\mathcal{P}'$ be the subset of patterns $P \in \mathcal{P}$ where the unique leftmost 1-entry ℓ of P is above the unique rightmost 1-entry r of P. It is easy to see that each $P \in \mathcal{P}'$ has the following form, where the boxes contain arbitrarily many 1-entries:

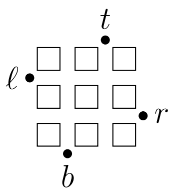

Since for each $P \in \mathcal{P} \setminus \mathcal{P}'$, we have $\operatorname{rev}(P) \in \mathcal{P}'$, Observation 1.8 implies that it is sufficient to prove that each $P \in \mathcal{P}'$ has a vertical witness.

Lemma 2.1. *Each $P \in \mathcal{P}'$ has a vertical witness.*

Proof. Let $P \in \mathcal{P}'$ be a $k_1 \times k_2$ pattern, let $\ell = (i, j)$ be the unique leftmost 1-entry in P, and let $r = (i', j')$ be the unique rightmost 1-entry in P. Note that $i < i'$.

We essentially use the construction shown in Figure 1.4 *d*. Let P_{L} and P_{R} be the submatrices of P obtained by removing the rightmost, resp. leftmost, column. Note that in P_{L}, the i'-th row is empty, and in P_{R}, the i-th row is empty. As described in Section 1.5, the idea is to place a copy of P_{L} to the left of P_{R}, so that the two empty rows coincide. More formally, obtain L from P_{L} by appending $i' - i > 0$ rows (at the bottom), obtain R from P_{R} by prepending $i' - i > 0$ rows (at the top), and define $S(P)$ as the horizontal concatenation (L, R). Note that $S(P)$ is a $(k_1 + i' - i) \times (2k_2 - 2)$ matrix, and that the i'-th row of $S(P)$ is empty. In the following, we use L and R interchangeably with the corresponding subsets of $E(S(P))$.

We claim that the i'-th row is P-expandable. Indeed, adding a 1-entry in the i'-th row in the first $k-1$ columns (to the left of R) completes an occurrence of P with R, and adding a 1-entry in the last $k-1$ columns (to the right of L) completes an occurrence of P with L.

It remains to show that $S(P)$ avoids P. Suppose $S(P)$ contains P, so there is an embedding ϕ of P into $S(P)$. Let $t, b \in E(P)$ be the unique topmost, respectively bottommost, 1-entry in P.

Suppose first that $\phi(b) \in L$. Since $\mathrm{height}(L) = \mathrm{d}^{\mathrm{v}}(t, b) = k - 1$, and the i'-th row of P is empty, we have $\mathrm{height}_\phi(L) < \mathrm{d}^{\mathrm{v}}(t, b)$. This implies that $\phi(t)$ is above L. But $S(P)$ has no 1-entries above L, a contradiction.

Otherwise, $\phi(b) \in R$. Since t is to the right of b, this implies that $\phi(t) \in R$. But a similar argument as above shows that $\mathrm{height}_\phi(R) < \mathrm{d}^{\mathrm{v}}(t, b)$, a contradiction. □

3. Spanning oscillations starting with t

In this section, we prove:

Lemma 3.1. *Each permutation matrix P with a wide spanning oscillation of length $m \geqslant 5$ that starts with t_P has a vertical witness.*

In Section 3.1, we present a construction of (possible) witnesses, which we use for the case $m = 5$ in Section 3.2, and for the case $m > 5$ in Section 3.3.

3.1. Witness construction

Let P be an $k \times k$ permutation matrix such that $\ell = \ell_P$ is above $r = r_P$, and let $q = (i_q, j_q) \in E(P)$, such that q is above ℓ. We first construct a matrix $S'(P, q)$ with a P-expandable row, in the way shown in Figure 1.4 *j*. Then, we modify $S'(P, q)$ to obtain the matrix $S(P, q)$, which retains the expandable row and will be shown to avoid P if P has a wide spanning oscillation $(t_P, \ell_P, x_3, x_4, \ldots, x_m)$ with $m \geqslant 5$ and we choose $q = x_3$.

Let P_{R} (P_{L}) be the submatrix of P obtained by removing the leftmost (rightmost) column. Both P_{R} and P_{L} have an empty row. To start the construction of $S'(P, q)$, we place a copy of P_{R} to the *left* of a copy of P_{L}, such that the two copies do not intersect, and the empty rows are aligned. We denote the copy of P_{R} in the construction with R and the copy of P_{L} with L. Note that, compared to the construction in Section 2, L and R switch places.

Let P'_{L} consist of all columns of P to the left of q, and let P'_{R} consist of all columns of P to the right of q. To finish the construction of $S'(P, q)$, we place a copy of P'_{L} to the left of R and a copy of P'_{R} to the right of L, such that the empty i_q-th rows of P'_{L} and P'_{R} are aligned with the empty row in R and L. Denote the copies of P'_{L} and P'_{R} as L' and R' and let $P' = L' \cup R'$.

Clearly, the empty row in $S'(P, q)$ is expandable: Adding a 1-entry to the left of R will complete the partial occurrence R of P, adding a 1-entry to the right of L will complete L, and adding a 1-entry within R or L will complete P'.

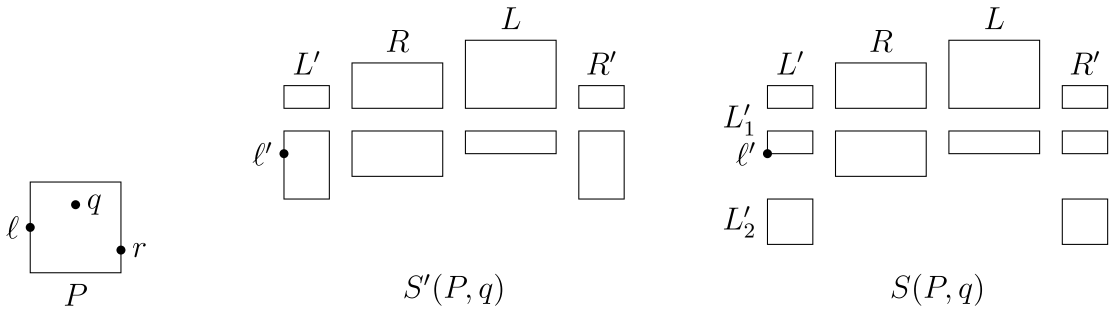

Figure 3.1: A sketch of P and the two witness constructions $S'(P,q)$ and $S(P,q)$.

We modify $S'(P,q)$ to obtain $S(P,q)$ as follows.[8] Let B be the set of entries in $P' = L' \cup R'$ that are below the leftmost 1-entry in P' (the copy of ℓ in P'). Move B down by a fixed number of rows, such that each 1-entry in B is lower than all 1-entries in $R \cup L$. Clearly, the expandable row stays expandable after this change.

Figure 3.1 sketches the constructions. In the following sections, we denote the 1-entries in $S(P,q)$ as follows. If x is a 1-entry in P, then let x^{R} be the copy of x in R, let x^{L} be the copy of x in L, and let x' be the copy of x in P'. For subsets $X \subseteq E(P)$, we use X^{R}, X^{L} and X' similarly.

We now show a property of $S(P,q)$ that is useful in both of the following subsections.

Lemma 3.2. *Let P be a $k \times k$ permutation matrix and $q \in E(P)$ such that $q <_{\mathrm{v}} \ell_P <_{\mathrm{v}} r$ and t_P is to the left of b_P. If ϕ is an embedding of P into $S(P,q)$, then $\phi(t_P) \notin L'$ and $\phi(b_P) \in R'$.*

Proof. We write ℓ, t, b, r for ℓ_P, t_P, b_P, r_P. Let L'_2 denote the portion of L' below ℓ', and let $L'_1 = L' \setminus L'_2$.

We first show that $\phi(t) \notin L'$. Suppose $\phi(t) \in L'$. Then also $\phi(\ell) \in L'$. Since $\mathrm{height}(L'_2) < \mathrm{d}^{\mathrm{v}}(\ell, b)$, and there are no nonempty rows below L'_2, we know that $\phi(\ell) \notin L'_2$, and therefore $\phi(t), \phi(\ell) \in L'_1$. But $\mathrm{height}_\phi(L'_1) \leqslant \mathrm{d}^{\mathrm{v}}(t,\ell) - 1$, a contradiction.

$\phi(t) \notin L'$ already shows that $\phi(b) \notin L'$, since b is to the right of t. It remains to show that $\phi(b) \notin R \cup L$. First, suppose that $\phi(b) \in L$. Then there are at most $k-2$ nonempty rows above $\phi(b)$, but $\mathrm{d}^{\mathrm{v}}(t,b) = k-1$, a contradiction.

Second, suppose that $\phi(b) \in R$. Then also $\phi(t) \in R$, because t is to the left of b and $\phi(t) \notin L'$. But $\mathrm{height}_\phi(R) \leqslant \mathrm{d}^{\mathrm{v}}(t,b) - 1$, a contradiction. □

3.2. Spanning oscillations of length five

Lemma 3.3. *Let P be a permutation matrix and $X = (t_P, x_2, x_3, x_4, x_5)$ be a a spanning oscillation of P. Then $S(P, x_3)$ avoids P.*

Proof. Let $q = x_3$, and note that $x_2 = \ell_P$ and $x_4 = b_P$, so q is above ℓ_P and to the right of b_P. Suppose ϕ is an embedding of P into $S(P,q)$. By Lemma 3.2, $\phi(b_P) \in R'$. But $\mathrm{width}(R') = \mathrm{d}^{\mathrm{h}}(q, r_P) - 1 < \mathrm{d}^{\mathrm{h}}(b_P, r_P)$, a contradiction. □

[8]This idea comes from Geneson's construction. [Gen21]

3.3. Longer spanning oscillations

We now consider the case where P has a wide spanning oscillation $(t_P, x_2, \ldots, x_m)$ of length greater than five. We first prove a useful property of long spanning oscillations.

Lemma 3.4. *Let P be a permutation matrix and $X = (t_P, x_2, \ldots, x_m)$ be a spanning oscillation of P with $m \geqslant 6$. Then, removing $t = t_P$, the columns to the left of t, and the rows above x_3 (as well as all newly created rows or columns) does not make P decomposable.*

Proof. Suppose it does, and let P_0 be the resulting decomposable pattern. Since x_3 is the highest 1-entry in P_0 (slightly abusing notation), and x_3 is above $r = r_P$ and to the left of $b = b_P$, we know that P_0 has the form $\left(\begin{smallmatrix} A & \mathbf{0} \\ \mathbf{0} & B \end{smallmatrix}\right)$, where x_3 lies in A and r, b lie in B. This means that x_4 lies in A, since $t <_{\mathrm{h}} x_4 <_{\mathrm{h}} x_3$. Let P_1 be the matrix obtained from P_0 by further removing all columns to the right of x_4. Clearly, P_1 is decomposable, but $(x_3, x_4, \ldots, x_m)$ is a spanning oscillation of P_1, a contradiction. □

We are now ready to prove the main result of this subsection.

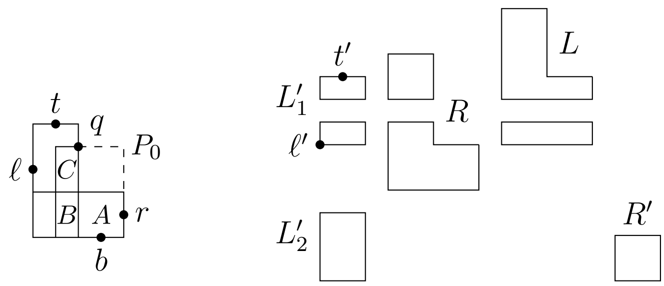

Figure 3.2: P and $S(P, q)$ in the case of Lemma 3.5.

Lemma 3.5. *Let $X = (t_P, x_2, \ldots, x_m)$ be a wide spanning oscillation of P with $m \geqslant 6$. Then P has a vertical witness.*

Proof. We write ℓ, t, b, r for ℓ_P, t_P, b_P, r_P in the following. Let $q = x_3$, and let P_0 be the set of 1-entries of P that are to the right of t and not above q. By Lemma 3.4, P_0 does not correspond to a decomposable pattern. Let A denote the set of 1-entries to the right of q. Note that $b, r \in A$, and, by wideness of X, all 1-entries in A are below ℓ. Let x be the highest 1-entry in A, and let B be the set of 1-entries below x, to the left of q and to the right of t. Then $B \neq \varnothing$, otherwise P_0 would be decomposable. Finally, $C = P_0 \setminus (A \cup B)$ consists of the 1-entries to the right of t, not above q, and above x. Figure 3.2 shows a sketch of P and $S(P, X)$. Note that $A' = R'$ (recall that A' denotes the copy of A in P').

Suppose ϕ is an embedding of P into $S(P, q)$. By Lemma 3.2, $\phi(b) \in R'$ and $\phi(t) \notin L'$. Since all 1-entries in B are to the right of t, this implies $\phi(y) \notin L'$ for each $y \in B$. Moreover, $\mathrm{width}(R') = \mathrm{d}^{\mathrm{h}}(q, r) - 1 < \mathrm{d}^{\mathrm{h}}(y, r)$ for each $y \in B$, so we have $\phi(B) \subseteq L \cup R$.

Let L_2' denote the portion of L' below ℓ' and let $L_1' = L' \setminus L_2'$. Note that L_2' is below all 1-entries in $L \cup R$. Since all 1-entries in C are above all 1-entries in B, and all 1-entries in A are to the right of all 1-entries in B, we have $\phi(P_0) = \phi(A \cup B \cup C) \subseteq L_1' \cup L \cup R \cup R'$. Since $R' = A'$,

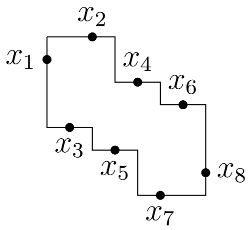

Figure 4.1: A tall traversal. The solid lines indicate the boundary of possible positions for other 1-entries.

all 1-entries in R' are to the right and below all 1-entries in $L'_1 \cup L \cup R$, so $L'_1 \cup L \cup R \cup R'$ can be decomposed into the two blocks $L'_1 \cup L \cup R$ and R'. Further, $\phi(b) \in R'$ by Lemma 3.2, and since $\mathrm{height}(R') < \mathrm{d}^{\mathrm{v}}(q, b)$, we have $\phi(q) \notin R'$. This means that P_0 is decomposable, a contradiction. □

4. Even-length spanning oscillations starting with ℓ

In this section, we prove:

Lemma 4.1. *Each permutation matrix P with a tall spanning oscillation of even length $m \geqslant 6$ that starts with ℓ_P has a vertical witness.*

For our witness construction to work, we need to define a substructure that generalizes (tall) spanning oscillations of even length that start with ℓ_P. We call that substructure a *traversal*. Defining our witness construction for traversals instead of spanning oscillations allows us to make a maximality assumption that is required later in the proof.

4.1. Traversals

Let P be a permutation matrix and let $m \geqslant 4$. A *traversal* of P is a sequence X of distinct 1-entries $x_1, x_2, \ldots, x_m$ such that

(i) $x_1 = \ell_P$, $x_2 = t_P$, $x_{m-1} = b_P$, $x_m = r_P$;

(ii) $x_1 <_{\mathrm{h}} x_3 <_{\mathrm{h}} x_2 <_{\mathrm{h}} x_5 <_{\mathrm{h}} x_4 <_{\mathrm{h}} \ldots <_{\mathrm{h}} x_{m-1} <_{\mathrm{h}} x_{m-2} <_{\mathrm{h}} x_m$;

(iii) $\ell_P <_{\mathrm{v}} x_4 <_{\mathrm{v}} x_6 <_{\mathrm{v}} \ldots <_{\mathrm{v}} x_m$;

(iv) $x_3 <_{\mathrm{v}} x_5 <_{\mathrm{v}} \ldots <_{\mathrm{v}} x_{m-3} <_{\mathrm{v}} r_P$; and

(v) x_s is below x_{s+1} for each odd $s \in [m-1]$.

Intuitively, property (ii) enforces the same horizontal order on the 1-entries as an even-length spanning oscillation. Vertically, however, we are allowed to arrange the 1-entries more freely. There are still *upper* (even) and *lower* (odd) 1-entries as in Observation 1.12 (this is implied by (iii), (iv), (v)), and we keep the order within the upper, resp. lower, 1-entries with (iii), (iv). But

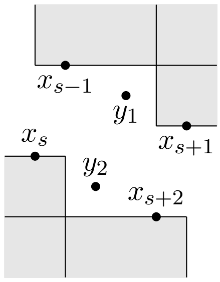

Figure 4.2: Arrangement of $x_{s-1}, x_s, y_1, y_2, x_{s+1}, x_{s+2}$ in Lemma 4.3. The shaded areas must be empty, since X is tall.

we drop the condition that x_i is above x_j for each odd $i \leqslant m-3$ and even $j \geqslant i+3$. This means that we are allowed to "move" some upper 1-entries upwards, and some lower 1-entries downwards, as long as the vertical order among upper (lower) 1-entries is kept intact. (iii), (iv) additionally ensure that we cannot move any 1-entries above ℓ_P or below r_P. Figure 4.1 shows the shortest traversal that is not an oscillation.

We say a traversal $(x_1, x_2, \dots, x_m)$ is *tall* if it satisfies the following two properties for each even $2 \leqslant i \leqslant m-2$.

(vi) P has no 1-entry that is below x_{i+1} and to the left of x_i.

(vii) P has no 1-entry that is above x_i and to the right of x_{i+1}.

Observation 4.2. *Each tall spanning oscillation of even length that starts with ℓ is a tall traversal.* □

4.2. Maximality assumption

Let P be a permutation matrix with a tall traversal X. We can assume that X is maximal in the sense that no tall traversal of P has X as a proper subsequence. We now show that such a *maximally tall* traversal also cannot be extended to a larger non-tall traversal in the following sense. Call a traversal $(x_1, x_2, \dots, x_m)$ *extendable* if there is an odd s with $5 \leqslant s \leqslant m-5$, and two 1-entries y_1, y_2 in P such that $(x_1, x_2, \dots, x_s, y_1, y_2, x_{s+1}, \dots, x_m)$ is a traversal of P.

Lemma 4.3. *Let $X = (x_1, x_2, \dots, x_m)$ be a maximally tall traversal of the permutation matrix P. Then X is non-extendable.*

Proof. Suppose X is extendable. Then there exists an odd s with $5 \leqslant s \leqslant m-5$ and 1-entries $y_1, y_2 \in E(P)$ such that $Y = (x_1, x_2, \dots, x_s, y_1, y_2, x_{s+1}, \dots, x_m)$ is a traversal of P. We show that then P has a tall traversal of length $m+2$ with X as a subsequence. This contradicts our assumption that X is maximally tall.

Note that property (v) of X implies that x_{s+1} is above x_s. Further using properties (ii), (iii), (iv) of Y, it follows that the relative positions of x_{s-1}, x_s, y_1, y_2, x_{s+1}, and x_{s+2} are fixed as shown in Figure 4.2.

Let y_1' and y_2' be 1-entries in P such that

(a) y_2' is to the left of y_1';

(b) y_1' is above or equal to y_1 and y_2' is below or equal to y_2; and

(c) $\mathrm{d}^{\mathrm{v}}(y_1', y_2')$ is maximal under the previous two conditions.

Let $Y' = (x_1, x_2, \dots, x_s, y_1', y_2', x_{s+1}, \dots, x_m)$. We first show that Y' is a traversal. Y' clearly satisfies (i). Since y_1' is not below y_1, it is above x_{s+1}, so tallness of X implies that y_1' is to the left of x_{s+2}. Symmetrically, y_2' is to the right of x_{s-1}, so (a) implies $x_{s-1} <_{\mathrm{h}} y_2' <_{\mathrm{h}} y_1' <_{\mathrm{h}} x_{s+2}$, and thus Y' satisfies (ii).

Since y_1' is to the right of x_{s-1}, tallness of X implies that y_1' is below x_{s-1}. We already observed that y_1' is above x_{s+1}, so we have $x_{s-1} <_{\mathrm{v}} y_1' <_{\mathrm{v}} x_{s+1}$. Similarly, we have $x_s <_{\mathrm{v}} y_s' <_{\mathrm{v}} x_{s+2}$. Together with $x_{s+1} <_{\mathrm{v}} x_s$, this implies the remaining traversal properties (iii), (iv), (v).

It remains to show that Y' is tall. Suppose Y violates tallness property (vi). Since X is tall, the only way this can happen is if there is a 1-entry z below y_2' and to the left of y_1'. Then z is also below y_2, but $\mathrm{d}^{\mathrm{v}}(y_1', z) > \mathrm{d}^{\mathrm{v}}(y_1', y_2')$, violating our assumption (c). A symmetric argument shows that Y satisfies (vii). □

4.3. Construction

Fix a $k \times k$ permutation matrix P. Throughout this subsection, we write ℓ, b, t, r for ℓ_P, b_P, t_P, r_P. For a 1-entry $x = (i, j) \in E(P)$, denote by P_x^{L} the submatrix of P consisting of all columns to the left of x (i.e., the leftmost $j-1$ columns), and denote by P_x^{R} the submatrix of P consisting of all columns to the right of x (i.e., the rightmost $k-j$ columns). Note that in both P_x^{L} and P_x^{R}, the i-th row is empty. Also note that the constructions in Sections 2 and 3 implicitly used $P_x^{\mathrm{L}}, P_x^{\mathrm{R}}$, with $x \in \{\ell, r, q\}$.

We first construct a matrix with an expandable row. For this, we take P_x^{L} and P_x^{R} for (almost) every 1-entry x in a traversal of P and arrange them like in Figure 1.4 k. Then, we vertically move parts of to arrive at our final construction. A formal description follows.

Let $X = (x_1, x_2, \dots, x_m)$ be a traversal of P with $m \geqslant 6$, and write $(i_s, j_s) = x_s$ for $s \in [m]$. Then the $(2k-1) \times (m-2)k$ matrix $S'(P, X)$ is constructed as follows. Let L_s' be the $(2k-1) \times (j_s - 1)$ matrix consisting of a copy of $P_{x_s}^{\mathrm{L}}$ that is shifted down by $k - i_s$ rows (i.e., we prepend $k - i_s$ rows and append $i_s - 1$ rows to $P_{x_s}^{\mathrm{L}}$). Similarly, let R_s' be the $(2k-1) \times (k - j_s)$ matrix consisting of a copy of $P_{x_s}^{\mathrm{R}}$ that is shifted down by $k - i_s$ rows. Note that the empty i_s-th row of $P_{x_s}^{\mathrm{L}}$ ($P_{x_s}^{\mathrm{R}}$) corresponds to the k-th row of L_s' (R_s'). Finally, let $S'(P, X)$ be the following horizontal concatenation[9] of matrices:

$$S'(P, X) = (L_3', R_1', L_4', R_3', L_5', R_4', \dots, L_{m-3}', R_{m-4}', L_{m-2}', R_{m-3}', L_m', R_{m-2}').$$

Note the irregularities at the beginning and the end. Notably, $L_2', R_2', L_{m-1}', R_{m-1}'$ are not used in the construction. L_1' and R_m' are not used, either, but they are empty anyway, since $x_1 = \ell$ and $x_m = r$. See Figure 4.3 for an example.

[9]See Page 4 for the definition of horizontal concatenations.

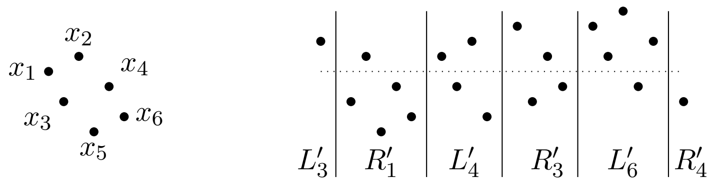

Figure 4.3: A matrix P consisting of a 6-traversal X, and the corresponding construction $S'(P, X)$. Some empty columns in $S'(P, X)$ have been omitted. The dotted line indicates the expandable row.

We claim that the k-th row of $S'(P, X)$ is expandable. Indeed, for each i with $3 \leqslant i \leqslant m-2$, adding a 1-entry in the k-th row between L'_i and R'_i will complete a copy of P with L'_i and R'_i. Moreover, adding a 1-entry in the k-th row to the left of R'_1 or to the right of L'_m will complete a copy of P. By construction, this covers the whole k-th row.

As in the Section 3.1, we will not directly use $S'(P, X)$, but rather a modified construction that preserves the expandable row. We do this to avoid that two different parts of the construction (i.e., $L'_i \cup R'_i$ for $i \in [m]$) overlap vertically. This reduces the number of ways P could appear in the constructed matrix, and thus makes the analysis much easier.

In the following, we will slightly abuse the notation by writing L'_s (R'_s) for the subsets of $E(S'(P, X))$ that correspond to L'_s (R'_s).

Let $S(P, X)$ be a $((2m-6)k+1) \times (m-2)k$ matrix, constructed as follows. Start with a copy of $S'(P, X)$, shifted down by $(m-4)k$ rows, such that the expandable k-th row of $S'(P, X)$ corresponds to the $(m-3)k$-th row of $S(P, X)$. Now, for each $s \in \{5, 6, \dots, m-1, m-2, m\}$, take all 1-entries in $L'_s \cup R'_s$ that are above the $((m-3)k-1)$-th row (i.e., at least two rows above the expandable row), and move them up by $(s-4)k$ rows. Similarly, for each $s \in \{1, 3, 4, \dots, m-4\}$, take all 1-entries in $L'_s \cup R'_s$ that are below the $((m-3)k+1)$-th row (i.e., at least two rows below the expandable row), and move them down by $(m-s-3)k$ rows. Figure 4.4 shows the rough structure of $S(P, X)$ when $m = 12$ and X is tall.

Let L_s (R_s) denote the the modified set of entries in $S(P, X)$ corresponding to L'_s (R'_s). Clearly, L_s and R_s still form a partial occurrence of P with a single 1-entry missing between them in the $(m-3)k$-th row. Similarly, R_1 and L_m form occurrences when adding a 1-entry in the left- or rightmost part of that row. Thus:

Lemma 4.4. *If X is a traversal of P, then $S(P, X)$ has an expandable row.*

Note that the construction used in Section 2 can be seen as a special case of both $S(P, X)$ and $S'(P, X)$ when $m = 4$.

The remainder of this paper is dedicated to the proof that if X is a non-extendable tall traversal of a permutation matrix P, then $S(P, X)$ avoids P, implying that $S(P, X)$ is a vertical witness of P. We first fix some notations and make a few observations about $S(P, X)$. Let T denote the set of 1-entries that are above row $(m-3)k-1$ (at least two rows above the expandable row). Similarly, let B denote the set of 1-entries that are below row $(m-3)k+1$, and let M denote the remaining 1-entries. For a subset $A \subseteq E(S(P, X))$, let $A^{\mathrm{T}} = A \cap T$, let $A^{\mathrm{B}} = A \cap B$ and let $A^{\mathrm{M}} = A \cap M$. For a 1-entry $p \neq x_s$, let p^s denote the copy of p in $L_s \cup R_s$.

Observation 4.5. *Let $s, u \in \{1, 3, 4, \dots, m-3, m-2, m\}$ with $s < u$. If $u \geqslant 5$, then every 1-entry in $L_s^{\mathrm{T}} \cup R_s^{\mathrm{T}}$ is below every 1-entry in $L_u^{\mathrm{T}} \cup R_u^{\mathrm{T}}$. Moreover, if $s \leqslant m-4$, then every 1-entry in $L_s^{\mathrm{B}} \cup R_s^{\mathrm{B}}$ is below every 1-entry in $L_u^{\mathrm{B}} \cup R_u^{\mathrm{B}}$.* □

Since X is tall, there are no 1-entries below and to the left of x_s if s is odd, or above and to the right of x_s if s is even. This implies:

Observation 4.6. *Let s be odd with $3 \leqslant s \leqslant m-3$. Then L_s contains no 1-entries below the expandable row, and R_{s+1} contains no 1-entries above the expandable row. In particular, $L_s^{\mathrm{B}} = \varnothing$ and $R_{s+1}^{\mathrm{T}} = \varnothing$.* □

We now consider the width and height of relevant parts of $S(P, X)$.

Observation 4.7. *For each $s \in \{1, 3, 4, \dots, m-3, m-2, m\}$,*

- $\mathrm{width}(L_s) = \mathrm{d}^{\mathrm{h}}(\ell, x_s) - 1$*;*
- $\mathrm{width}(R_s) = \mathrm{d}^{\mathrm{h}}(x_s, r) - 1$*;*
- $\mathrm{height}(L_s^{\mathrm{T}} \cup R_s^{\mathrm{T}}) = \mathrm{d}^{\mathrm{v}}(t, x_s) - 2$*, if $L_s^{\mathrm{T}} \cup R_s^{\mathrm{T}} \neq \varnothing$;*
- $\mathrm{height}_\phi(L_s^{\mathrm{M}} \cup R_s^{\mathrm{M}}) \leqslant 1$*; and*
- $\mathrm{height}(L_s^{\mathrm{B}} \cup R_s^{\mathrm{B}}) = \mathrm{d}^{\mathrm{v}}(x_s, b) - 2$*, if $L_s^{\mathrm{B}} \cup R_s^{\mathrm{B}} \neq \varnothing$.* □

Let $3 \leqslant s \leqslant m-3$ be odd. Since X is tall, there are no 1-entries in P above x_{s-1} and to the right of x_s. Thus, x_{s-1}^s is the topmost 1-entry in R_s. Similarly, x_{s+2}^{s+1} is the bottommost 1-entry in L_{s+1}. This implies the following improved bounds:

Observation 4.8. *For each odd $s \in \{3, 4, \dots, m-2\}$:*

- $\mathrm{height}(R_s^{\mathrm{T}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s-1}, x_s) - 2$*, if $R_s^{\mathrm{T}} \neq \varnothing$; and*
- $\mathrm{height}(L_{s+1}^{\mathrm{B}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, x_{s+2}) - 2$*, if $L_{s+1}^{\mathrm{B}} \neq \varnothing$.* □

4.4. $S(P, X)$ avoids P

In this section, we show:

Lemma 4.9. *Let P be a permutation matrix, $m \geqslant 6$ be even and let $X = (x_1, x_2, \dots, x_m)$ be a non-extendable tall traversal of P. Then $S(P, X)$ avoids P.*

Together with Observation 4.2 and Lemmas 4.3 and 4.4, this implies Lemma 1.17. For the remainder of this section, fix P and X as in Lemma 4.9, and write ℓ, b, t, r for ℓ_P, b_P, t_P, r_P. We use the notation for parts of $S(P, X)$ as defined in Section 4.3. Suppose ϕ is an embedding of P into $S(P, X)$. Our overall strategy is to distinguish cases based on the location of $\phi(t)$, and derive a contradiction in each case. While the full proof is long and technical, it only uses a handful of simple arguments that are combined and applied to various situations.

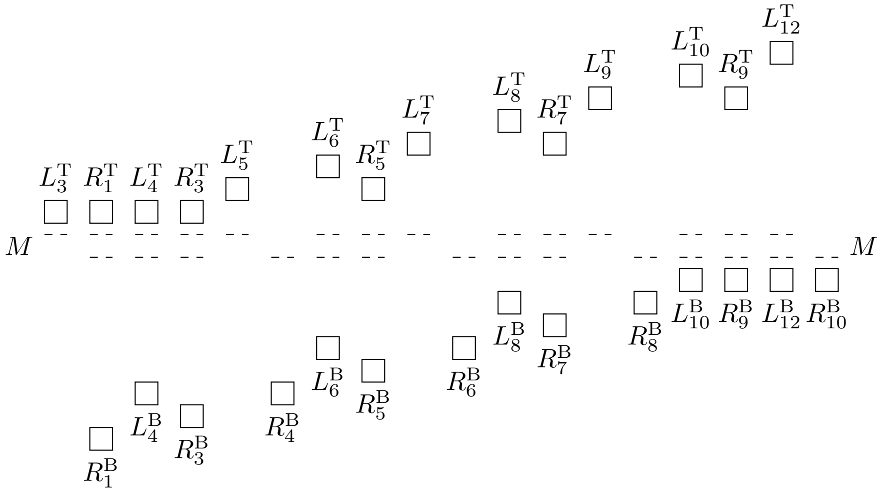

Figure 4.4: A sketch of the block structure of $S(P, X)$ with $|X| = 12$.

Note that we make no further assumptions on P, X, ϕ, so each lemma or corollary in this section holds on its own for every choice of P, X, ϕ (we only fix P, X, ϕ for brevity). This allows us to make use of the following symmetry argument. $S(P, X)$ is not usually symmetric, in the sense that its 180-degree rotation $\mathrm{rot}^2(S(P, X))$ is equal to $S(P, X)$. However, it is easy to see that $\mathrm{rot}^2(S(P, X))$ is equal to $S(\mathrm{rot}^2(P), \mathrm{rot}^2(X))$. Now, in Lemma 4.10, for example, we show that $\phi(t) \notin L_3$ for each choice of P, X, ϕ, in particular also for every embedding ϕ' of $\mathrm{rot}^2(P)$ into $\mathrm{rot}^2(S(P, X))$.

We also get $\phi(b) \notin R_{m-2}$, since R_{m-2} in $S(P, X)$ corresponds to L_3 in $\mathrm{rot}^2(S(P, X)) = S(\mathrm{rot}^2(P), \mathrm{rot}^2(X))$, b in P corresponds to t in $\mathrm{rot}^2(P)$, and ϕ corresponds to some embedding ϕ' of $\mathrm{rot}^2(P)$ into $\mathrm{rot}^2(S(P, X))$.

4.4.1 $\phi(t)$ in the front or the back

In this section, we show $\phi(t)$ and $\phi(b)$ cannot lie in the leftmost or rightmost few "blocks" of $S(P, X)$. The precise results are Corollary 4.18 and Lemma 4.19 at the end of the section. The proofs in this section also serve as a warm-up for the more complex later proofs. Most techniques used in Sections 4.4.2 and 4.4.3 already appear here, where we explain them thoroughly.

Lemma 4.10. *$\phi(t) \notin L_3$ and $\phi(b) \notin R_{m-2}$.*

Proof. By symmetry, it suffices to show $\phi(t) \notin L_3$. Suppose $\phi(t) \in L_3$. Then also $\phi(\ell) \in L_3$, since $S(P, X)$ contains no 1-entries to the left of L_3. But $\mathrm{width}(L_3) = \mathrm{d}^{\mathrm{h}}(\ell, x_3) - 1 < \mathrm{d}^{\mathrm{h}}(\ell, t) - 1$, thus we cannot have both $\phi(\ell)$ and $\phi(t)$ in L_3, a contradiction. □

Lemma 4.11. *$\phi(t) \notin R_1$ and $\phi(b) \notin L_m$.*

Proof. By symmetry, it suffices to show $\phi(t) \notin R_1$. Suppose $\phi(t) \in R_1$. Note that $\mathrm{height}(R_1^{\mathrm{T}} \cup M) \leqslant \mathrm{d}^{\mathrm{v}}(t, \ell) + 1$. Since this bound counts the (empty) expandable row,

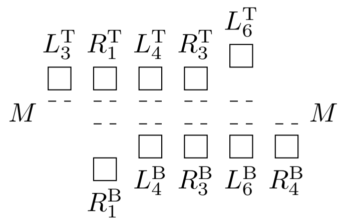

Figure 4.5: A sketch of the block structure of $S(P,X)$ with $|X|=6$.

we have $\mathrm{height}_\phi(R_1^{\mathrm{T}} \cup M) \leqslant \mathrm{d}^{\mathrm{v}}(t,\ell)$. This implies that $\phi(\ell)$ must lie in the lowest row of M or lower, and thus $\phi(\ell)$ is below the expandable row.

Since x_3 is below ℓ, this also implies that x_3 at least two rows below the expandable row, so $\phi(x_3) \in B$. Further, x_3 is to the right of t and $L_3^{\mathrm{B}} = \varnothing$, so we have $\phi(x_3) \in R_1^{\mathrm{B}}$. As r is below x_3, and all 1-entries in $S(P,X)$ that are to the right of R_1 are above R_1^{B}, we have $\phi(r) \in R_1^{\mathrm{B}}$. Since $\mathrm{width}(R_1) < \mathrm{d}^{\mathrm{h}}(\ell,r)$, this implies that $\phi(\ell)$ is to the left of R_1. We now know that $\phi(\ell)$ is below the expandable row and to the left of R_1^{T}. But $S(P,X)$ has no such 1-entry, a contradiction. □

If $m=6$ (see Figure 4.5), then the only remaining possibility is $\phi(t),\phi(b) \in L_4 \cup R_3$, which implies $\phi(t) \in L_4$ or $\phi(b) \in R_3$ (since t is to the left of b). Thus, the following lemma concludes the case $m=6$.

Lemma 4.12. *If $m=6$, then $\phi(t) \notin L_4$ and $\phi(b) \notin R_3$.*

Proof. By symmetry, it suffices to show $\phi(t) \notin L_4$. This can be done with essentially the same argument as in the proof of Lemma 2.1. Suppose $\phi(t) \in L_4$. Then $\phi(t)$ is not above $t^4 \in L_4$. By Lemmas 4.10 and 4.11, $\phi(b) \in L_4 \cup R_3$. The lowest 1-entry in $L_4 \cup R_3$ is b^4, so $\phi(b)$ is not below b^4. But $\mathrm{d}_\phi^{\mathrm{v}}(t^4,b^4) < \mathrm{d}^{\mathrm{v}}(t,b)$ (note the empty expandable row), a contradiction. □

We now continue with the case $m \geqslant 8$.

Lemma 4.13. *If $m \geqslant 8$, then $\phi(t) \notin L_4$ and $\phi(b) \notin R_{m-3}$.*

Proof. By symmetry, it suffices to show $\phi(t) \notin L_4$. Suppose $\phi(t) \in L_4$. We have $\mathrm{height}_\phi(L_4^{\mathrm{T}} \cup M) \leqslant \mathrm{d}^{\mathrm{v}}(t,x_4) < \mathrm{d}^{\mathrm{v}}(t,x_3)$, implying that $\phi(x_3) \in B$. More precisely, we have $\phi(x_3) \in R_1^{\mathrm{B}} \cup L_4^{\mathrm{B}}$, because x_3 is to the left of t.

Since r is below x_3 and to the left of t, we have $\phi(r) \in L_4^{\mathrm{B}} \cup R_3^{\mathrm{B}} \cup R_4^{\mathrm{B}}$. This means that ϕ maps no 1-entry to the right of R_4, and thus maps no 1-entry into the rows between M and $L_4^{\mathrm{B}} \cup R_4^{\mathrm{B}}$. This is a very useful observation, since it essentially allows us to pretend that M is directly above $L_4^{\mathrm{B}} \cup R_4^{\mathrm{B}}$. Similar observations will be used frequently in subsequent proofs.

From the above, we get $\mathrm{height}_\phi(L_4 \cup R_4) < \mathrm{d}^{\mathrm{v}}(t,b)$ (note that L_4^{T} is directly above M), so $\phi(b)$ is below $L_4 \cup R_4$, and thus $\phi(b) \in R_3^{\mathrm{B}}$. Moreover, by tallness of X, we have $\mathrm{height}_\phi(L_4) < \mathrm{d}^{\mathrm{v}}(t,x_5)$, so $\phi(x_5)$ is below L_4. Since x_5 is to the left of b, this means that $\phi(x_5) \in R_3^{\mathrm{B}}$.

Consider now $\phi(x_4)$. Since $x_5 <_{\mathrm{h}} x_4 <_{\mathrm{h}} b$, we have $\phi(x_4) \in R_3$. Since $\mathrm{height}(R_3^{\mathrm{B}}) < \mathrm{d}^{\mathrm{v}}(x_3,b) < \mathrm{d}^{\mathrm{v}}(x_4,b)$, we have $\phi(x_4) \in R_3^{\mathrm{T}} \cup R_3^{\mathrm{M}}$.

We conclude the proof with a case distinction. First, assume that $\phi(t) \neq t^4$. Since $\phi(t) \in L_4$ and t^4 is the highest 1-entry in L_4, this means that $\phi(t)$ is below t^4, and thus $\mathrm{d}^{\mathrm{v}}_{\phi}(\phi(t), \phi(x_4)) < \mathrm{d}^{\mathrm{v}}_{\phi}(t^4, \phi(x_4)) \leqslant \mathrm{height}_{\phi}(L_4 \cup M) \leqslant \mathrm{d}^{\mathrm{v}}(t, x_4)$, a contradiction.

Second, assume that $\phi(t) = t_4$. Recall that ϕ maps no 1-entries between M and L_4^{B}. Because of this and the fact that the expandable row is empty, we have that $\mathrm{d}^{\mathrm{v}}_{\phi}(t^4, x_3^4) < \mathrm{d}^{\mathrm{v}}(t, x_3)$, implying that $\phi(x_3)$ is below x_3^4. By tallness of X, this also implies that $\phi(x_3)$ is to the right of x_3^4. However, since x_3^4 is to the left of t^4, this means that $\mathrm{d}^{\mathrm{h}}(\phi(x_3), \phi(t)) < \mathrm{d}^{\mathrm{h}}(x_3^4, t^4) = \mathrm{d}^{\mathrm{h}}(x_3, t)$, a contradiction. □

Lemma 4.14. *Let $m \geqslant 8$. If $\phi(t) \in R_3$, then $\phi(b)$ is to the right of R_4. Moreover, if $\phi(b) \in L_{m-2}$, then $\phi(t)$ is to the left of L_{m-3}.*

Proof. By symmetry, proving the first statement suffices. Let $\phi(t) \in R_3$ and suppose $\phi(b)$ not to the right of R_4. Since $\phi(R_3^{\mathrm{T}}) \leqslant \mathrm{d}^{\mathrm{v}}(t, x_3)$, we know that $\phi(x_3)$ is below the expandable row. Let q_3 be the 1-entry directly below x_3 in P. Clearly, $\phi(q_3), \phi(b), \phi(r) \in B$, and since $\phi(b)$ is to the right of $\phi(t)$ and not to the right of R_4, we have $\phi(b) \in R_3^{\mathrm{B}} \cup R_4^{\mathrm{B}}$. We separately consider three cases.

Case 1: $\phi(r) \in R_3^{\mathrm{B}}$. Since X is tall, q_3 is to the right of t, so $\phi(q_3) \in R_3^{\mathrm{B}}$. But $\mathrm{height}(R_3^{\mathrm{B}}) = \mathrm{d}^{\mathrm{v}}(x_3, b) - 2 = \mathrm{d}^{\mathrm{v}}(q_3, b) - 1$, a contradiction.

Case 2: $\phi(r) \in R_4^{\mathrm{B}}$. Consider x_5. Since x_5 is below x_3, we have $\phi(x_5) \in B$. Since x_5 is to the right of t, and above and to the left of r, we have $\phi(x_5) \in R_4^{\mathrm{B}}$. But $\mathrm{width}(R_4) = \mathrm{d}^{\mathrm{h}}(x_4, r) - 1 < \mathrm{d}^{\mathrm{h}}(x_5, r)$, a contradiction.

Case 3: $\phi(r)$ is to the right of R_4. Then $\phi(r)$ is also above $L_4^{\mathrm{B}} \cup R_4^{\mathrm{B}}$. Consider again x_5. We know that $\phi(x_5)$ is below M and above $L_4^{\mathrm{B}} \cup R_4^{\mathrm{B}}$. Since x_5 is to the left of b, we also know that $\phi(x_5)$ is not to the right of R_4. But there are no such 1-entries in $S(P, X)$, a contradiction. □

We proceed with some more special cases, showing that $\phi(t)$ also cannot lie in the rightmost few blocks of $S(P, X)$.

Lemma 4.15. *Let $m \geqslant 8$. Then, $\phi(t)$ lies to the left of L_{m-2}, and $\phi(b)$ lies to the right of R_3.*

Proof. By symmetry, it suffices to prove that $\phi(t)$ lies to the left of L_{m-2}. If $\phi(b)$ lies to the left of L_{m-2}, then $\phi(t)$ does, too. $\phi(b) \notin R_{m-3} \cup L_m \cup R_{m-2}$ by Lemmas 4.10, 4.11 and 4.13. The only remaining possibility is that $\phi(b) \in L_{m-2}$, where Lemma 4.14 implies that $\phi(t)$ lies to the left of L_{m-3}, and thus to the left of L_{m-2}. □

To show that $\phi(t) \notin L_{m-3} \cup R_{m-4}$, we use the following more general lemma. Figure 4.6 is useful to visualize the proof.

Lemma 4.16. *Let s be odd with $5 \leqslant s \leqslant m-3$. If $\phi(t) \in L_s \cup R_{s-1}$, then $\phi(b)$ lies to the right of R_{s-1}.*

Proof. Suppose not. Then, $\phi(t), \phi(b) \in L_s \cup R_{s-1}$.

Case 1: $\phi(\ell) \notin L_s \cup R_{s-1}$. Since ℓ is to the left of t, this means that $\phi(\ell)$ is to the left of L_s. This implies that $\phi(\ell)$ is also below L_s^{T}, and thus $\phi(x_4)$ is below L_s^{T}. Since x_4 is to the right of t, we have $\phi(x_4) \in M \cup B$, which implies $\phi(x_5) \in B$, as $\mathrm{height}_\phi(M) \leqslant 1 < \mathrm{d^v}(x_4, x_5)$. Since x_5 is to the right of t and to the left of b, we further know $\phi(x_5) \in R_{s-1}^{\mathrm{B}}$. Since $\mathrm{width}(R_{s-1}) < \mathrm{d^h}(x_5, r)$, this implies that $\phi(r)$ is to the right of R_{s-1}. But then $\phi(r)$ is above $\phi(x_5)$, a contradiction.

Case 2: $\phi(\ell) \in L_s \cup R_{s-1}$. Then ϕ maps no 1-entry to the left of L_s. Since P is indecomposable, there must be some $y, z \in E(P)$ such that $\phi(y) \in L_s$, and $\phi(z)$ is above $\phi(y)$ and to the right of L_s. Note that L_s contains no 1-entries below the expandable row (by tallness of X), so $\phi(z) \in T$. Further, $\phi(t) \in L_s$ implies that $\phi(z) \in R_s^{\mathrm{T}}$. Since $\phi(b) \in L_s \cup R_{s-1}$, we know that b is to the left of z. Tallness of X implies that z is not above x_{m-2}. Now consider x_{s-1}. We know $x_{s-1} \leqslant_{\mathrm{h}} x_{m-4} <_{\mathrm{h}} b$ and $x_{s-1} \leqslant_{\mathrm{v}} x_{m-4} <_{\mathrm{v}} x_{m-2} \leqslant_{\mathrm{v}} z$. Thus, $\phi(x_{s-1}) \in L_s^{\mathrm{T}}$. But $\mathrm{width}(L_s^{\mathrm{T}}) < \mathrm{d^v}(\ell, x_s) < \mathrm{d^v}(\ell, x_{s-1})$, a contradiction. □

Corollary 4.17. *If $m \geqslant 8$, then $\phi(t) \notin L_{m-3} \cup R_{m-4}$ and $\phi(b) \notin L_5 \cup R_4$.*

Proof. By symmetry, it suffices to prove that $\phi(t) \notin L_{m-3} \cup R_{m-4}$. Suppose $\phi(t) \in L_{m-3} \cup R_{m-4}$. By Lemmas 4.10, 4.11, 4.13 and 4.14, $\phi(b)$ cannot lie in L_{m-2} or further right. This contradicts Lemma 4.16. □

We now consolidate and reformulate the above results. For the more involved proofs in Sections 4.4.2 and 4.4.3, it will be convenient to organize the "middle" blocks L_i, R_i of $S(P, X)$ into two sets of groups, as follows. For each odd s with $5 \leqslant s \leqslant m-5$, let $G_s = L_s \cup R_{s-1} \cup L_{s+1} \cup R_s$, and let $H_s = L_{s+1} \cup R_s \cup L_{s+2} \cup R_{s+1}$. Figure 4.6 illustrates G_s and H_s. Combining Lemmas 4.10, 4.11, 4.13 and 4.15 and Corollary 4.17 yields:

Corollary 4.18. *If $m \geqslant 8$, then:*

- *$\phi(t)$ lies to the right of L_4 and to the left of L_{m-3}. In other words, $\phi(t) \in R_3$ or $\phi(t) \in G_s$ for some odd s with $5 \leqslant s \leqslant m-5$; and*
- *$\phi(b)$ lies to the right of R_4 and to the left of R_{m-3}. In other words, $\phi(b) \in L_{m-2}$ or $\phi(b) \in H_s$ for some odd s with $5 \leqslant s \leqslant m-5$.*

At this stage, we cannot easily show that both $\phi(t) \notin R_3$ and $\phi(b) \notin L_{m-2}$, but we can show that at least one of the two must be true.

Lemma 4.19. *If $m \geqslant 8$, then $\phi(t) \notin R_3$ or $\phi(b) \notin L_{m-2}$*

Proof. Suppose $\phi(t) \in R_3$ and $\phi(b) \in L_{m-2}$. Since $\mathrm{height}_\phi(R_3^{\mathrm{T}} \cup M) \leqslant \mathrm{d^v}(t, x_3) < \mathrm{d^v}(t, x_5)$, we have $\phi(x_5) \in B$. More precisely, as $b \in L_{m-2}$, we have $\phi(x_5) \in L_{m-2}^{\mathrm{B}} \cup R_{m-3}^{\mathrm{B}} \cup L_m^{\mathrm{B}} \cup R_{m-2}^{\mathrm{B}}$. Similarly, $\phi(x_{m-4}) \in L_3^{\mathrm{T}} \cup R_1^{\mathrm{T}} \cup L_4^{\mathrm{T}} \cup R_3^{\mathrm{T}}$. In particular, $\phi(x_5)$ is to the right of $\phi(x_{m-4})$. But $x_5 <_{\mathrm{h}} x_4 \leqslant_{\mathrm{h}} x_{m-4}$, a contradiction. □

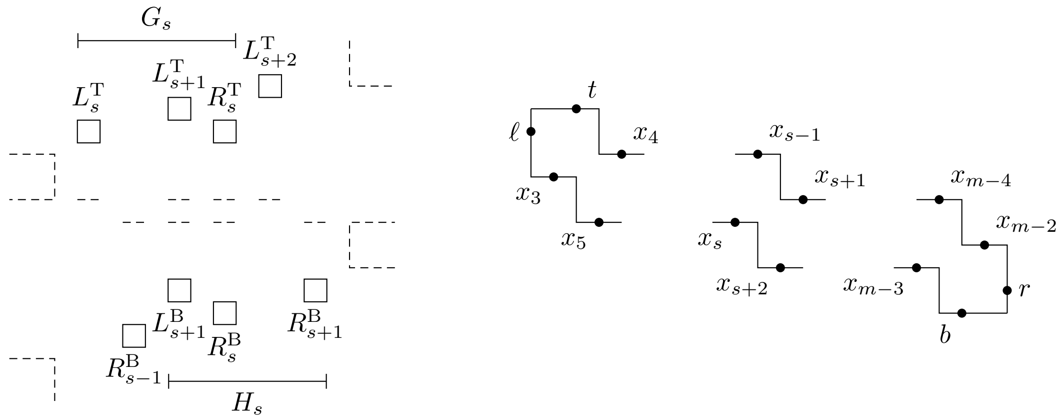

Figure 4.6: *(left)* A sketch of parts of $S(P, X)$. Here, s is odd and $5 \leqslant s \leqslant m - 5$. The dashed lines and open rectangles indicate M and the rest of $S(P, X)$. *(right)* Sketches of three (not necessarily disjoint) parts of P. The solid lines illustrate tallness.

Note that Corollary 4.18 and Lemma 4.19 completely resolve the case $m = 8$.

In the following two subsections, we show that the remaining possibilities also lead to a contradiction. In Section 4.4.2, we treat the easier case, where $\phi(t) \in G_s$ for some odd s with $5 \leqslant s \leqslant m - 5$, and $\phi(b)$ is to the right of R_{s+1} (i.e., to the right of H_s). This also handles the symmetric case where $\phi(b) \in H_s$ and $\phi(t)$ is to the left of G_s. In Section 4.4.3, we consider the case where $\phi(t) \in G_s$ and $\phi(b) \in H_s$.

4.4.2 $\phi(t)$, $\phi(b)$ in the middle and far from each other

The following lemma is central to this subsection and will also be useful later on.

Lemma 4.20. *Let s be odd with $5 \leqslant s \leqslant m - 5$ such that $\phi(t) \in G_s$. Then $\phi(x_s)$ is below the expandable row, or $\phi(\ell), \phi(t), \phi(x_s) \in L_{s+1}$.*

Proof. Assume that $\phi(\ell), \phi(t), \phi(x_s) \in L_{s+1}$ does not hold. We show that then $\phi(x_s)$ is below the expandable row. Note that $\phi(t) \in G_s$ implies that ϕ maps no 1-entry into G^{T}_u for $u > s$.

Case 1: $\phi(\ell) \notin L_s \cup L_{s+1} \cup R_s$. Then $\phi(\ell)$ is below $L^{\mathrm{T}}_s \cup R^{\mathrm{T}}_s$. Since x_4 is to the right of t and below ℓ, this implies that $x_4 \in M \cup B$. Since x_4 is above x_s, this means that $\phi(x_s)$ is in the bottom row of M or further below, so $\phi(x_s)$ is below the expandable row.

Case 2: $\phi(t) \notin L_s \cup L_{s+1} \cup R_s$. Then $\phi(t) \in R_{s-1}$, so $\phi(t)$ is below the expandable row, implying the same for $\phi(x_s)$.

Case 3: $\phi(\ell), \phi(t) \in L_s \cup R_s$. Since ϕ does not map any 1-entry to a position below $L^{\mathrm{T}}_s \cup R^{\mathrm{T}}_s$ and above M, we have $\mathrm{height}_\phi(L^{\mathrm{T}}_s \cup L^{\mathrm{M}}_s \cup R^{\mathrm{T}}_s \cup R^{\mathrm{M}}_s) \leqslant \mathrm{d}^{\mathrm{v}}(t, x_s)$. Thus, $\phi(x_s)$ is in the bottom row of M or further below.

Case 4: $\phi(\ell) \in L_s$ and $\phi(t) \in L_{s+1}$. Since x_4 is below ℓ and to the right of t, we have either $\phi(x_4) \in M \cup B$ or $\phi(x_4) \in R^{\mathrm{T}}_s$. In the former case, we are done, as in case 1. In the latter case, note that ϕ does not map any 1-entry of P into a row below $L_S^{\mathrm{T}} \cup R_s^{\mathrm{T}}$ and above M, thus $\mathrm{height}_\phi(R_s^{\mathrm{T}} \cup R_s^{\mathrm{M}}) = \mathrm{d}^{\mathrm{v}}(x_{s-1}, x_s) \leqslant \mathrm{d}^{\mathrm{v}}(x_4, x_s)$ by Observation 4.8. Thus, $\phi(x_s)$ is below the expandable row.

Case 5: $\phi(\ell), \phi(t) \in L_{s+1}$ and $\phi(x_s) \notin L_{s+1}$. Since x_s is to the right of t, this means that $\phi(x_s)$ is to the right of L_{s+1}. Suppose x_s is above the expandable row. Then all 1-entries in $L_{s+1}^{\mathrm{M}} \cup L_{s+1}^{\mathrm{B}}$ are below or in the same row as x_s. Tallness of X, together with the fact that ϕ maps no two 1-entries into the same row, implies that ϕ maps no 1-entries into $L_{s+1}^{\mathrm{M}} \cup L_{s+1}^{\mathrm{B}}$. But then ϕ maps every 1-entry of P either to L_{s+1}^{T} or to the right and below L_{s+1}^{T}, and $\phi(t) \in L_{s+1}^{\mathrm{T}}, \phi(b) \notin L_{s+1}^{\mathrm{T}}$. This means that P is decomposable, a contradiction. □

Lemma 4.21. *For each odd s with $5 \leqslant s \leqslant m-5$, if $\phi(t) \in G_s$ and $\phi(b)$ is to the right of H_s, then $\phi(x_s)$ is below the expandable row.*

Proof. Suppose $\phi(x_s)$ is above the expandable row. By Lemma 4.20, $\phi(\ell), \phi(t), \phi(x_s) \in L_{s+1}$. Since $\mathrm{height}(L_{s+1}^{\mathrm{T}}) < \mathrm{d}^{\mathrm{v}}(t, x_s)$, we know that $\phi(x_s)$ is below L_{s+1}^{T}, so $\phi(x_s)$ must be in the top row of L_{s+1}^{M}.

x_{s+1} is above and to the right of x_s, implying that $\phi(x_{s+1}) \in L_{s+1}^{\mathrm{T}} \cup R_s^{\mathrm{T}}$. Further, x_{s+2} is below x_s, so $\phi(x_{s+2}) \in M \cup B$, and $x_s <_{\mathrm{h}} x_{s+2} <_{\mathrm{h}} x_{s+1}$, so $\phi(x_{s+2}) \in L_{s+1} \cup R_s$. Since $\phi(b)$ is to the right of H_s by assumption, we know that ϕ maps no 1-entry to H_s^{B}. Thus, $\phi(x_{s+2}) \in L_{s+1}^{\mathrm{M}} \cup R_s^{\mathrm{M}}$.

But now $\phi(x_s), \phi(x_{s+2}) \in L_{s+1}^{\mathrm{M}} \cup R_s^{\mathrm{M}}$, so ϕ maps no further 1-entries to M. Therefore, ϕ maps every 1-entry either to $A = L_{s+1}^{\mathrm{T}} \cup L_{s+1}^{\mathrm{M}} \cup R_s^{\mathrm{T}} \cup R_s^{\mathrm{M}}$, or below and to the right of A (and $\phi(t) \in A$, $\phi(b) \notin A$). This means P is decomposable, a contradiction. □

We now consider a simple special case.

Lemma 4.22. *If $\phi(t) \in G_s$ for some odd s with $5 \leqslant s \leqslant m-5$, then $\phi(b) \notin L_{m-2}$.*
Moreover, if $\phi(b) \in H_s$ for some odd s with $5 \leqslant s \leqslant m-5$, then $\phi(t) \notin R_3$.

Proof. By symmetry, it suffices to show the first statement. Suppose $\phi(b) \in L_{m-2}$. Then, $\phi(b)$ is to the right of R_{m-4}, and thus to the right of R_{s+1}, so Lemma 4.21 implies that $\phi(x_s)$ is below the expandable row. Since x_s is to the left and above b, we have $\phi(x_s) \in L_{m-2}$. Since $x_s <_{\mathrm{h}} x_{s+1} <_{\mathrm{h}} b$, and x_{s+1} is below t, we have $\phi(x_{s+1}) \in L_{m-2}^{\mathrm{M}} \cup L_{m-2}^{\mathrm{B}}$. But $\mathrm{height}_\phi(L_{m-2}^{\mathrm{M}} \cup L_{m-2}^{\mathrm{B}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{m-2}, b) < \mathrm{d}^{\mathrm{v}}(x_{m-4}, b) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, b)$, a contradiction. □

We proceed with the main case of this subsection.

Lemma 4.23. *Let s, u be odd such that $5 \leqslant s < u \leqslant m-5$. If $\phi(t) \in G_s$, then $\phi(b) \notin H_u$.*

Proof. Suppose $\phi(b) \in H_u$. Note that then ϕ maps no 1-entry to G_s^{B} or H_u^{T}.

Lemma 4.21 implies that $\phi(x_s)$ is below the expandable row. Since x_u is below x_s, we have $\phi(x_u) \in B$. Moreover, $\phi(x_u)$ is to the left and above $\phi(b) \in H_u$, so $\phi(x_u) \in H_u^{\mathrm{B}}$.

Since $x_u <_{\mathrm{h}} x_{u-1} <_{\mathrm{h}} x_{u+1} <_{\mathrm{h}} b$, we have $\phi(x_{u-1}), \phi(x_{u+1}) \in H_u$. Note that ϕ maps nothing to H_u^{T}, so $\phi(x_{u-1}) \in H_u^{\mathrm{M}} \cup H_u^{\mathrm{B}}$, and $\phi(x_{u+1})$ is below the expandable row, as x_{u+1} is below x_{u-1}.

Further, $x_u <_{\mathrm{v}} r <_{\mathrm{v}} b$, so $\phi(r) \in H_u^{\mathrm{B}}$. Thus, ϕ does not map any 1-entries to the rows between M and H_u^{B}, so $\mathrm{height}_\phi(M \cup L_{u+1}^{\mathrm{B}} \cup R_{u+1}^{\mathrm{B}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{u+1}, b)$. Since $\phi(x_{u+1})$ is not in the top row of M, this means that $\phi(b)$ is below $L_{u+1}^{\mathrm{B}} \cup R_{u+1}^{\mathrm{B}}$, so $\phi(b) \in R_u^{\mathrm{B}}$.

Consider now $\phi(x_{u+2})$. Since $\phi(x_{u+1})$ is below the expandable row and, by Observations 4.7 and 4.8, $\mathrm{height}_\phi(L_{u+1}^{\mathrm{M}} \cup L_{u+1}^{\mathrm{B}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{u+1}, x_{u+2})$, we know that $\phi(x_{u+2})$ is below L_{u+1}^{B}. Further, x_{u+2} is to the left of b, so $\phi(x_{u+2}) \in R_u^{\mathrm{B}}$. Since r is below x_{u+2}, this implies $\phi(r) \in R_u^{\mathrm{B}}$.

This means that $\phi(x_{u+1}) \in L_{u+1} \cup R_u$. Since $\mathrm{height}(R_u^{\mathrm{B}}) < \mathrm{d}^{\mathrm{v}}(x_u, b) < \mathrm{d}^{\mathrm{v}}(x_{u+1}, b)$, we have $\phi(x_{u+1})$ above R_u^{B}, so $\phi(x_{u+1}) \in L_{u+1}^{\mathrm{B}} \cup L_{u+1}^{\mathrm{M}} \cup R_u^{\mathrm{M}}$. We distinguish two cases:

Case 1: $\phi(x_{u+1}) \in L_{u+1}^{\mathrm{B}} \cup L_{u+1}^{\mathrm{M}}$. Since $\phi(x_{u+1})$ is below the expandable row, tallness of X implies that ϕ maps no 1-entry to R_u^{M}. But then ϕ maps all 1-entries to R_u^{B} or above and to the left of R_u^{B} (recall that ϕ maps no 1-entries to H_u^{T}). Thus, P is decomposable (since $\phi(b) \in R_u^{\mathrm{B}}$, $\phi(t) \notin R_u^{\mathrm{B}}$), a contradiction.

Case 2: $\phi(x_{u+1}) \in R_u^{\mathrm{M}}$. Since $\phi(x_{u-1}) \in H_u^{\mathrm{M}} \cup H_u^{\mathrm{B}}$ is above $\phi(x_{u+1})$, this means that $\phi(x_{u-1}) \in L_{u+1}^{\mathrm{M}} \cup R_u^{\mathrm{M}}$. Note that ϕ cannot map any further 1-entries to M. But this means that ϕ maps all 1-entries either to $H_u^{\mathrm{M}} \cup H_u^{\mathrm{B}}$ or above and to the left of $H_u^{\mathrm{M}} \cup H_u^{\mathrm{B}}$, again contradicting that P is indecomposable. □

Corollary 4.24. *There is some odd s with $5 \leqslant s \leqslant m-5$ such that $\phi(t) \in G_s$ and $\phi(b) \in H_s$.*

Proof. Suppose first that $\phi(t) \in R_3$. Then Corollary 4.18 and Lemma 4.19 imply that $\phi(b) \in H_s$ for some odd s with $5 \leqslant s \leqslant m-5$. But then Lemma 4.22 implies that $\phi(t) \notin R_3$, a contradiction. A similar argument shows that $\phi(b) \notin L_{m-2}$.

As such, there are odd s, u with $5 \leqslant s, u \leqslant m-5$ such that $\phi(t) \in G_s$ and $\phi(b) \in H_u$. Lemma 4.23 implies $u \leqslant s$. If $u < s$, then $\phi(t), \phi(b) \in L_s \cup R_{s-1}$, contradicting Lemma 4.16. Thus, $s = u$. □

4.4.3 $\phi(t), \phi(b)$ in the middle and close to each other

In this subsection, we show that Corollary 4.24 also leads to a contradiction, which shows that our assumption that $S(P, X)$ contains P must have been false. Figure 4.6 will be useful throughout this subsection. We start with the case $\phi(t) \in L_s$.

Lemma 4.25. *$\phi(t) \notin L_s$ and $\phi(b) \notin R_{s+1}$ for each odd s with $5 \leqslant s \leqslant m-5$.*

Proof. By symmetry, it suffices to prove the first statement. Suppose $\phi(t) \in L_s$. Lemma 4.20 implies that $\phi(x_s)$ is below the expandable row and thus to the right of L_s.

We will consider several possibilities for the location of $\phi(b)$ and $\phi(r)$. Before, we make some observations. Corollary 4.24 implies that $\phi(b) \in H_s$. Since $\phi(x_s)$ is below the expandable row, $\phi(x_{s+2}), \phi(b) \in B$, and thus $\phi(x_{s+2}), \phi(b) \in H_s^{\mathrm{B}} = L_{s+1}^{\mathrm{B}} \cup R_s^{\mathrm{B}} \cup R_{s+1}^{\mathrm{B}}$. Since $x_{s+2} \leqslant_{\mathrm{v}} x_{m-3} <_{\mathrm{v}} r$, this also implies $\phi(r) \in H_s^{\mathrm{B}}$. This means that ϕ does not map any 1-entry into the rows between M and $L_{s+1}^{\mathrm{B}} \cup R_{s+1}^{\mathrm{B}}$, so $\mathrm{height}_\phi(M \cup L_{s+1}^{\mathrm{B}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, x_{s+2})$ and $\mathrm{height}_\phi(M \cup R_{s+1}^{\mathrm{B}}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, b)$.

Case 1: $\phi(b) \in L^{\mathrm{B}}_{s+1}$. Then we have $\phi(r) \in L^{\mathrm{B}}_{s+1} \cup R^{\mathrm{B}}_{s+1}$. Since $\mathrm{height}_\phi(M \cup L^{\mathrm{B}}_{s+1}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, x_{s+2}) < \mathrm{d}^{\mathrm{v}}(x_{s+1}, b)$, we have $\phi(x_{s+1}) \in T$. Since x_{s+1} is to the left of b, we have $\phi(x_{s+1}) \in L^{\mathrm{T}}_s$. Moreover, $t <_{\mathrm{v}} \ell <_{\mathrm{v}} x_{s+1}$ implies $\phi(\ell) \in L_s$. But $\mathrm{width}(L_s) = \mathrm{d}^{\mathrm{h}}(\ell, x_s) - 1 < \mathrm{d}^{\mathrm{h}}(\ell, x_{s+1})$, a contradiction.

Case 2: $\phi(b) \in R^{\mathrm{B}}_{s+1}$. Then $\phi(r) \in R^{\mathrm{B}}_{s+1}$. Since $\mathrm{height}_\phi(M \cup R^{\mathrm{B}}_{s+1}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, b)$, we know that $\phi(x_{s+1})$ is above the expandable row, and therefore to the left of R_{s+1} (by Observation 4.6).

Since x_{s-1} is above x_{s+1}, we have $\phi(x_{s-1}) \in T$, implying $\phi(x_{s-1}) \in L^{\mathrm{T}}_s \cup R^{\mathrm{T}}_s$ and thus $\phi(\ell) \in L_s$. Further, $\mathrm{width}(L^{\mathrm{T}}_s) < \mathrm{d}^{\mathrm{h}}(\ell, x_{s-1})$, so $\phi(x_{s-1}) \in R^{\mathrm{T}}_s$.

Finally, since $\phi(x_{s+2}) \in B$ and x_{s+2} is to the left of x_{s+1}, we have $\phi(x_{s+2}) \in L^{\mathrm{B}}_{s+1}$. But then $\phi(x_{s+2})$ is to the left of $\phi(x_{s-1}) \in R^{\mathrm{T}}_s$, while x_{s+2} is to the right of x_{s-1}, a contradiction.

Case 3: $\phi(b), \phi(r) \in R^{\mathrm{B}}_s$. We consider the location of $\phi(x_{s-1})$. Note that $\phi(x_{s-1}) \in G_s$, since $t <_{\mathrm{h}} x_{s-1} <_{\mathrm{h}} r$.

First, suppose that $\phi(x_{s-1}) \in R_s$. Let q_s be the 1-entry of P in the row below x_s. We have $\phi(q_s) \in B$, because $\phi(x_s)$ is below the expandable row. Since X is tall, q_s is to the right of x_{s-1}, so $\phi(q_s) \in R^{\mathrm{B}}_s$. But $\mathrm{height}(R^{\mathrm{B}}_s) \leqslant \mathrm{d}^{\mathrm{v}}(x_s, b) - 2 = \mathrm{d}^{\mathrm{v}}(q_s, b) - 1$, a contradiction.

Second, suppose $\phi(x_{s-1}) \in L_{s+1}$. Since $\phi(t) \in L_s$, this means $\phi(x_{s-1}) \in M \cup B$. By tallness of X, there are no 1-entries in P that are above and to the right of x_{s-1}, so ϕ does not map any 1-entry to R^{T}_s. Note that ϕ must map some 1-entry y to R^{M}_s. Otherwise, ϕ maps all 1-entries to R^{B}_s or above and to the left of R^{B}_s (and $\phi(t) \notin R^{\mathrm{B}}_s$, $\phi(b) \in R^{\mathrm{B}}_s$), so P is decomposable.

By tallness of X, and since y is to the left of x_{s-1}, we know that y must be below x_{s-1}. Thus, $\phi(x_{s-1})$ is in the top row of M, and $\phi(y)$ is in the bottom row of M. But since M only consists of two rows, ϕ maps no further 1-entries to M, so ϕ maps all 1-entries either to $H^{\mathrm{M}}_s \cup H^{\mathrm{B}}_s$ or to the left and above $H^{\mathrm{M}}_s \cup H^{\mathrm{B}}_s$. This again implies that P is decomposable, a contradiction.

Third, suppose $\phi(x_{s-1}) \in R^{\mathrm{M}}_{s-1}$. Then $\phi(x_{s-1})$ is below the expandable row (by Observation 4.6), so $\phi(x_s) \in B$. But x_s also lies to the left of x_{s-1} and above b, a contradiction.

Finally, suppose $\phi(x_{s-1}) \in L_s$. Since $t <_{\mathrm{h}} x_s <_{\mathrm{h}} x_{s-1}$, this implies $\phi(x_s) \in L_s$. But $\phi(x_s)$ is below the expandable row, contradicting Observation 4.6.

Case 4: $\phi(b) \in R^{\mathrm{B}}_s$ and $\phi(r) \in R^{\mathrm{B}}_{s+1}$. Then $\phi(x_{s+2})$ is above and not to the right of R^{B}_s. Together with the fact that $\phi(x_{s+2}) \in B$, this implies $\phi(x_{s+2}) \in L^{\mathrm{B}}_{s+1}$.

Since $\mathrm{height}_\phi(L^{\mathrm{B}}_{s+1} \cup M) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, x_{s+2})$, we know that $\phi(x_{s+1})$ is above the expandable row, and thus $\phi(x_{s-1}) \in T$, implying $\phi(x_{s-1}) \in L^{\mathrm{T}}_s \cup R^{\mathrm{T}}_s$. Moreover, since $\mathrm{width}(L^{\mathrm{T}}_s) < \mathrm{d}^{\mathrm{h}}(\ell, x_{s-1})$, we have $\phi(x_{s-1}) \in R^{\mathrm{T}}_s$. Now $\phi(x_{s-1}) \in R^{\mathrm{T}}_s$ is to the right of $\phi(x_{s+2}) \in L^{\mathrm{B}}_{s+1}$, but x_{s-1} is to the left of x_{s+2}, a contradiction. □

The next few lemmas deal with the case that $\phi(t) \in L_{s+1}$.

Lemma 4.26. *Let s be odd with $5 \leqslant s \leqslant m-5$. If $\phi(t) \in L_{s+1}$, then $\phi(b) \notin L_{s+1}$.*

Proof. Suppose $\phi(t), \phi(b) \in L_{s+1}$. Since $\mathrm{width}(L_{s+1}) < \mathrm{d}^{\mathrm{h}}(\ell, x_{s+1}) < \mathrm{d}^{\mathrm{h}}(\ell, b)$, we know that $\phi(\ell)$ is to the left of L_{s+1}.

$t <_{\mathrm{h}} x_4 <_{\mathrm{h}} x_s <_{\mathrm{h}} x_{s+1} <_{\mathrm{h}} b$ implies that $\phi(x_4), \phi(x_s)\phi(x_{s+1}) \in L_{s+1}$. Moreover, $\phi(x_4)$ is below L^{T}_{s+1}, since x_4 is below ℓ. This implies that x_{s+1} is below the expandable row, which in turn implies that $x_s \in L^{\mathrm{B}}_{s+1}$.

Since r is below x_s and above b, we have $\phi(r) \in L^{\mathrm{B}}_{s+1} \cup R^{\mathrm{B}}_{s+1}$, implying that ϕ maps no 1-entry into the rows between M and $L^{\mathrm{B}}_{s+1} \cup R^{\mathrm{B}}_{s+1}$. But then $\mathrm{height}_\phi(L^{\mathrm{B}}_{s+1} \cup M) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s+1}, b)$, so $\phi(x_{s+1})$ is above the expandable row, a contradiction. □

Lemma 4.27. *Let s be odd with $5 \leqslant s \leqslant m-5$. If $\phi(t) \in L_{s+1}$, then $\phi(b) \in R_s$.*

Proof. Assume $\phi(t) \in L_{s+1}$. By Corollary 4.24, we have $\phi(b) \in H_s$. Lemmas 4.25 and 4.26 imply that $\phi(b) \notin L_{s+1} \cup R_{s+1}$. If $\phi(b) \in L_{s+2}$, then $\phi(b)$ is above the expandable row. But then $\phi(r)$ is to the right of R_s, below L^{T}_{s+2}, and in T, which is impossible. The only remaining possibility is that $\phi(b) \in R_s$. □

Lemma 4.28. *Let s be odd with $5 \leqslant s \leqslant m-5$. If $\phi(t) \in L_{s+1}$, then $\phi(x_{s-1}) \in L_{s+1}$ and $\phi(x_{s+2}) \in R_s$.*

Proof. By Lemma 4.27 and symmetry, it suffices to show that $\phi(x_{s-1}) \in L_{s+1}$. Suppose not. By Lemma 4.27, $\phi(b) \in R_s$, so $t <_{\mathrm{h}} x_{s-1} <_{\mathrm{h}} b$ implies that $\phi(x_{s-1}) \in R_s$. Let $q_s \in E(P)$ be the 1-entry of P in the row directly below x_s.

We claim that $\phi(x_s)$ is below the expandable row, and thus $\phi(q_s) \in B$. If $\phi(x_{s-1}) \in M \cup B$, then $\phi(x_s)$ is indeed below the expandable row, since x_s is below x_{s-1}. Otherwise, $\phi(x_{s-1}) \in R^{\mathrm{T}}_s$, which implies that $\phi(\ell) \in L^{\mathrm{T}}_s \cup L^{\mathrm{T}}_{s+1}$, so ϕ maps no 1-entry into the rows between $L^{\mathrm{T}}_s \cup R^{\mathrm{T}}_s$ and M. Thus, $\mathrm{height}_\phi(R^{\mathrm{T}}_s \cup R^{\mathrm{M}}_s) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s-1}, x_s)$, implying that $\phi(x_s)$ is below the expandable row. This proves the claim.

We have height $\mathrm{height}(R^{\mathrm{B}}_s) \leqslant \mathrm{d}^{\mathrm{v}}(x_s, b) - 2 = \mathrm{d}^{\mathrm{v}}(q_s, b) - 1$, which implies that $\phi(q_s) \notin R^{\mathrm{B}}_s$. Since X is tall, q_s is to the right of x_{s-1} and thus $\phi(q_s)$ is not to the left of R_s. Since $\phi(q_s) \in B \setminus R^{\mathrm{B}}_s$, this implies that $\phi(q_s)$ is to the right of R_s, so $\phi(r)$ is to the right of R_s, and thus above R^{B}_s.

Consider now x_{s+2}. First, $x_{s-1} <_{\mathrm{h}} x_{s+2} <_{\mathrm{h}} b$ implies that $\phi(x_{s+2}) \in R_s$. Since x_s is below the expandable row, $\phi(x_{s+2}) \in R^{\mathrm{B}}_s$. But then $\phi(x_{s+2})$ is below $\phi(r)$, a contradiction. □

Lemma 4.29. *Let s be odd with $5 \leqslant s \leqslant m-5$. If $\phi(t) \in L_{s+1}$, then $\phi(E(P)) \subseteq L_{s+1} \cup R_s$.*

Proof. We show that $\phi(\ell) \in L_{s+1}$ and $\phi(r) \in R_s$. By Lemma 4.27, we have $\phi(b) \in R_s$. Thus, by symmetry, it suffices to prove $\phi(\ell) \in L_{s+1}$. Suppose $\phi(\ell) \notin L_{s+1}$. Then $\phi(\ell)$ is below L^{T}_{s+1}. Since x_{s-1} is below ℓ and $\phi(x_{s-1}) \in L_{s+1}$ by Lemma 4.28, we have $\phi(x_{s-1}) \in L^{\mathrm{M}}_{s+1} \cup L^{\mathrm{B}}_{s+1}$.

Since $\mathrm{d}^{\mathrm{v}}(x_{s-1}, x_{s+2}) > 1$, this means $\phi(x_{s+2}) \in B$. More precisely, by Lemma 4.28, we have $\phi(x_{s+2}) \in R^{\mathrm{B}}_s$. Since r is below x_{s+2}, we also have $\phi(r) \in R^{\mathrm{B}}_s$.

Now consider $\phi(x_{s+1})$. Since $\mathrm{height}(R^{\mathrm{B}}_s) < \mathrm{d}^{\mathrm{v}}(x_{s+1}, b)$, we know that $\phi(x_{s+1})$ is above R^{B}_s. Further, $x_{s+2} <_{\mathrm{h}} x_{s+1} <_{\mathrm{h}} b$ implies $\phi(x_{s+1}) \in R^{\mathrm{T}}_s \cup R^{\mathrm{M}}_s$.

x_{s-1} is above x_{s+1}, so we have $\phi(x_{s-1}), \phi(x_{s+1}) \in M$. This implies that $\mathrm{d}^{\mathrm{v}}(x_{s-1}, x_{s+1}) \leqslant \mathrm{d}^{\mathrm{v}}_{\phi}(\phi(x_{s-1}), \phi(x_{s+1})) = 1$, so x_{s-1} is in the row directly above x_{s+1} in P. Note that this means that the top row of L^{M}_{s+1} contains precisely x^{s+1}_{s-1}, and thus $\phi(x_{s-1}) = x^{s+1}_{s-1}$.

Finally, consider $\phi(x_s)$. We know that $\phi(x_s)$ is not to the right of $x^{s+1}_{s} \in L_{s+1}$, otherwise $\mathrm{d}^{\mathrm{h}}(\phi(x_s), \phi(x_{s-1})) < \mathrm{d}^{\mathrm{h}}(x^{s+1}_{s}, x^{s+1}_{s-1}) = \mathrm{d}^{\mathrm{h}}(x_s, x_{s-1})$. Moreover, $\phi(x_s)$ must be below x^{s+1}_{s}. Indeed, $\phi(r) \in R_s$ implies that ϕ maps no 1-entries between $L^{\mathrm{B}}_{s+1} \cup R^{\mathrm{B}}_{s+1}$ and M, implying $\mathrm{d}^{\mathrm{v}}_{\phi}(x^{s+1}_{s-1}, x^{s+1}_{s}) \leqslant \mathrm{d}^{\mathrm{v}}(x_{s-1}, x_s) - 1$. So $\phi(x_s)$ is below and to the left of x^{s+1}_{s}. But then tallness of X implies that $\phi(x_s)$ is below $\phi(b) \in H^{\mathrm{B}}_{s}$, a contradiction. □

Lemma 4.30. *$\phi(t) \notin L_{s+1}$ and $\phi(b) \notin R_s$ for each odd s with $5 \leqslant s \leqslant m-5$.*

Proof. By symmetry, it suffices to show the first statement. Suppose $\phi(t) \in L_{s+1}$. By Lemmas 4.28 and 4.29, we have $\phi(x_{s-1}) \in L_{s+1}$ as well as $\phi(x_{s+2}), \phi(b) \in R_s$ and $\phi(P) \subseteq L_{s+1} \cup R_s$.

Since $t <_{\mathrm{h}} x_s <_{\mathrm{h}} x_{s-1}$, we also have $\phi(x_s) \in L_{s+1}$. Symmetrically, $\phi(x_{s+1}) \in R_s$. Let $p_{s+1} \in E(P)$ be the 1-entry of P in the row directly above x_{s+1}, and let $q_s \in E(P)$ be the 1-entry in the row directly below x_s. Since $\mathrm{height}(L^{\mathrm{T}}_{s+1}) = \mathrm{d}^{\mathrm{v}}(t, p_{s+1}) - 1$, we know that $\phi(p_{s+1})$ is below L^{T}_{s+1}, and, symmetrically, $\phi(q_s)$ is above R^{B}_{s}. With $p_{s+1} <_{\mathrm{v}} x_{s+1} <_{\mathrm{v}} x_s <_{\mathrm{v}} q_s$, we have $\phi(x_s), \phi(x_{s+1}), \phi(q_s), \phi(p_{s+1}) \in L^{\mathrm{M}}_{s+1} \cup L^{\mathrm{B}}_{s+1} \cup R^{\mathrm{T}}_{s} \cup R^{\mathrm{M}}_{s}$.

Our goal for the remainder of the proof is to find two 1-entries $y_1, y_2 \in E(P)$ such that the sequence $Y = x_1, x_2, \ldots, x_s, y_1, y_2, x_{s+1}, \ldots, x_m$ is a traversal of P. For this, we have to show that (i) $y_2 <_{\mathrm{h}} y_1$, as well as (ii) $y_1 <_{\mathrm{v}} x_{s+1}$, and (iii) $x_s <_{\mathrm{v}} y_2$ (note that then $x_{s-1} <_{\mathrm{h}} y_2 <_{\mathrm{h}} y_1 <_{\mathrm{h}} x_{s+2}$ and $x_{s-1} <_{\mathrm{v}} y_1 <_{\mathrm{v}} y_2 <_{\mathrm{v}} x_{s+2}$ follow from tallness of X). The existence of such a traversal implies that X is extendable, contradicting our assumption.

We consider two cases. First, assume that q_s is to the left of p_{s+1}. Then we simply choose $y_1 = p_{s+1}$ and $y_2 = q_s$. By definition, p_{s+1} is above x_{s+1} and q_s is below x_s, and (i) follows by assumption.

Second, assume that p_{s+1} is to the left of q_s. Then either $\phi(p_{s+1}) \in L_{s+1}$ or $\phi(q_s) \in R_s$. By symmetry, we can assume the former, which implies $\phi(p_{s+1}) \in L^{\mathrm{M}}_{s+1} \cup L^{\mathrm{B}}_{s+1}$. Since $\phi(x_{s+1}) \in R^{\mathrm{T}}_{s} \cup R^{\mathrm{M}}_{s}$ and x_{s+1} is below p_{s+1}, we have $\phi(p_{s+1}), \phi(x_{s+1}) \in M$. More precisely, $\phi(p_{s+1}) = p^{s+1}_{s+1} \in L^{\mathrm{M}}_{s+1}$ and $\phi(x_{s+1}) = q^{s}_{s} \in R^{\mathrm{M}}_{s}$.

Now choose $y_1, y_2 \in E(P)$ such that $y^{s+1}_{1} = \phi(x_{s-1})$ and $y^{s+1}_{2} = \phi(x_s)$. Note that this is well-defined, since $\phi(x_{s-1}), \phi(x_s) \in L_{s+1}$. We immediately have (i) from $x_s <_{\mathrm{h}} x_{s-1}$.

We now show (ii), that y_1 is above x_{s+1}. Since x_{s-1} is above x_{s+1} and $\phi(x_{s+1}) = q^{s}_{s}$, we know that $\phi(x_{s-1}) = y^{s+1}_{1}$ is above the expandable row. Since the expandable row in L_{s+1} corresponds to the row containing x_{s+1} in P, this means that y_1 is above x_{s+1}.

Finally, we show (iii), that y_2 is below x_s. Since $\phi(r) \in R_s$, we know that ϕ maps no 1-entry into the rows between M and L^{B}_{s+1}. Since ϕ also maps no 1-entry into the expandable row, we have $\mathrm{d}^{\mathrm{v}}_{\phi}(p^{s+1}_{s+1}, x^{s+1}_{s}) \leqslant \mathrm{d}^{\mathrm{v}}(p_{s+1}, x_s) - 1$. As $\phi(p_{s+1}) = p^{s+1}_{s+1}$, this means that $\phi(x_s) = y^{s+1}_{1}$ is below x^{s+1}_{s}, implying (iii). □

The remaining cases are now easy:

Lemma 4.31. *$\phi(t) \notin R_{s-1} \cup R_s$ and $\phi(b) \notin L_{s+1} \cup L_{s+2}$ for each odd s with $5 \leqslant s \leqslant m-5$.*

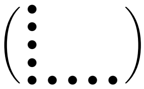

Figure 5.1: An indecomposable non-permutation matrix without a spanning oscillation.

Proof. By symmetry, it suffices to show the first statement. Suppose $\phi(t) \in R_{s-1} \cup R_s$. By Corollary 4.24 and Lemmas 4.25 and 4.30, we have $\phi(b) = L_{s+1} \cup L_{s+2}$.

Suppose first that $\phi(t) \in R_{s-1}$. Then $\phi(t)$ is below the expandable row, meaning that $\phi(\ell)$ is below M, but not to the right of R_{s-1}. But $\phi(b)$ is above R^{B}_{s-1}, implying that $\phi(\ell)$ is also above R^{B}_{s-1}, a contradiction.

Second, if $\phi(t) \in R_s$, then $\phi(b) \in L_{s+2}$, since b is to the right of t. A symmetric argument shows that $\phi(r)$ is above M, below L^{T}_{s+2}, and not to the right of L_{s+2}, a contradiction. □

Lemmas 4.25, 4.30 and 4.31 imply that $\phi(t) \notin G_s$, contradicting Corollary 4.24. As such, our assumption that ϕ is an embedding of P into $S(P, X)$ must be false. This concludes the proof of Lemma 4.9.

5. Conclusion and open problems

We showed that each indecomposable permutation matrix has bounded saturation function, thereby completing the classification of saturation functions of permutation matrices. Our proofs imply the upper bound $\mathrm{sat}(P, n) \leqslant 9k^4$ for an indecomposable $k \times k$ permutation matrix P (note that the largest witness $S(P, X)$ is not larger than $2k^2 \times k^2$, and Lemma 1.7 combines it with its 90-degree rotation, resulting in a $3k^2 \times 3k^2$ matrix). It would be interesting to improve this bound, especially if a simpler construction for patterns satisfying the conditions of Lemma 4.9 can be found. Note that for general patterns with bounded saturation functions, no upper bound for $\mathrm{sat}(P, n)$ in terms of P is known, as noted by Fulek and Keszegh [FK21].

We also characterized a large class of non-permutation patterns with bounded saturation function, including very dense matrices (Theorem 1.2). Still, a full characterization of the saturation functions of all matrices remains out of reach. Note that there are indecomposable patterns without spanning oscillations, see, e.g., Figure 5.1. Thus, new techniques are likely required to fully resolve this problem.

Our results trivially imply that every permutation matrix with a vertical witness also has a horizontal witness. It would be interesting to determine whether this is true for arbitrary patterns.

It is also possible to consider the saturation functions of *sets* of patterns. If $\mathcal{P}$ is a set of patterns, let a matrix M be $\mathcal{P}$-saturating if M avoids each $P \in \mathcal{P}$, and adding a single 1-entry in M creates an occurrence of some $P \in \mathcal{P}$ in M. Let $\mathrm{sat}(\mathcal{P}, n)$ be the minimum weight of $\mathcal{P}$-saturating matrices. Since our witnesses for $k \times k$ permutation matrices have size at most $3k^2 \times 3k^2$, they avoid all patterns with one side of side length more than $3k^2$. Thus, if $\mathcal{P}$ contains one indecomposable permutation matrix, and arbitrarily many much larger patterns, our results imply that $\mathrm{sat}(\mathcal{P}, n) \in \mathcal{O}(1)$.

It would be interesting to determine the saturation functions for, say, all pairs of two permutation matrices of the same size. Gerbner, Nagy, Patkós, and Vizer [GNPV22] observed that certain saturation problems for two-dimensional posets can be reduced to saturation problems for sets of matrix patterns. However, these sets usually contain both permutation matrices and non-permutation matrices (of similar size).

A. Proof of Lemma 1.19

Lemma A.1. *Let P, M be matrices, and let P have no empty rows or columns. Then P is contained in M if and only if there is an embedding of P into M.*

Proof. Say P is $q \times s$ and M is $m \times n$. Suppose $P = (p_{i,j})_{i,j}$ is contained in M. Then there are rows $r_1 < r_2 < \cdots < r_q$ and columns $c_1 < c_2 < \cdots < c_r$ such that $p_{i,j} \leqslant m_{r_i,c_j}$ for each $i \in [q], j \in [r]$. Now simply define $\phi(i,j) = (r_i, c_j)$. Clearly, $\phi(E(P)) \subseteq E(M)$. Moreover, consider $(i,j),(i',j') \in E(P)$. We have $i < i'$ if and only if $r_i < r_{i'}$, and $j < j'$ if and only if $r_j < r_{j'}$. Thus ϕ is an embedding of P into M.

Now suppose $\phi\colon E(P) \to E(M)$ is an embedding of P into M. Note that $x, y \in E(P)$ are in the same row (resp. column) if and only if $\phi(x), \phi(y)$ are in the same row (resp. column). Thus, $\phi(E(P))$ intersects exactly q rows and s columns. Let $r_1 < r_2 < \cdots < r_q$ be those rows and $c_1 < c_2 < \cdots < c_r$ be those columns. We show that $\phi(i,j) = (r_i, c_j)$ for each $(i,j) \in E(P)$. Let $x_1, x_2, \dots, x_m \in E(P)$ such that x_i is in the i-th row for each $i \in [m]$, and let r'_i be the row of M containing $\phi(x_i)$. Clearly $r'_1 \geqslant r_1$. By induction, we further have $r'_i \geqslant r_i$ for each $i \in [m]$. Similarly, $r'_m \leqslant r_m$, and, again by induction, $r'_i \leqslant r_i$ for each $i \in [m]$. This implies that $\phi(i,j)$ is in the r_i-th row of M for every $(i,j) \in E(P)$. An analogous argument shows that $\phi(i,j)$ is in the c_j-th column of M.

Since ϕ is an embedding, we have $(r_i, c_j) = \phi(i,j) \in E(M)$ for each $(i,j) \in E(P)$. Thus, $p_{i,j} \leqslant m_{r_i,c_j}$ for each $(i,j) \in [q] \times [s]$, so P is contained in M. □

Acknowledgements

The author would like to thank Jesse Geneson, László Kozma, and the anonymous reviewers for their helpful comments.

References

[ARSz99] Noga Alon, Lajos Rónyai, and Tibor Szabó. Norm-graphs: Variations and applications. *Journal of Combinatorial Theory, Series B*, 76(2):280–290, 1999. doi:10.1006/jctb.1999.1906.

[BC21] Richard A. Brualdi and Lei Cao. Pattern-avoiding (0,1)-matrices and bases of permutation matrices. *Discrete Applied Mathematics*, 304:196–211, 2021. doi:10.1016/j.dam.2021.07.039.

[Ber20] Benjamin Aram Berendsohn. Matrix patterns with bounded saturation function. 2020. arXiv:2012.14717v1.

[BG91] Dan Bienstock and Ervin Györi. An extremal problem on sparse 0-1 matrices. *SIAM Journal on Discrete Mathematics*, 4(1):17–27, 1991. doi:10.1137/0404002.

[BRV08] Robert Brignall, Nik Ruškuc, and Vincent Vatter. Simple permutations: Decidability and unavoidable substructures. *Theoretical Computer Science*, 391(1):150–163, 2008. doi:10.1016/j.tcs.2007.10.037.

[CGK+15] Parinya Chalermsook, Mayank Goswami, Laszlo Kozma, Kurt Mehlhorn, and Thatchaphol Saranurak. Pattern-avoiding access in binary search trees. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 410–423, 2015. doi:10.1109/FOCS.2015.32.

[FH92] Zoltán Füredi and Péter Hajnal. Davenport-schinzel theory of matrices. *Discrete Mathematics*, 103(3):233–251, 1992. doi:10.1016/0012-365X(92)90316-8.

[FK21] Radoslav Fulek and Balázs Keszegh. Saturation problems about forbidden 0-1 submatrices. *SIAM Journal on Discrete Mathematics*, 35(3):1964–1977, 2021. doi:10.1137/20M1376327.

[Ful09] Radoslav Fulek. Linear bound on extremal functions of some forbidden patterns in 0–1 matrices. *Discrete Mathematics*, 309(6):1736–1739, 2009. doi:10.1016/j.disc.2008.02.040.

[Für90] Zoltán Füredi. The maximum number of unit distances in a convex n-gon. *J. Comb. Theory, Ser. A*, 55(2):316–320, 1990. doi:10.1016/0097-3165(90)90074-7.

[Gen09] Jesse T. Geneson. Extremal functions of forbidden double permutation matrices. *Journal of Combinatorial Theory, Series A*, 116(7):1235–1244, 2009. doi:10.1016/j.jcta.2009.03.004.

[Gen21] Jesse Geneson. Almost all permutation matrices have bounded saturation functions. *The Electronic Journal of Combinatorics*, 28(2):P2.16, 2021. doi:10.37236/10124.

[GNPV22] Dániel Gerbner, Dániel T. Nagy, Balázs Patkós, and Máté Vizer. Forbidden subposet problems in the grid. *Discrete Mathematics*, 345(3):112720, 2022. doi:10.1016/j.disc.2021.112720.

[Kesz09] Balázs Keszegh. On linear forbidden submatrices. *Journal of Combinatorial Theory, Series A*, 116(1):232–241, 2009. doi:10.1016/j.jcta.2008.05.006.

[Kla00] Martin Klazar. The Füredi-Hajnal conjecture implies the stanley-wilf conjecture. In *Formal power series and algebraic combinatorics*, pages 250–255. Springer, 2000. doi:10.1007/978-3-662-04166-6_22.

[Kla01] Martin Klazar. *Enumerative and extremal combinatorics of a containment relation of partitions and hypergraphs.* Habilitation thesis, 2001.

[Mit87] Joseph S. B. Mitchell. Shortest rectilinear paths among obstacles. Technical report, Cornell University Operations Research and Industrial Engineering, 1987.

[MT04] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley–Wilf conjecture. *Journal of Combinatorial Theory, Series A*, 107(1):153–160, 2004. doi:10.1016/j.jcta.2004.04.002.

[Pet10] Seth Pettie. Applications of forbidden 0–1 matrices to search tree and path compression-based data structures. In *Proceedings of the 2010 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1457–1467, 2010. doi:10.1137/1.9781611973075.118.

[Pet11a] Seth Pettie. Degrees of nonlinearity in forbidden 0–1 matrix problems. *Discrete Mathematics*, 311(21):2396–2410, 2011. doi:10.1016/j.disc.2011.06.020.

[Pet11b] Seth Pettie. Generalized Davenport–Schinzel sequences and their 0-1 matrix counterparts. *J. Comb. Theory Ser. A*, 118(6):1863–1895, 2011. doi:10.1016/j.jcta.2011.02.011.

[Pra73] Vaughan R. Pratt. Computing permutations with double-ended queues, parallel stacks and parallel queues. In *Proceedings of the Fifth Annual ACM Symposium on Theory of Computing*, pages 268–277, New York, NY, USA, 1973. ACM. doi:10.1145/800125.804058.

[PT06] János Pach and Gábor Tardos. Forbidden paths and cycles in ordered graphs and matrices. *Israel Journal of Mathematics*, 155(1):359–380, 2006. doi:10.1007/BF02773960.

[Tar05] Gábor Tardos. On 0–1 matrices and small excluded submatrices. *Journal of Combinatorial Theory, Series A*, 111(2):266–288, 2005. doi:10.1016/j.jcta.2004.11.015.

[Vat11] Vincent Vatter. Small permutation classes. *Proceedings of the London Mathematical Society*, 103(5):879–921, 2011. doi:10.1112/plms/pdr017.

[Vat14] Vincent Vatter. Permutation classes. 2014. arXiv:1409.5159v3.