

A framework for minimal hereditary classes of graphs of unbounded clique-width

R. Brignall D. Cocks
School of Mathematics and Statistics
The Open University, UK

March 29, 2022

Abstract

We create a framework for hereditary graph classes \mathcal{G}^δ built on a two-dimensional grid of vertices and edge sets defined by a triple $\delta = \{\alpha, \beta, \gamma\}$ of objects that define edges between consecutive columns, edges between non-consecutive columns (called bonds), and edges within columns. This framework captures all previously proven minimal hereditary classes of graph of unbounded clique-width, and many new ones, although we do not claim this includes all such classes.

We show that a graph class \mathcal{G}^δ has unbounded clique-width if and only if a certain parameter \mathcal{N}^δ is unbounded. We further show that \mathcal{G}^δ is minimal of unbounded clique-width (and, indeed, minimal of unbounded linear clique-width) if another parameter \mathcal{M}^β is bounded, and also δ has defined recurrence characteristics. Both the parameters \mathcal{N}^δ and \mathcal{M}^β are properties of a triple $\delta = (\alpha, \beta, \gamma)$, and measure the number of distinct neighbourhoods in certain auxiliary graphs.

Throughout our work, we introduce new methods to the study of clique-width, including the use of Ramsey theory in arguments related to unboundedness, and explicit (linear) clique-width expressions for subclasses of minimal classes of unbounded clique-width.

1 Introduction

Until 4 years ago only a couple of examples of minimal hereditary classes of unbounded clique-width had been identified, see Lozin [10]. However, more recently many more such classes have been identified, in Atminas, Brignall, Lozin and Stacho [2], Collins, Foniok, Korpelainen, Lozin and Zamaraev [5], and most recently the current authors demonstrated an uncountably infinite family of minimal hereditary classes of unbounded clique-width in [3].

This paper brings together all these examples into a single consistent framework. The framework consists of hereditary graph classes constructed by taking the finite induced subgraphs of an infinite graph \mathcal{P}^δ whose vertices form a two-dimensional array and whose edges are defined by three objects, collectively denoted as a triple $\delta = (\alpha, \beta, \gamma)$. Though we defer full definitions until Section 2, the components of the triple define edges between consecutive columns (α), between non-consecutive columns (β ‘bonds’), and within columns (γ) as follows.

- (a) α is an infinite word from the alphabet $\{0, 1, 2, 3\}$. The four types of α -edge set between consecutive columns can be described as a matching (0), the complement of a matching (1), a chain (2) and the complement of a chain (3).
- (b) β is a symmetric subset of pairs of natural numbers (x, y) . If $(x, y) \in \beta$ then every vertex in column x is adjacent to every vertex in column y .
- (c) γ is an infinite binary word. If the j -th letter of γ is 0 then vertices in column j form an independent set and if it is 1 they form a clique.

We show that these hereditary graph classes \mathcal{G}^δ have unbounded clique-width if and only if a parameter \mathcal{N}^δ measuring the number of distinct neighbourhoods between any two rows of the grid, is unbounded – see Theorem 3.16. We denote Δ as the set of δ -triples for which \mathcal{G}^δ has unbounded clique-width.

Furthermore, we define a subset $\Delta_{\min} \subset \Delta$ such that if $\delta \in \Delta_{\min}$ the hereditary graph class \mathcal{G}^δ is minimal both of unbounded clique-width and of unbounded *linear* clique-width – see Theorem 4.10. These ‘minimal’ δ -triples are characterised by

- (a) δ is recurrent (i.e. any factor δ^* of δ repeats an infinite number of times),
- (b) for any factor δ^* of δ , the subgraphs induced on the columns between two consecutive copies of δ^* (the δ -factor ‘gap’) have bounded clique-width (always true for almost periodic δ), and
- (c) a bound on a parameter \mathcal{M}^β defined by the bond set β which is a measure of the number of distinct neighbourhoods between intervals of a single row.

All hereditary graph classes previously shown to be minimal of unbounded clique-width fit this framework i.e. they are defined by a δ -triple in Δ_{\min} . This is demonstrated in Table 1 which shows their corresponding $\delta = (\alpha, \beta, \gamma)$ values from the framework.

Name	α	β ($x, y \in \mathbb{N}$)	γ
Bipartite permutation [10]	2^∞	\emptyset	0^∞
Unit interval [10]	2^∞	\emptyset	1^∞
Bichain [2]	$(23)^\infty$	$(2x, 2x + 2y + 1)$	0^∞
Split permutation [2]	$(23)^\infty$	$(2x, y) : y > 2x + 1$	$(01)^\infty$
$\alpha \in \{0, 1\}$ [5]	periodic	\emptyset	0^∞
$\alpha \in \{0, 1, 2, 3\}$ [3]	recurrent ¹	\emptyset	0^∞

Table 1: Hereditary graph classes proven to be minimal of unbounded clique-width

¹A set of minimal classes Γ defined by the infinite word α which is recurrent over the alphabet $\{0, 1, 2, 3\}$ and for which the ‘gap’ factors have bounded weight (including all almost periodic α)

The unified viewpoint provided by our framework offers a fuller understanding of the landscape of (the uncountably many known) minimal hereditary classes of unbounded clique-width. This landscape is in stark contrast to the situation for downwards-closed sets of graphs under different orderings and with respect to other parameters. For example, planar graphs are the unique minimal minor-closed class of graphs of unbounded treewidth (see Robertson and Seymour [12]), and circle graphs are the unique minimal vertex-minor-closed class of unbounded rank-width (or, equivalently, clique-width) – see Geelen, Kwon, McCarty and Wollan [9]. Nevertheless, clique-width is more compatible with hereditary classes of graphs than treewidth: if H is an induced subgraph of G , then the clique-width of H is at most the clique-width of G , but the same does not hold in general for treewidth.

Our focus on the minimal classes of unbounded clique-width is due to the following fact: any graph property expressible in MSO_1 logic has a linear time algorithm on graphs with bounded clique-width, see Courcelle, Makowsky and Rotics [7]. As it happens, any proper subclass of a minimal class from our framework also has bounded *linear* clique-width. However, beyond our framework there do exist classes that have bounded clique-width but unbounded linear clique-width, see [1] and [4].

After introducing the necessary definitions in Section 2, the rest of this paper is organised as follows.

We set out in Section 3 our proof determining which hereditary classes \mathcal{G}^δ have unbounded clique-width. Proving a class has unbounded clique-width is done from first principles, using a new method, by identifying a lower bound for the number of labels required for a clique-width expression for an $n \times n$ square graph, using distinguished coloured vertex sets and showing such sets always exist for big enough n using Ramsey theory. For those classes which have bounded clique-width we prove this by providing a general clique-width expression for any graph in the class, using a bounded number of labels.

In Section 4 we prove that the class \mathcal{G}^δ is minimal of unbounded clique-width if $\delta \in \Delta_{\min}$. To do this we introduce an entirely new method of ‘veins and slices’, partitioning the vertices of an arbitrary graph in a proper subclass of \mathcal{G}^δ into sections we call ‘panels’ using vertex colouring. We then create a recursive linear clique-width expression to construct these panels in sequence, allowing recycling of labels each time a new panel is constructed, so that an arbitrary graph can be constructed with a bounded number of labels.

Previous papers on minimal hereditary graph classes of unbounded clique-width have focused mainly on bipartite graphs. The introduction of β -bonds and γ -cliques has significantly broadened the scope of proven minimal classes.

In Section 5 we provide some examples of new hereditary graph classes that are minimal of unbounded clique-width revealed by this approach. Finally, in Section 6, we discuss where the investigation of minimal classes of unbounded clique-width might go next.

2 Preliminaries

2.1 Graphs - General

A graph $G = (V, E)$ is a pair of sets, vertices $V = V(G)$ and edges $E = E(G) \subseteq V(G) \times V(G)$. Unless otherwise stated, all graphs in this paper are simple, i.e. undirected, without loops or multiple edges.

If vertex u is adjacent to vertex v we write $u \sim v$ and if u is not adjacent to v we write $u \not\sim v$. We denote $N(v)$ as the neighbourhood of a vertex v , that is, the set of vertices adjacent to v . A set of vertices is *independent* if no two of its elements are adjacent and is a *clique* if all the vertices are pairwise adjacent. We denote a clique with r vertices as K^r and an independent set of r vertices as \bar{K}^r . A graph is *bipartite* if its vertices can be partitioned into two independent sets, V_1 and V_2 , and is *complete bipartite* if, in addition, each vertex of V_1 is adjacent to each vertex of V_2 .

We will use the notation $H \leq G$ to denote graph H is an *induced subgraph* of graph G , meaning $V(H) \subseteq V(G)$ and two vertices of $V(H)$ are adjacent in H if and only if they are adjacent in G . We will denote the subgraph of $G = (V, E)$ induced by the set of vertices $U \subseteq V$ by $G[U]$. If graph G does not contain an induced subgraph isomorphic to H we say that G is *H-free*.

A class of graphs \mathcal{C} is *hereditary* if it is closed under taking induced subgraphs, that is $G \in \mathcal{C}$ implies $H \in \mathcal{C}$ for every induced subgraph H of G . It is well known that for any hereditary class \mathcal{C} there exists a unique (but not necessarily finite) set of minimal forbidden graphs $\{H_1, H_2, \dots\}$ such that $\mathcal{C} = \text{Free}(H_1, H_2, \dots)$ (i.e. any graph $G \in \mathcal{C}$ is H_i -free for $i = 1, 2, \dots$). We will use the notation $\mathcal{C} \subseteq \mathcal{G}$ to denote that \mathcal{C} is a *hereditary subclass* of hereditary graph class \mathcal{G} ($\mathcal{C} \subsetneq \mathcal{G}$ for a proper subclass).

An *embedding* of graph H in graph G is an injective map $\phi : V(H) \rightarrow V(G)$ such that the subgraph of G induced by the vertices $\phi(V(H))$ is isomorphic to H . In other words, $vw \in E(H)$ if and only if $\phi(v)\phi(w) \in E(G)$. If H is an induced subgraph of G then this can be witnessed by one or more embeddings.

Given a graph $G = (V, E)$ and a subset of vertices $U \subseteq V$, two vertices of $V \setminus U$ will be called *$V \setminus U$ -similar* if they have the same neighbourhood in $V \setminus U$. Thus $V \setminus U$ -similarity is an equivalence relation. The number of such equivalence classes of $V \setminus U$ in G will be denoted $\mu(G, U)$. A special case is when all the equivalence classes are singletons when we call U a *distinguished vertex set*.

A *distinguished pairing* $\{U, W\}$ of size r of a graph $G = (V, E)$ is a pair of vertex subsets $U = \{u_i\} \subseteq V$ and $W = \{w_i\} \subseteq V \setminus U$ with $|U| = |W| = r$ such that the vertices in U have pairwise different neighbourhoods in W (but not necessarily vice-versa). A distinguished pairing is *matched* if the vertices of U and W can be paired (u_i, w_i) so that $u_i \sim w_i$ for each i , and is *unmatched* if the vertices of U and W can be paired (u_i, w_i) so that $u_i \not\sim w_i$ for each i . Clearly the set U of a distinguished pairing $\{U, W\}$ is a distinguished vertex set of $G[U \cup W]$ which gives us the following:

Proposition 2.1. *If $\{U, W\}$ is a distinguished pairing of size r in graph G then $\mu(G[U \cup W], U) = r$.*

2.2 \mathcal{G}^δ hereditary graph classes

The graph classes we consider are all formed by taking the set of finite induced subgraphs of an infinite graph defined on a grid of vertices. We start by defining an infinite empty graph \mathcal{P} with vertices

$$V(\mathcal{P}) = \{v_{i,j} : i, j \in \mathbb{N}\}.$$

We use Cartesian coordinates throughout this paper. Hence, we think of \mathcal{P} as an infinite two-dimensional array in which $v_{i,j}$ represents the vertex in the i -th column (counting from the left) and the j -th row (counting from the bottom). Hence vertex $v_{1,1}$ is in the bottom left corner of the grid and the grid extends infinitely upwards and to the right. The i -th column of \mathcal{P} is the set $C_i = \{v_{i,j} : j \in \mathbb{N}\}$, and the j -th row of \mathcal{P} is the set $R_j = \{v_{i,j} : i \in \mathbb{N}\}$. Likewise, the collection of vertices in columns i to j is denoted $C_{[i,j]}$.

We will add edges to \mathcal{P} using a triple δ of objects that define the edges between consecutive columns, edges between non-consecutive columns and edges within each column.

We refer to a (finite or infinite) sequence of letters chosen from a finite alphabet as a *word*. We denote by ω_i the i -th letter of the word ω . A *factor* of ω is a contiguous subword $\omega_{[i,j]}$ being the sequence of letters from the i -th to the j -th letter of ω . If a is a letter from the alphabet we will denote a^∞ as the infinite word $aaa\dots$, and if $a_1 \dots a_n$ is a finite sequence of letters from the alphabet then we will denote $(a_1 \dots a_n)^\infty$ as the infinite word consisting of the infinite repetition of this factor.

The *length* of a word (or factor) is the number of letters the word contains.

An infinite word ω is *recurrent* if each of its factors occurs in it infinitely many times. We say that ω is *almost periodic* (sometimes called *uniformly recurrent* or *minimal*) if for each factor $\omega_{[i,j]}$ of ω there exists a constant $\mathcal{L}(\omega_{[i,j]})$ such that every factor of ω of length at least $\mathcal{L}(\omega_{[i,j]})$ contains $\omega_{[i,j]}$ as a factor. Finally, ω is *periodic* if there is a positive integer p such that $\omega_k = \omega_{k+p}$ for all k . Clearly, every periodic word is almost periodic, and every almost periodic word is recurrent.

A *bond-set* β is a symmetric subset of $\{(x, y) \in \mathbb{N}^2, |x - y| > 1\}$. For a set $Q \subseteq \mathbb{N}$ we write β_Q to mean the subset of β -bonds $\{(x, y) \in \beta : x, y \in Q\}$. For instance, $\beta_{[i,j]} = \{(x, y) \in \beta : i \leq x, y \leq j\}$.

Let α be an infinite word over the alphabet $\{0, 1, 2, 3\}$, β be a bond set and γ be an infinite binary word. We denote the triple $\delta = (\alpha, \beta, \gamma)$.

We define an infinite graph \mathcal{P}^δ with vertices $V(\mathcal{P})$ and with edges defined by δ as follows:

- (a) α -edges between consecutive columns determined by the letters of the word α . For each $i = 1, 2, \dots$, the edges between C_i and C_{i+1} are:

- (i) $\{(v_{i,j}, v_{i+1,j}) : j \in \mathbb{N}\}$ if $\alpha_i = 0$ (i.e. a matching);
- (ii) $\{(v_{i,j}, v_{i+1,k}) : j \neq k; j, k \in \mathbb{N}\}$ if $\alpha_i = 1$ (i.e. the bipartite complement² of a matching);

²The *bipartite complement* \hat{G} of a bipartite graph G has the same independent vertex sets V_1 and V_2 as G where vertices $v_1 \in V_1$ and $v_2 \in V_2$ are adjacent in \hat{G} if and only if they are not adjacent in G .

- (iii) $\{(v_{i,j}, v_{i+1,k}) : j \geq k; j, k \in \mathbb{N}\}$ if $\alpha_i = 2$;
 - (iv) $\{(v_{i,j}, v_{i+1,k}) : j < k; j, k \in \mathbb{N}\}$ if $\alpha_i = 3$ (i.e. the bipartite complement of a 2).
- (b) β -edges defined by the bond-set β such that $v_{i,x} \sim v_{j,y}$ for all $x, y \in \mathbb{N}$ when $(i, j) \in \beta$ (i.e. a complete bipartite graph between C_i and C_j), and
- (c) γ -edges defined by the letters of the binary word γ such that for any $j, k \in \mathbb{N}$ we have $v_{i,j} \sim v_{i,k}$ if and only if $\gamma_i = 1$ (i.e. C_i forms a clique if $\gamma_i = 1$ and an independent set if $\gamma_i = 0$).

The hereditary graph class \mathcal{G}^δ is the set of all finite induced subgraphs of \mathcal{P}^δ .

Any graph $G \in \mathcal{G}^\delta$ can be witnessed by an embedding $\phi(G)$ into the infinite graph \mathcal{P}^δ . To simplify the presentation we will associate G with a particular embedding in \mathcal{P}^δ depending on the context. We will be especially interested in the induced subgraphs of G that occur in consecutive columns: in particular, an α_j -link is the induced subgraph of G on the vertices of $G \cap C_{[j,j+1]}$, and will be denoted by $G_{[j,j+1]}$. More generally, an induced subgraph of G on the vertices of $G \cap C_{[j,k]}$ will be denoted $G_{[j,k]}$.

For $k \geq 2$ we denote the triple $\delta_{[j,j+k-1]} = (\alpha_{[j,j+k-2]}; \beta_{[j,j+k-1]}; \gamma_{[j,j+k-1]})$ as a k -factor of δ . Thus for a graph $G \in \mathcal{G}^\delta$ with a particular embedding in \mathcal{P}^δ , the induced subgraph $G_{[j,j+k-1]}$ has edges defined by the k -factor $\delta_{[j,j+k-1]}$.

We say that two k -factors $\delta_{[x,x+k]}$ and $\delta_{[y,y+k]}$ are the same if

- (i) for all $i \in [0, k-1]$, $\alpha_{x+i} = \alpha_{y+i}$, and
- (ii) for all $i, j \in [0, k]$, $(x+i, x+j) \in \beta$ if and only if $(y+i, y+j) \in \beta$, and
- (iii) for all $i \in [0, k]$, $\gamma_{x+i} = \gamma_{y+i}$.

We say that a δ -triple is *recurrent* if every k -factor occurs in it infinitely many times. We say that δ is *almost periodic* if for each k -factor $\delta_{[j,k]}$ of δ there exists a constant $\mathcal{L}(\delta_{[j,k]})$ such that every factor of δ of length $\mathcal{L}(\delta_{[j,k]})$ contains $\delta_{[j,k]}$ as a factor.

A *couple set* P is a subset of \mathbb{N} such that if $x, y \in P$ then $|x - y| > 2$. Such a set is used to identify sets of links that have no α -edges between them. We say that two elements $x, y \in P$ are β -dense if both $(x, y+1)$ and $(x+1, y)$ are in β and they are β -sparse when neither of these bonds is in β .

We say the bond-set β is *sparse* in P if every pair in P is β -sparse and is *not sparse* in P if there are no β -sparse pairs in P . Likewise, β is *dense* in P if every pair in P is β -dense and is *not dense* in P if there are no β -dense pairs in P . Clearly it is possible for two elements in P to be neither β -sparse nor β -dense (i.e. when only one of the required bonds is in β). These ideas are used to identify matched and unmatched distinguished pairings (see Lemmas 3.7 and 3.8).

2.3 Clique-width and linear clique-width

Clique-width is a graph width parameter introduced by Courcelle, Engelfriet and Rozenberg in the 1990s [6]. The clique-width of a graph is denoted $\text{cwd}(G)$ and is defined as the minimum number of labels needed to construct G by means of the following four graph operations:

- (a) creation of a new vertex v with label i (denoted $i(v)$),
- (b) adding an edge between every vertex labelled i and every vertex labelled j for distinct i and j (denoted $\eta_{i,j}$),
- (c) giving all vertices labelled i the label j (denoted $\rho_{i \rightarrow j}$), and
- (d) taking the disjoint union of two previously-constructed labelled graphs G and H (denoted $G \oplus H$).

The *linear clique-width* of a graph G denoted $\text{lcw}(G)$ is the minimum number of labels required to construct G by means of only the first three of the operations described above.

Every graph can be defined by an algebraic expression τ using the four (or three for the linear version) operations above, which we will refer to as a (*linear*) *clique-width expression*. This expression is called a k -expression if it uses k different labels.

Alternatively, any clique-width expression τ defining G can be represented as a rooted tree, $\text{tree}(\tau)$, whose leaves correspond to the operations of vertex creation, the internal nodes correspond to the \oplus -operation, and the root is associated with G . The operations η and ρ are assigned to the respective edges of $\text{tree}(\tau)$. In the case of a linear clique-width expression the tree becomes a *caterpillar tree*, that is, a tree that becomes a path after the removal of the leaves.

Clearly from the definition, $\text{lcw}(G) \geq \text{cwd}(G)$. Hence, a graph class of unbounded clique-width is also a class of unbounded linear clique-width. Likewise, a class with bounded linear clique-width is also a class of bounded clique-width.

3 \mathcal{G}^δ graph classes with unbounded clique-width

This section identifies which hereditary classes \mathcal{G}^δ have unbounded clique-width. We prove this is determined by a neighbourhood parameter \mathcal{N}^δ derived from a graph induced on any two rows of the graph \mathcal{P}^δ . We show that \mathcal{G}^δ has unbounded clique-width if and only if \mathcal{N}^δ is unbounded (Theorem 3.16).

3.1 The two-row graph and \mathcal{N}^δ

We will show that the boundedness of clique-width for a graph class \mathcal{G}^δ is determined by the adjacencies between the first two rows of \mathcal{P}^δ (it could, in fact, be any two rows). Hence, the following graph is useful.

A *two-row graph* $T^\delta(Q) = (V, E)$ is the subgraph of \mathcal{P}^δ induced on the vertices $V = R_1(Q) \cup R_2(Q)$ where $R_1(Q) = \{v_{i,1} : i \in Q\}$ and $R_2(Q) = \{v_{j,2} : j \in Q\}$ for finite subset $Q \subseteq \mathbb{N}$.

We define the parameter $\mathcal{N}^\delta(Q) = \mu(T^\delta(Q), (R_1(Q)))$.

Lemma 3.1. *For any fixed $j \in \mathbb{N}$, $\mathcal{N}^\delta([1, n])$ is bounded as $n \rightarrow \infty$ if and only if $\mathcal{N}^\delta([j, n])$ is bounded as $n \rightarrow \infty$.*

Proof. It is easy to see that if there exists N such that $\mathcal{N}^\delta([1, n]) < N$ for all $n \in \mathbb{N}$ then $\mathcal{N}^\delta([j, n]) < N$ for all $n \in \mathbb{N}$.

On the other hand, if $\mathcal{N}^\delta([j, n]) < N$ then $\mathcal{N}^\delta([j-1, n]) < 2N + 1$ since by adding the extra column each 'old' equivalence class could at most be split in 2 and there is one new vertex in each row. By induction we have $\mathcal{N}^\delta([1, n]) < 2^j N + \sum_{i=0}^{j-1} 2^i$ for all $n \in \mathbb{N}$. \square

We will say \mathcal{N}^δ is *unbounded* if $\mathcal{N}^\delta([j, n])$ is unbounded as $n \rightarrow \infty$ for some fixed $j \in \mathbb{N}$. In many cases it is simple to check that \mathcal{N}^δ is unbounded – e.g. the following δ -triples have unbounded \mathcal{N}^δ :

$$(1^\infty, \emptyset, 0^\infty), (2^\infty, \emptyset, 0^\infty), (3^\infty, \emptyset, 0^\infty), (0^\infty, \emptyset, 1^\infty)$$

In Lemma 3.13 we show that \mathcal{N}^δ is unbounded whenever α contains an infinite number of $2s$ or $3s$.

3.2 Clique-width expression and colour partition for an $n \times n$ square graph

We will denote $H_{i,j}^\delta(m, n)$ as the $m(\text{cols}) \times n(\text{rows})$ induced subgraph of \mathcal{P}^δ formed from the rectangular grid of vertices $\{v_{x,y} : x \in [i, i+m-1], y \in [j, j+n-1]\}$. See Figure 1.




Figure 1: $H_{1,1}^\delta(9, 6)$ where $\alpha = 01230123 \dots$ (β and γ edges not shown)

We can calculate a lower bound for the clique-width of the $n \times n$ square graph $H_{j,1}^\delta(n, n)$ (shortened to $H(n, n)$ when δ, j and 1 are clearly implied), by demonstrating a minimum number of labels needed to construct it using the allowed four graph operations, as follows.

Let τ be a clique-width expression defining $H(n, n)$ and $\text{tree}(\tau)$ the rooted tree representing τ . The subtree of $\text{tree}(\tau)$ rooted at a node \oplus_a corresponds to a subgraph of $H(n, n)$ we will call H_a .

We denote by \oplus_{red} and \oplus_{blue} the two children of \oplus_a in $\text{tree}(\tau)$. Let us colour the vertices of H_{red} and H_{blue} red and blue, respectively, and all the other vertices in $H(n, n)$ white. Let $\text{colour}(v)$ denote the colour of a vertex $v \in H(n, n)$ as described above, and $\text{label}(v)$ denote the label of vertex v (if any) at node \oplus_a (If v has no label we say $\text{label}(v) = \epsilon$).

Our identification of a minimum number of labels needed to construct $H(n, n)$ relies on the following observation regarding this vertex colour partition.

Observation 3.2. *Suppose u_1, u_2, w are three vertices in $H(n, n)$ such that u_1 and u_2 are non-white, $u_1 \sim w$ but $u_2 \not\sim w$, and $\text{colour}(w) \neq \text{colour}(u_1)$. Then u_1 and u_2 must have different labels at node \oplus_a .*

This is true because the edge u_1w still needs to be created, whilst respecting the non-adjacency of u_2 and w . We now focus on sets of blue and sets of nonblue vertices (Equally, we could have chosen red-nonred). Observation 3.2 leads to the following key lemma which is the basis of much which follows.

Lemma 3.3. *For graph $H(n, n)$ let U and W be two disjoint vertex sets with induced subgraph $H = H(n, n)[U \cup W]$ such that $\mu(H, U) = r$. Then if the vertex colouring described above gives $\text{colour}(u) = \text{blue}$ for all $u \in U$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in W$ then the clique-width expression τ requires at least r labels at node \oplus_a .*

Proof. Choose one representative vertex from each equivalence class in U . For any two such representatives u_1 and u_2 there must exist a w in W such that $u_1 \sim w$ but $u_2 \not\sim w$ (or vice versa). By Observation 3.2 u_1 and u_2 must have different labels at node \oplus_a . This applies to any pair of representatives u_1, u_2 and hence all r such vertices must have distinct labels. \square

Note that from Proposition 2.1 a distinguished pairing gives us the sets U and W required for Lemma 3.3. The following lemmas identify structures in $H(n, n)$ that give us these distinguished pairings.

We denote by $H_{[y, y+1]}$ the α_y -link $H(n, n) \cap C_{[y, y+1]}$ where $y \in [j, j + n - 2]$. We will refer to a (adjacent or non-adjacent) *blue-nonblue pair* to mean two vertices, one of which is coloured blue and one non-blue, such that they are in consecutive columns, where the blue vertex could be to the left or the right of the nonblue vertex. If we have a set of such pairs with the blue vertex on the same side (i.e. on the left or right) then we say the pairs in the set have the same *polarity*.

Lemma 3.4. *Suppose that $H_{[y, y+1]}$ contains a horizontal pair b_1, b_2 of blue vertices and at least one nonblue vertex n_1, n_2 in each column, but not on the top or bottom row (see Figure 2).*

- (a) *If $\alpha_y \in \{0, 2, 3\}$ then $H_{[y, y+1]}$ contains a non-adjacent blue-nonblue pair.*
- (b) *If $\alpha_y \in \{1, 2, 3\}$ then $H_{[y, y+1]}$ contains an adjacent blue-nonblue pair.*




Figure 2: Horizontal blue-blue pair in $H_{[y, y+1]}$ (nonblue vertices in yellow)

Proof. If $\alpha_y = 0$ then both $\{b_1, n_1\}$ and $\{b_2, n_2\}$ form a non-adjacent blue-nonblue pair (Figure 2 A). If $\alpha_y = 1$ then both $\{b_1, n_1\}$ and $\{b_2, n_2\}$ form an adjacent blue-nonblue pair (Figure 2 B).

If $\alpha_y \in \{2, 3\}$ and the nonblue vertices $\{n_1, n_2\}$ in each column are either both above or both below the horizontal blue pair $\{b_1, b_2\}$ then it can be seen that one of the pairs $\{b_1, n_1\}$ or $\{b_2, n_2\}$ forms an adjacent blue-nonblue pair and the other forms a non-adjacent blue-nonblue pair (Figure 2 C). If the nonblue vertices in each column are either side of the blue pair (one above and one below) then the pairs $\{b_1, n_1\}$ and $\{b_2, n_2\}$ will both be adjacent (or non-adjacent) blue-nonblue pairs (See Figure 2 D). In this case we need to appeal to a 5-th vertex s which will form a non-adjacent (or adjacent) set with either n_1 or b_2 depending on its colour. Thus we always have both a non-adjacent and adjacent blue-non-blue pair when $\alpha_y \in \{2, 3\}$. \square

Lemma 3.5. Suppose $H_{[y, y+1]}$ contains $r \geq 3$ sets of horizontal blue-nonblue pairs of vertices $\{b_i, n_i\}$, $i = 1, \dots, r$, with the same polarity (see Figure 3). Then, irrespective of the value of α_y , it contains

- (a) a matched distinguished pairing $\{\mathcal{U}, \mathcal{W}\}$ of size $r - 1$ such that $\text{colour}(u) = \text{blue}$ for all $u \in \mathcal{U}$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in \mathcal{W}$, and
- (b) an unmatched distinguished pairing $\{\mathcal{U}', \mathcal{W}'\}$ of size $r - 1$ such that $\text{colour}(u') = \text{blue}$ for all $u' \in \mathcal{U}'$ and $\text{colour}(w') \neq \text{blue}$ for all $w' \in \mathcal{W}'$.

Proof. This is easily observable from Figure 3 for $r = 3$. If we set $\mathcal{U} = \{b_1, b_2\}$, $\mathcal{W} = \{n_1, n_2\}$, $\mathcal{U}' = \{b_2, b_3\}$ and $\mathcal{W}' = \{n_1, n_2\}$ then one of $\{\mathcal{U}, \mathcal{W}\}$ and $\{\mathcal{U}', \mathcal{W}'\}$ is a matched distinguished pairing of size 2 and the other is an unmatched distinguished pairing of size 2, irrespective of the value of α_y . Simple induction establishes this for all $r \geq 3$. \square

Lemma 3.6. Suppose $H_{[y, y+1]}$ contains a horizontal blue-nonblue pair of vertices $\{b_1, n_1\}$, not the top or bottom row, and at least one nonblue vertex n_2 in the same column as b_1 . Then $H_{[y, y+1]}$ contains both an adjacent and a non-adjacent blue-nonblue pair of vertices, irrespective of the value of α_y (see Figure 4).




Figure 3: 3 horizontal blue-nonblue pairs in $H_{[y,y+1]}$ (nonblue vertices in yellow)




Figure 4: 1 horizontal blue-nonblue pair in $H_{[y,y+1]}$ (nonblue vertices in yellow)

Proof. If $\alpha_y \in \{0, 2\}$ then the horizontal blue-nonblue pair $\{b_1, n_1\}$ is adjacent, and given a non-blue vertex n_2 in the same column as b_1 , we can find a vertex s in the same column as n_1 that forms a non-adjacent pairing with either b_1 or n_2 depending on its colour (See Figure 4 A and C). If $\alpha_y \in \{1, 3\}$ then the horizontal blue-nonblue pair $\{b_1, n_1\}$ is non-adjacent, and given a nonblue vertex n_2 in the same column as b_1 , we can find a vertex s in the same column as n_1 that forms an adjacent pairing with either b_1 or n_2 depending on its colour (See Figure 4 B and D). \square

In Lemmas 3.4, 3.5 and 3.6 we identified blue-nonblue pairs within a particular link $H_{[y,y+1]}$. The next two lemmas identify distinguished pairings across link-sets. Let $P \subset [j, j + n - 2]$ be a couple set (see definition on page 6) of size r with corresponding α_y -links $H_{[y,y+1]} \leq H(n, n)$ for each $y \in P$.

Lemma 3.7. *If β is not dense in P and each $H_{[y,y+1]}$ has an adjacent blue-nonblue pair with the same polarity, then we can combine these pairs to form a matched distinguished pairing $\{U, W\}$ of size r where*




Figure 5: Adjacent blue-nonblue vertex pairs, β not dense (nonblue vertices in yellow)




Figure 6: Non-adjacent blue-nonblue vertex pairs, β not sparse (nonblue vertices in yellow)

the vertices of \mathcal{U} are blue and the vertices of \mathcal{W} nonblue.

Proof. Suppose $s, t \in P$ such that (v_s, v_{s+1}) and (v_t, v_{t+1}) are two adjacent blue-nonblue pairs in different links, with $v_s, v_t \in \mathcal{U}$ and $v_{s+1}, v_{t+1} \in \mathcal{W}$. Consider the two possible β bonds (v_s, v_{t+1}) and (v_{s+1}, v_t) . If neither of these bonds exist then v_s is distinguished from v_t by both v_{s+1} and v_{t+1} (see Figure 5 (i)). If one of these bonds exists then v_s is distinguished from v_t by either v_{s+1} or v_{t+1} (see Figure 5 (ii) and (iii)). Both bonds cannot exist as β is not dense in P . Note that the bonds (v_s, v_t) and (v_{s+1}, v_{t+1}) are not relevant in distinguishing v_s from v_t since, if they exist, they connect blue to blue and nonblue to nonblue.

So any two blue vertices $v_s, v_t \in \mathcal{U}$ are distinguished by the two nonblue vertices $v_{s+1}, v_{t+1} \in \mathcal{W}$ and hence $\{\mathcal{U}, \mathcal{W}\}$ is a matched distinguished pairing of size r . \square

Lemma 3.8. *If β is not sparse in P and each $H_{[y, y+1]}$ has a non-adjacent blue-nonblue pair with the same polarity, then we can combine these pairs to form an unmatched distinguished pairing $\{\mathcal{U}, \mathcal{W}\}$ of size r where the vertices of \mathcal{U} are blue and the vertices of \mathcal{W} nonblue.*

Proof. This is very similar to the proof of Lemma 3.7 and is demonstrated in Figure 6. \square

3.3 Two colour partition cases to consider

Having identified structures that give us a lower bound on labels required for a clique-width expression for $H(n, n)$, we now apply this knowledge to the following subtree of $\text{tree}(\tau)$.

Let \oplus_α be the lowest node in $\text{tree}(\tau)$ such that H_α contains all the vertices in rows 2 to $(n-1)$ in a column of $H(n, n)$. We reserve rows 1 and n so that we may apply Lemmas 3.4 and 3.6.

Thus $H(n, n)$ contains at least one column where vertices in rows 2 to $(n - 1)$ are non-white but no column has entirely blue or red vertices in rows 2 to $(n - 1)$. Let C_b be a non-white column. Without loss of generality we can assume that the number of blue vertices in column C_b between rows 2 and $(n - 1)$ is at least $(n/2) - 1$ otherwise we could swap red for blue.

Now consider rows 2 to $(n - 1)$. We have two possible cases:

Case 1 Either none of the rows with a blue vertex in column C_b has blue vertices in every column to the right of C_b , or none of the rows with a blue vertex in column C_b has blue vertices in every column to the left of C_b . Hence, we have at least $\lceil n/2 \rceil - 1$ rows that have a horizontal blue-nonblue pair with the same polarity.

Case 2 One row R_r has a blue vertex in column C_b and blue vertices in every column to the right of C_b and one row R_l has a blue vertex in column C_b and blue vertices in every column to the left of C_b . Hence, either on row R_r or row R_l , we must have a horizontal set of consecutive blue vertices of size at least $\lceil n/2 \rceil + 1$.

To prove unboundedness of clique-width we will show that for any $r \in \mathbb{N}$ we can find an $n \in \mathbb{N}$ so that any clique-width expression τ for $H(n, n)$ requires at least r labels in $\text{tree}(\tau)$, whether this is a 'Case 1' or 'Case 2' scenario.

To address both cases we will need the following classic result:

Theorem 3.9 (Ramsey [11] and Diestel [8]). *For every $r \in \mathbb{N}$, every graph of order at least 2^{2r-3} contains either K^r or \overline{K}^r as an induced subgraph.*

We handle first Case 1, for all values of $\delta = (\alpha, \beta, \gamma)$.

Lemma 3.10. *For any $\delta = (\alpha, \beta, \gamma)$ and any $r \in \mathbb{N}$, if $n \geq 9 \times 2^{4r-1}$ and τ is a clique-width expression for $H(n, n)$ that results in Case 1 at node \oplus_α , then τ requires at least r labels to construct $H(n, n)$.*

Proof. In Case 1 we have, without loss of generality, at least $\lceil n/2 \rceil - 1$ horizontal blue-nonblue vertex pairs but we don't know which links these fall on.

If there are at least $\sqrt{n/2}$ such pairs on the same link then using Lemma 3.5 we have a matched distinguished pairing $\{U, W\}$ of size $\sqrt{n/2} - 1 > r$ such that $\text{colour}(u) = \text{blue}$ for all $u \in U$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in W$.

If there is no link with $\sqrt{n/2}$ such pairs then there must be at least one such pair on at least $\sqrt{n/2}$ different links. From Lemma 3.6 each such link contains both an adjacent and non-adjacent blue-nonblue pair. It follows from the pigeonhole principle that there is a subset of these of size $\sqrt{n/2}/4$ where the adjacent blue-nonblue pairs have the same polarity and also the non-adjacent blue-nonblue pairs have the same polarity. We use this subset (Note, the following argument applies whether the blue vertex is on the left or right for the adjacent and non-adjacent pairs). If we take the index of the first column in each remaining link, and then take every third one of these, we have a couple set P where $|P| \geq \sqrt{n/2}/12$, with corresponding

link set $S_L = \{H_{[y, y+1]} : y \in P\}$, such that the adjacent blue-nonblue pair in each link has the same polarity and the non-adjacent blue-nonblue pair in each link has the same polarity.

Define the graph G_P so that $V(G_P) = P$ and for $x, y \in V(G_P)$ we have $x \sim y$ if and only if they are β -dense (see definition on page 6). From Theorem 3.9 for any r , as $|P| \geq \sqrt{n/2}/12 \geq 2^{2r-3}$ then there exists a couple set $Q \subseteq P$ such that G_Q is either K^r or $\overline{K^r}$.

If G_Q is $\overline{K^r}$, it follows that β is not dense in Q , and S_L contains a link set of size r corresponding to the couple set Q where each link has an adjacent blue-nonblue pair with the same polarity. Applying Lemma 3.7 this gives us a matched distinguished pairing $\{U, W\}$ of size r such that $\text{colour}(u) = \text{blue}$ for all $u \in U$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in W$.

If G_Q is K^r , it follows that β is not sparse in Q , and S_L contains a link set of size r corresponding to the couple set Q where each link has a non-adjacent blue-nonblue pair with the same polarity. Applying Lemma 3.8 this gives us an unmatched distinguished pairing $\{U, W\}$ of size r such that $\text{colour}(u) = \text{blue}$ for all $u \in U$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in W$.

In each case we can construct a distinguished pairing $\{U, W\}$ of size r such that $\text{colour}(u) = \text{blue}$ for all $u \in U$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in W$. Hence, from Lemma 3.3 τ uses at least r labels to construct $H(n, n)$. \square

3.4 When α has an infinite number of 2s or 3s

For Case 2 we must consider different values for α separately. We denote $m_{23}(n)$ to be the total number of 2s and 3s in $\alpha_{[1, n-1]}$.

Lemma 3.11. *Let $\delta = (\alpha, \beta, \gamma)$ be a triple such that α contains an infinite number of 2s or 3s. Then for any β and γ and any $r \in \mathbb{N}$, if $m_{23}(n) \geq 3 \times 2^{2r}$ and τ is a clique-width expression for $H(n, n)$ that results in Case 2 at node \oplus_α , then τ requires at least r labels to construct $H(n, n)$.*

Proof. Remembering that C_b is the non-white column, without loss of generality we can assume that there are at least $(m_{23}(n)/2)$ 2- or 3-links to the right of C_b , since otherwise we could reverse the order of the columns. In Case 2 each link has a horizontal blue-blue vertex pair with at least one nonblue vertex in each column, so using Lemma 3.4 we have both an adjacent and non-adjacent blue-nonblue pair in each of these links.

It follows from the pigeonhole principle that there is a subset of these of size $(m_{23}(n)/8)$ where the adjacent blue-nonblue pairs have the same polarity and also the non-adjacent blue-nonblue pairs have the same polarity. We use this subset. If we take the index of the first column in each remaining link, and then take every third one of these, we have a couple set P where $|P| \geq (m_{23}(n)/24)$, with corresponding link set $S_L = \{H_{[y, y+1]} : y \in P\}$, such that the adjacent blue-nonblue pair in each link has the same polarity and the non-adjacent blue-nonblue pair in each link has the same polarity.

As in the proof of Lemma 3.10, we define a graph G_P so that $V(G_P) = P$ and for $x, y \in V(G_P)$ we have $x \sim y$ if and only if they are β -dense. From Theorem 3.9 for any r , as $|P| \geq (m_{23}(n))/24 \geq 2^{2r-3}$ then there exists a couple set $Q \subseteq P$ such that G_Q is either K^r or $\overline{K^r}$.

We now proceed in an identical way to Lemma 3.10 to show that we can always construct a distinguished pairing $\{U, W\}$ of size r such that $\text{colour}(u) = \text{blue}$ for all $u \in U$ and $\text{colour}(w) \neq \text{blue}$ for all $w \in W$. Hence, from Lemma 3.3 τ uses at least r labels to construct $H(n, n)$. \square

Corollary 3.12. *For any triple $\delta = (\alpha, \beta, \gamma)$ such that α has an infinite number of 2s or 3s the hereditary graph class \mathcal{G}^δ has unbounded clique-width.*

Proof. This follows directly from Lemma 3.10 for Case 1 and Lemma 3.11 for Case 2, since for any $r \in \mathbb{N}$ we can choose n big enough so that $n \geq 9 \times 2^{4r-1}$ and $m_{23}(n) \geq 3 \times 2^{2r}$ so that whether we are in Case 2 or Case 1 at node \oplus_α we require at least r labels for any clique-width expression for $H(n, n)$. \square

We are aiming to state our result in terms of unbounded \mathcal{N}^δ so we also require the following.

Lemma 3.13. *For any triple $\delta = (\alpha, \beta, \gamma)$ such that α has an infinite number of 2s or 3s the parameter \mathcal{N}^δ is unbounded.*

Proof. If there is an infinite number of 2s in α we can create a couple set P of any required size such that $\alpha_x = 2$ for every $x \in P$, so that in the two-row graph (see Section 3.1) $v_{1,x} \not\sim v_{2,x+1}$ and $v_{2,x} \sim v_{1,x+1}$ (i.e. we have both an adjacent and non-adjacent pair in the α_x -link).

We now apply the same approach as in Lemmas 3.10 and 3.11, applying Ramsey theory to the graph G_P defined in the same way as before. Then for any r we can set $|P| \geq 2^{2r-3}$ so that there exists a couple set $Q \subseteq P$ where G_Q is either K^r or \overline{K}^r .

If G_Q is \overline{K}^r it follows that β is not dense in Q . So for any $x, y \in Q$, $v_{1,x+1}$ and $v_{1,y+1}$ have different neighbourhoods in $R_2(Q)$ since they are distinguished by either $v_{2,x}$ or $v_{2,y}$. Hence, if n is the highest natural number in Q then $\mathcal{N}^\delta([1, n+1]) \geq r$.

If G_Q is K^r it follows that β is not sparse in Q . So for any $x, y \in Q$, $v_{1,x}$ and $v_{1,y}$ have different neighbourhoods in $R_2(Q)$ since they are distinguished by either $v_{2,x+1}$ or $v_{2,y+1}$. Hence, $\mathcal{N}^\delta([1, n+1]) \geq r$.

Either way, we have $\mathcal{N}^\delta([1, n+1]) \geq r$, but r can be arbitrarily large, so \mathcal{N}^δ is unbounded.

A similar argument applies if there is an infinite number of 3s. \square

3.5 When α has a finite number of 2s and 3s

If α contains only a finite number of 2s and 3s then there exists $J \in \mathbb{N}$ such that $\alpha_j \in \{0, 1\}$ for $j > J$. In Case 2 we are interested in the adjacencies of the blue vertices to the non-blue vertices in each column. Although the nonblue vertices could be in any row, in fact, if α is over the alphabet $\{0, 1\}$, the row index of the non-blue vertices does not alter the blue-nonblue adjacencies.

In Case 2, let Q be the set of column indices of the horizontal set of consecutive blue vertices in row R_r of $H(n, n)$ and let $U_1 = \{v_{i,r} : i \in Q\}$ be this horizontal set of blue vertices. Let

$U_2 = \{u_j : j \in Q\}$ be the corresponding set of nonblue vertices such that $u_j \in C_j$. We have the following:

Lemma 3.14. *In Case 2, with U_1 and U_2 defined as above, if α is a word over the alphabet $\{0, 1\}$ then for any $i, j \in Q$, $v_{i,r} \sim u_j$ in \mathcal{P}^δ if and only if $v_{i,1} \sim v_{j,2}$ in the two-row graph $T^\delta(Q)$.*

Proof. Considering the vertex sets $U_1 \cup U_2$ of \mathcal{P}^δ and $R_1(Q) \cup R_2(Q)$ of $T^\delta(Q)$ (see Section 3.1) we have:

- (a) For $i = j$ both $v_{j,r} \sim u_j$ and $v_{j,1} \sim v_{j,2}$ if and only if $\gamma_j = 1$.
- (b) For $|i - j| > 1$ both $v_{i,r} \sim u_j$ and $v_{i,1} \sim v_{j,2}$ if and only if $(i, j) \in \beta$.
- (c) For $j = i + 1$ both $v_{i,r} \sim u_j$ and $v_{i,1} \sim v_{j,2}$ if and only if $\alpha_i = 1$.

Hence $v_{i,r} \sim u_j$ if and only if $v_{i,1} \sim v_{j,2}$. □

Lemma 3.15. *If $\delta = (\alpha, \beta, \gamma)$ where α is an infinite word over the alphabet $\{0, 1, 2, 3\}$ with a finite number of 2s and 3s, then the hereditary graph class \mathcal{G}^δ has unbounded clique-width if and only if \mathcal{N}^δ is unbounded.*

Proof. First, we prove that \mathcal{G}^δ has unbounded clique-width if \mathcal{N}^δ is unbounded.

As α has a finite number of 2s and 3s there exists a $J \in \mathbb{N}$ such that $\alpha_j \in \{0, 1\}$ if $j > J$.

As \mathcal{N}^δ is unbounded this means that from Lemma 3.1 for any $r \in \mathbb{N}$ there exist $N_1, N_2 \in \mathbb{N}$ such that, setting $Q_1 = [J + 1, J + N_1]$ and $Q_2 = [J + N_1 + 1, J + N_1 + N_2]$, then $\mathcal{N}^\delta(Q_1) \geq r$ and $\mathcal{N}^\delta(Q_2) \geq r$.

Denote the $n \times n$ graph $H'(n, n) = H_{J+1,1}^\delta(n, n) \in \mathcal{G}^\delta$. As described in Section 3.3 we again consider the two possible cases for a clique-width expression τ for $H'(n, n)$ at a node \oplus_α which is the lowest node in $\text{tree}(\tau)$ such that H_α contains a column of $H'(n, n)$.

Case 1 is already covered by Lemma 3.10 for $n \geq 9 \times 2^{4r-1}$.

In Case 2, one row R_r of $H'(N_1 + N_2, N_1 + N_2)$ has a blue vertex in column C_b and blue vertices in every column to the right of C_b and one row R_l has a blue vertex in column C_b and blue vertices in every column to the left of C_b .

If $b \leq J + N_1$ then consider the graph to the right of C_b . We know every column has a blue vertex in row R_r and a non-blue vertex in a row other than R_r . The column indices to the right of C_b includes Q_2 . It follows from Lemma 3.14 that the neighbourhoods of the blue set to the non-blue set are identical to the neighbourhoods in graph $T^\delta(Q_2)$ between the vertex sets $R_1(Q_2)$ and $R_2(Q_2)$.

On the other hand if $b > N_1$ we can make an identical claim for the graph to the left of C_b which now includes the column indices for Q_1 . It follows from Lemma 3.14 that the neighbourhoods of the blue set to the non-blue set are identical to the neighbourhoods in graph $T^\delta(Q_1)$ between the vertex sets $R_1(Q_1)$ and $R_2(Q_1)$.

As both $\mathcal{N}^\delta(Q_1) = \mu(T^\delta(Q_1), R_1(Q_1)) \geq r$ and $\mathcal{N}^\delta(Q_2) = \mu(T^\delta(Q_2), R_2(Q_2)) \geq r$ it follows from Lemma 3.3 that any clique-width expression for $H'(n, n)$ with $n \geq (N_1 + N_2)$ resulting in Case 2 requires at least r labels.

For any $r \in \mathbb{N}$ we can choose n big enough so that $n \geq \max\{9 \times 2^{4r-1}, (N_1 + N_2)\}$ so that whether we are in Case 1 or Case 2 at node \oplus_a we require at least r labels for any clique-width expression for $H'(n, n)$. Hence, \mathcal{G}^δ has unbounded clique-width if \mathcal{N}^δ is unbounded.

Secondly, suppose that \mathcal{N}^δ is bounded, so that there exists $N \in \mathbb{N}$ such that $\mathcal{N}^\delta([J + 1, n]) = \mu(T^\delta([J + 1, n]), R_1([J + 1, n])) < N$ for all $n > J$.

We claim $\text{lcwd}(\mathcal{G}^\delta) \leq 2J + N + 2$. For we can create a linear clique-width expression using no more than $2J + N + 2$ labels that constructs any graph in \mathcal{G}^δ row by row, from bottom to top and from left to right.

For any graph $G \in \mathcal{G}^\delta$ let it have an embedding in the grid \mathcal{P} between columns 1 and $M > J$.

We will use the following set of $2J + N + 2$ labels:

- 2 current vertex labels: a_1 and a_2 ;
- J current row labels for first J columns: $\{c_y : y = 1, \dots, J\}$;
- J previous row labels for first J columns: $\{p_y : y = 1, \dots, J\}$;
- N partition labels: $\{s_y : y = 1, \dots, N\}$.

We allocate a default partition label s_y to each column of $G_{[J+1, M]}$ according to the $R_2([J + 1, M])$ -similar equivalence classes of the vertex set $R_1([J + 1, M])$ in $T^\delta([J + 1, M])$. There are at most N partition sets $\{S_y\}$ of $R_1([J + 1, M])$, and if vertex $v_{i,1}$ is in S_y , $1 \leq y \leq N$, then the default partition label for vertices in column i is s_y . It follows that for two default column labels, s_x and s_y , vertices in columns with label s_y are either all adjacent to vertices in columns with label s_x or they are all non-adjacent (except the special case of vertices in consecutive columns and the same row, which will be dealt with separately in our clique-width expression).

Carry out the following row-by-row linear iterative process to construct each row j , starting with row 1.

- (i) Construct the first J vertices in row j , label them c_1 to c_J and build any edges between them as necessary.
- (ii) Insert required edges from each vertex labelled c_1, \dots, c_J to vertices in lower rows in columns 1 to J . This is possible because the vertices in lower rows in column i ($1 \leq i \leq J$) all have label p_i and have the same adjacency with the vertices in the current row.
- (iii) Relabel vertices labelled c_1, \dots, c_J to p_1, \dots, p_{J-1}, a_2 respectively.
- (iv) Construct and label subsequent vertices in row j (columns $J + 1$ to M), as follows.
 - (a) Construct the next vertex in column i and label it a_1 (or a_2).

- (b) If $\alpha_{i-1} = 0$ then insert an edge from the current vertex $v_{i,j}$ (label a_1) to the previous vertex $v_{i-1,j}$ (label a_2).
 - (c) Insert edges to vertices that are adjacent as a result of the partition $\{S_y\}$ described above. This is possible because all previously constructed vertices with a particular default partition label s_y are either all adjacent or all non-adjacent to the current vertex.
 - (d) Insert edges from the current vertex to vertices labelled p_j ($1 \leq j \leq J$) as necessary.
 - (e) Relabel vertex $v_{i,j-1}$ to its default partition label s_y .
 - (f) Create the next vertex in row i and label it a_2 (or a_1 alternating).
- (v) When the end of the row is reached, repeat for the next row.

Hence we can construct any graph in the class with at most $2J + N + 2$ labels so the clique-width of \mathcal{G}^δ is bounded if \mathcal{N}^δ is bounded. \square

Corollary 3.12, Lemma 3.13 and Lemma 3.15 give us the following:

Theorem 3.16. *For any triple $\delta = (\alpha, \beta, \gamma)$ the hereditary graph class \mathcal{G}^δ has unbounded clique-width if and only if \mathcal{N}^δ is unbounded.*

We will denote Δ as the set of all δ -triples for which the class \mathcal{G}^δ has unbounded clique-width.

4 \mathcal{G}^δ graph classes that are minimal of unbounded clique-width

To show that some given \mathcal{G}^δ is a minimal class of unbounded clique-width we must show that any proper hereditary subclass \mathcal{C} has bounded clique-width. If \mathcal{C} is a hereditary graph class such that $\mathcal{C} \subsetneq \mathcal{G}^\delta$ then there must exist a non-trivial finite forbidden graph F that is in \mathcal{G}^δ but not in \mathcal{C} . In turn, this graph F must be an induced subgraph of some $H_{j,1}^\delta(k, k)$ for some j and $k \in \mathbb{N}$, and thus $\mathcal{C} \subseteq \text{Free}(H_{j,1}^\delta(k, k))$.

We know that for a minimal class, δ must be recurrent, because if it contains a k -factor $\delta_{[j, j+k-1]}$ that either does not repeat, or repeats only a finite number of times, then \mathcal{G}^δ cannot be minimal, as forbidding the induced subgraph $H_{j,1}^\delta(k, k)$ would leave a proper subclass that still has unbounded clique-width. Therefore, we will only consider recurrent δ for the remainder of the paper.

4.1 The bond-graph

To study minimality we will use the following graph class. A *bond-graph* $B^\beta(Q) = (V, E)$ for finite $Q \subseteq \mathbb{N}$ has vertices $V = Q$ and edges $E = \beta_Q$.

Let $\mathcal{B}^\beta = \{B^\beta(Q) : Q \subseteq \mathbb{N} \text{ finite}\}$. Note that \mathcal{B}^β is a hereditary subclass of \mathcal{G}^δ because

- (a) if $Q' \subseteq Q$ then $B^\beta(Q')$ is also a bond-graph, and
- (b) $B^\beta(Q)$ is an induced subgraph of \mathcal{P}^δ since if $Q = \{y_1, y_2, \dots, y_n\}$ with $y_1 < y_2 < \dots < y_n$ then it can be constructed from \mathcal{P}^δ by taking one vertex from each column y_j in turn such that there is no α or γ edge to previously constructed vertices.

We define a parameter (for $n \geq 2$)

$$\mathcal{M}^\beta(n) = \sup_{m < n} \mu(B^\beta([1, n]), [1, m]).$$

The bond-graphs can be characterised as the sub-class of graphs on a single row (although missing the α -edges) with the parameter \mathcal{M}^β measuring the number of distinct neighbourhoods between intervals of a single row.

We will say that the bond-set β has *bounded* \mathcal{M}^β if there exists M such that $\mathcal{M}^\beta(n) < M$ for all $n \in \mathbb{N}$.

The following proposition will prove useful later in creating linear clique-width expressions.

Proposition 4.1. *Let $n, m, m' \in \mathbb{N}$ satisfy $m < m' < n$. Then for graph $B^\beta([1, n])$, in any partition of $[1, m]$ into $[m+1, n]$ -similar sets $\{S_i : 1 \leq i \leq k\}$ and $[1, m']$ into $[m'+1, n]$ -similar sets $\{S'_j : 1 \leq j \leq k'\}$ for every $\ell \in [1, k]$ there exists $\ell' \in [1, k']$ such that $S_\ell \subseteq S'_{\ell'}$.*

Proof. As two vertices x and y in S_ℓ have the same neighbourhood in $[m+1, n]$ then it follows they have the same neighbourhood in $[m'+1, n]$ since $m < m'$ so x and y must sit in the same $[m'+1, n]$ -similar set $S'_{\ell'}$ for some $\ell' \in [1, k']$. \square

Proposition 4.2. *For any $\delta = (\alpha, \beta, \gamma)$ and any $n \in \mathbb{N}$,*

$$\mathcal{M}^\beta(n) \leq \mathcal{N}^\delta([1, n]) + 1.$$

Proof. In the two-row graph $T^\delta([1, n])$ partition $R_1([1, n])$ into $R_2([1, n])$ -similar equivalence classes $\{W_i\}$ so that two vertices $v_{x,1}$ and $v_{y,1}$ are in the same set W_i if they have the same neighbourhood in $R_2([1, n])$. By definition the number of such sets is $\mu(T^\delta([1, n]), R_1([1, n])) = \mathcal{N}^\delta([1, n])$. For $m < n$ partition $[1, m]$ into s sets $\{P_i\}$ such that $P_i = \{j : v_{j,1} \in W_i\}$. Then s is no more than the number of sets in $\{W_i\}$ by definition, but no less than $\mu(B^\beta([1, n]), [1, m]) - 1$, the number of equivalence classes that are $[m+1, n]$ -similar (excluding, possibly, vertex m). This holds for all $m < n$, so

$$\mathcal{M}^\beta(n) - 1 = \sup_{m < n} \mu(B^\beta([1, n]), [1, m]) - 1 \leq \mu(T^\delta([1, n]), R_1([1, n])) = \mathcal{N}^\delta([1, n]).$$

\square

4.2 Veins and Slices

We will start by considering only graph classes \mathcal{G}^δ for $\delta = (\alpha, \beta, \gamma)$ in which α is an infinite word from the alphabet $\{0, 2\}$ and then extend to the case where α is an infinite word from the alphabet $\{0, 1, 2, 3\}$.

Consider a specific embedding of a graph $G = (V, E) \in \mathcal{C}$ in \mathcal{P}^δ , and recall that the induced subgraph of G on the vertices $V \cap C_{[j, j+k-1]}$ is denoted $G_{[j, j+k-1]}$.

Let α be an infinite word over the alphabet $\{0, 2\}$. A *vein* \mathcal{V} of $G_{[j, j+k-1]}$ is a set of $t \leq k$ vertices $\{v_s, \dots, v_{s+t-1}\}$ in consecutive columns such that $v_y \in V \cap C_y$ for each $y \in \{s, \dots, s+t-1\}$ and for which $v_y \sim v_{y+1}$ for all $y \in \{s, \dots, s+t-2\}$.

We will call a vein of length k a *full vein* and a vein of length $< k$ a *part vein*. Note that as α comes from the alphabet $\{0, 2\}$, for a vein $\{v_s, \dots, v_{s+t-1}\}$, v_{y+1} is no higher than v_y for each $y \in \{s, \dots, s+t-2\}$. A horizontal row of k vertices in $G_{[j, j+k-1]}$ is a full vein.

As G is $\text{Free}(H_{j,1}^\delta(k, k))$ we know that no set of vertices of G induces $H_{j,1}^\delta(k, k)$. We consider this in terms of disjoint full veins of $G_{[j, j+k-1]}$. Note that k rows of vertices between column j and column $j+k-1$ are a set of k disjoint full veins and are isomorphic to $H_{j,1}^\delta(k, k)$. There are other sets of k disjoint full veins that form a graph isomorphic to $H_{j,1}^\delta(k, k)$, but some sets of k full veins do not. Our first task is to clarify when a set of k full veins has this property.

Let $\{v_j, \dots, v_{j+k-1}\}$ be a full vein such that each vertex v_x has coordinates (x, u_x) in \mathcal{P} , observing that $u_{x+1} \leq u_x$ for $x \in [j, j+k-2]$. We construct an *upper border* to be a set of vertical coordinates $\{w_j, \dots, w_{j+k-1}\}$ using the following procedure:

- (1) Set $w_j = u_j$,
- (2) Set $x = j + 1$,
- (3) if $\alpha_{x-1} = 2$ set $w_x = u_{x-1}$,
- (4) if $\alpha_{x-1} = 0$ set $w_x = w_{x-1}$,
- (5) set $x = x + 1$,
- (6) if $x = j + k$ terminate the procedure, otherwise return to step (3).

Given a full vein $\mathcal{V} = \{v_j, \dots, v_{j+k-1}\}$, define the *fat vein* $\mathcal{V}^f = \{v_{x,y} \in V(G_{[j, j+k-1]}) : x \in [j, j+k-1], y \in [u_x, w_x]\}$ (See examples shown in Figure 7).

Let \mathcal{V}_1 and \mathcal{V}_2 be two full veins. Then we say they are *independent* if $\mathcal{V}_1^f \cap \mathcal{V}_2^f = \emptyset$ i.e. their corresponding fat veins are disjoint.

Proposition 4.3. $G_{[j, j+k-1]}$ cannot contain more than $(k-1)$ independent full veins.

Proof. We claim that k independent full veins $\{\mathcal{V}_1, \dots, \mathcal{V}_k\}$ induce the forbidden graph $H_{j,1}^\delta(k, k)$.

Remembering $v_{x,y}$ is the vertex in the grid \mathcal{P} in the x -th column and y -th row, let $w_{x,y}$ be the vertex in the y -th full vein \mathcal{V}_y in column x . We claim the mapping $\phi(w_{x,y}) \rightarrow v_{x,y}$ is an isomorphism.

Consider vertices $w_{x,y} \in \mathcal{V}_y$ and $w_{s,t} \in \mathcal{V}_t$ for $t \geq y$. Then

- (a) If $t = y$ (i.e the vertices are on the same vein) then both $w_{x,y} \sim w_{s,t}$ and $v_{x,y} \sim v_{s,t}$ if and only if $|x - s| = 1$ or $(x, s) \in \beta$,
- (b) If $t > y$ and $x = s$ then both $w_{x,y} \sim w_{s,t}$ and $v_{x,y} \sim v_{s,t}$ if and only if $\gamma_x = 1$,
- (c) If $t > y$ and $s = x + 1$ then both $w_{x,y} \not\sim w_{s,t}$ and $v_{x,y} \not\sim v_{s,t}$,
- (d) If $t > y$ and $s = x - 1$ then both $w_{x,y} \sim w_{s,t}$ and $v_{x,y} \sim v_{s,t}$ if and only if $\alpha_s = 2$,
- (e) If $t > y$ and $|s - x| > 1$ then both $w_{x,y} \sim w_{s,t}$ and $v_{x,y} \sim v_{s,t}$ if and only if $(x, s) \in \beta$.

Hence, $w_{x,y} \sim w_{s,t}$ if and only if $v_{x,y} \sim v_{s,t}$ and ϕ is an isomorphism from k independent full veins to $H_{j,1}^\delta(k, k)$. \square

4.3 Vertex colouring

Our objective is to identify conditions on (recurrent) δ that make \mathcal{G}^δ a minimal class of unbounded linear clique-width. For such a δ we must show that any graph G in a proper hereditary subclass \mathcal{C} has bounded linear clique-width. In order to do this we will partition G into manageable sections (which we will call "panels"), the divisions between the panels chosen so that they can be built separately and then 'stuck' back together again, using a linear clique-width expression requiring only a bounded number of labels. In this section we describe a vertex colouring that will lead (in Section 4.5) to the construction of these panels.

As previously observed, for any subclass \mathcal{C} there exist j and k such that $\mathcal{C} \subseteq \text{Free}(H_{j,1}^\delta(k, k))$. As δ is recurrent, if we let $\delta^* = \delta_{[j, j+k-1]}$ be the k -factor that defines the forbidden graph $H_{j,1}^\delta(k, k)$, we can find δ^* in δ infinitely often, and we will use these instances of δ^* to divide our embedded graph G into the required panels.

Firstly, we will construct a maximal set \mathbb{B} of independent full veins for $G_{[j, j+k-1]}$, a section of G that by Proposition 4.3 cannot have more than $(k - 1)$ independent full veins. We start with the lowest full vein (remembering that the rows of the grid \mathcal{P} are indexed from the bottom) and then keep adding the next lowest independent full vein until the process is exhausted.

Note that the next lowest independent full vein is unique because if we have two full veins $\mathcal{V}_1, \mathcal{V}_2$ with vertices $\{v_j, \dots, v_{j+k-1}\}$ and $\{v'_j, \dots, v'_{j+k-1}\}$ respectively then they can be combined to give $\{\min(v_j, v'_j), \dots, \min(v_{j+k-1}, v'_{j+k-1})\}$ which is a full vein with a vertex in each column at least as low as the vertices of \mathcal{V}_1 and \mathcal{V}_2 .

Let \mathbb{B} contain $b < k$ independent full veins, numbered from the bottom as $\mathcal{V}_1, \dots, \mathcal{V}_b$ such that any other full vein not in \mathbb{B} must have a vertex in common with a fat vein \mathcal{V}_y^f corresponding to one of the veins \mathcal{V}_y of \mathbb{B} .

Let $u_{x,y}$ be the lowest vertical coordinate and $w_{x,y}$ the highest vertical coordinate of vertices in $\mathcal{V}_y^f \cap C_x$. We define $\mathcal{S}_0 = \{v_{x,y} \in V(G_{[j,j+k-1]}) : x \in [j, j+k-1], y < u_{x,1}\}$, $\mathcal{S}_b = \{v_{x,y} \in V(G_{[j,j+k-1]}) : x \in [j, j+k-1], y > w_{x,b}\}$, and for $y = 1, \dots, b-1$ we define:

$$\mathcal{S}_i = \{v_{x,y} \in V(G_{[j,j+k-1]}) : x \in [j, j+k-1], w_{x,i} < y < u_{x,i+1}\}$$

This gives us $b+1$ slices $\{\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_b\}$.

We partition the vertices in the fat veins and the slices into sets which have similar neighbourhoods, which will facilitate the division of G into panels. We colour the vertices of $G_{[j,j+k-1]}$ so that each slice has green/pink vertices to the left and red vertices to the right of the partition, and each fat vein has blue vertices (if any) to the left and yellow vertices to the right. Examples of vertex colourings are shown in Figure 7.

Colour the vertices of each slice \mathcal{S}_i as follows:

- Colour any vertices in the left-hand column green. Now colour green any remaining vertices in the slice that are connected to one of the green left-hand column vertices by a part vein that does not have a vertex in common with any of the fat veins corresponding to the full veins in \mathbb{B} .
- Locate the column t of the right-most green vertex in the slice. If there are no green vertices set $t = s = j$. If $t > j$ then choose s in the range $j \leq s < t$ such that s is the highest column index for which $\alpha_s = 2$. If there are no columns before t for which $\alpha_s = 2$ then set $s = j$. Colour pink any vertices in the slice (not already coloured) in columns j to s which are below a vertex already coloured green.
- Colour any remaining vertices in the slice red.

Note that no vertex in the right-hand column can be green because if there was such a vertex then this would contradict the fact that there can be no full veins other than those which have a vertex in common with one of the fat veins corresponding to the full veins in \mathbb{B} . Furthermore, no vertex in the right hand column can be pink as this would contradict the fact that every pink vertex must lie below a green vertex in the same slice.

Colour the vertices of each fat vein \mathcal{V}_i^f as follows:

- Let s be the column as defined above for the slice immediately above the fat vein. If $s = j$ colour the whole fat vein yellow. If $s > j$ colour vertices of the fat vein in columns j to s blue and the rest of the vertices in the fat vein yellow.

When we create a clique-width expression we will be particularly interested in the edges between the blue and green/pink vertices to the left and the red and yellow vertices to the right.




Figure 7: Examples of vein and slice colouring – a 222222, a 222000 and a 222000022 factor, with vertices coloured blue, green, pink, red and yellow as described. The only edges shown are the veins (bold blue), other edges in the fat veins (blue), part veins that start on the left column but do not reach the right column (green) and related pink rows.

Proposition 4.4. *Let v be a red vertex in column x and slice \mathcal{S}_i .*

If u is a blue, green or pink vertex in column $x - 1$ then

$$uv \in E(G) \text{ if and only if } \alpha_{x-1} = 2 \text{ and } u \in \mathcal{V}_{i+1}^f \cup \mathcal{S}_{i+1} \cup \dots \cup \mathcal{V}_b^f \cup \mathcal{S}_b.$$

Similarly, if u is a blue, green or pink vertex in column $x + 1$ then

$$uv \in E(G) \text{ if and only if } \alpha_x = 2 \text{ and } u \in \mathcal{S}_0 \cup \mathcal{V}_1^f \cup \mathcal{S}_1 \cup \dots \cup \mathcal{V}_i^f \cup \mathcal{S}_i.$$

Proof. Note that as u and v are in consecutive columns we need only consider α -edges.

If u is green in column $x - 1$ of \mathcal{S}_i then red v in column x of \mathcal{S}_i cannot be adjacent to u as this would place red v on a green part-vein which is a contradiction. Likewise, if u is green in column $x + 1$ of \mathcal{S}_i then red v in column x of \mathcal{S}_i must be adjacent to u since if it was not adjacent to such a green vertex in the same slice then this implies the existence of a green vertex above the red vertex in the same column which contradicts the colouring rule to colour pink any vertex in columns j to s below a vertex coloured green.

The other adjacencies are straightforward. □

Proposition 4.5. *Let v be a yellow vertex in column x and fat vein \mathcal{V}_i^f .*

If u is a blue, green or pink vertex in column $x - 1$ then

$$uv \in E(G) \text{ if and only if } \alpha_{x-1} = 2 \text{ and } u \in \mathcal{V}_i^f \cup \mathcal{S}_i \cup \dots \cup \mathcal{V}_b^f \cup \mathcal{S}_b.$$

Similarly, if u is a blue, green or pink vertex in column $x + 1$ then

$$uv \in E(G) \text{ if and only if } \alpha_x = 2 \text{ and } u \in \mathcal{S}_0 \cup \mathcal{V}_1^f \cup \mathcal{S}_1 \cup \dots \cup \mathcal{V}_{i-1}^f \cup \mathcal{S}_{i-1}.$$

Proof. Note that as u and v are in consecutive columns we need only consider α -edges.

If u is blue in column $x - 1$ of \mathcal{V}_i^f then yellow v in column x of \mathcal{V}_i^f must be adjacent to u from the definition of a fat vein. Equally, from the colouring definition for a fat vein there cannot be a blue vertex in column $x + 1$ of \mathcal{V}_i^f if there is a yellow vertex in column x of \mathcal{V}_i^f .

The other adjacencies are straightforward. □

Having established these propositions, as the pink and green vertices in a particular slice and column have the same adjacencies to the red and yellow vertices, we now combine the green and pink sets and simply refer to them all as *green*.

4.4 Extending α to the 4-letter alphabet

Our analysis so far has been based on α being a word from the alphabet $\{0, 2\}$. We now use the following lemma to extend our colouring to the case where α is a word over the 4-letter alphabet $\{0, 1, 2, 3\}$.

Let α be an infinite word over the alphabet $\{0, 1, 2, 3\}$ and α^+ be the infinite word over the alphabet $\{0, 2\}$ such that for each $x \in \mathbb{N}$,

$$\alpha_x^+ = \begin{cases} 0 & \text{if } \alpha_x = 0 \text{ or } 1, \\ 2 & \text{if } \alpha_x = 2 \text{ or } 3, \end{cases}$$

Denoting $\delta = (\alpha, \beta, \gamma)$ and $\delta^+ = (\alpha^+, \beta, \gamma)$, let $G = (V, E)$ be a graph in the class \mathcal{G}^δ with a particular embedding in the vertex grid $V(\mathcal{P})$. We will refer to $G^+ = (V, E^+)$ as the graph with the same vertex set V as G from the class \mathcal{G}^{δ^+} .

Lemma 4.6. *For any subset of vertices $U \subseteq V$, 2 vertices of U in the same column of $V(\mathcal{P})$ are $V \setminus U$ -similar in G if and only if they are $V \setminus U$ -similar in G^+ .*

Proof. Let u_1 and u_2 be two vertices in U in the same column x and v be a vertex of $V \setminus U$ in column y . If $x = y$ then v is in the same column as u_1 and u_2 and is either adjacent to both or neither depending on whether there is a γ -clique on column x , which is the same in both G and G^+ . If $|x - y| > 1$ then v is adjacent to both u_1 and u_2 if and only if there is a bond (x, y) in β , which is the same in both G and G^+ .

If $y = x + 1$ then the adjacency of v to u_1 and u_2 is determined by α_x in G and α_x^+ in G^+ . If $\alpha_x = \alpha_x^+$ (i.e. both 0 or both 2) then the adjacencies are the same in G and G^+ . If $\alpha_x = 1$ and $\alpha_x^+ = 0$, then u_1 and u_2 are both adjacent to v in G if and only if they are both non-adjacent to v in G^+ . If $\alpha_x = 3$ and $\alpha_x^+ = 2$, then u_1 and u_2 are both adjacent to v in G if and only if they are both non-adjacent to v in G^+ .

Hence u_1 and u_2 have the same neighbourhood in $V \setminus U$ in G if and only if they have the same neighbourhood in $V \setminus U$ in G^+ . \square

4.5 Panel construction

We now construct the panels of G based on our embedding of G in \mathcal{P}^δ .

To recap, $\delta^* = \delta_{[j, j+k-1]}$ is the k -factor that defines the forbidden graph $H_{j,1}^\delta(k, k)$ and we will use the repeated instances of δ^* to divide our embedded graph G into panels.

Define $t_0, t_1, \dots, t_\omega$ where t_0 is the index of the column before the first column of the embedding of G , t_ω is the index of the last column of the embedding of G and t_i ($0 < i < \omega$) represents the rightmost letter index of the i -th copy of δ^* in δ , such that $t_i > k + t_{i-1}$ to ensure the copies are disjoint. Hence, the i -th copy of δ^* in δ corresponds to columns $C_{[t_i-k+1, t_i]}$ of \mathcal{P}^δ and we denote the induced graph on these columns $G_i^* = G_{[t_i-k+1, t_i]}$.

Colour the vertices of G_i^* blue, yellow, green or red as described in Section 4.3. Call these vertex sets U_i^b, U_i^y, U_i^g and U_i^r respectively. Denote U_1^w as the vertices in $G_{[t_0+1, t_1-k]}$, and for $1 < i < \omega$ denote U_i^w the set of vertices in $G_{[t_{i-1}+1, t_i-k]}$ and colour the vertices in each U_i^w white.

We now create a sequence of *panels*, the first panel is $P_1 = U_1^w \cup U_1^g \cup U_1^b$, and subsequent panels given by

$$P_i = U_{i-1}^y \cup U_{i-1}^r \cup U_i^w \cup U_i^g \cup U_i^b.$$

These panels create a disjoint partition of the vertices of our embedding of G .

The following lemma will be used to put a bound on the number of labels required in a linear clique-width expression to create edges between panels. We denote $\mathbb{P}_i = \cup_{s=1}^i P_s$, and $\mathcal{V}_{i,x}^f$ and $\mathcal{S}_{i,x}$ as the x -th fat vein and x -th slice respectively of G_i^* .

Lemma 4.7. *Let (α, β, γ) be a recurrent δ -triple where α is an infinite word over the alphabet $\{0, 1, 2, 3\}$, γ is an infinite binary word and β is a bond set which has bounded \mathcal{M}^β , so that $\mathcal{M}^\beta(n) < M$ for all $n \in \mathbb{N}$.*

Then for any graph $G = (V, E) \in \mathcal{G}^\delta \cap \text{Free}(H_{j,1}^\delta(k, k))$ for some $j, k \in \mathbb{N}$ with vertices V partitioned into panels $\{P_1, \dots, P_\omega\}$ and $1 \leq i \leq \omega$,

$$\mu(G, \mathbb{P}_i) < M + 2k^2.$$

Proof. Using Lemma 4.6 with $U = \mathbb{P}_i$, we only need to prove the case when $\alpha = \alpha^+$, a word over the alphabet $\{0, 2\}$.

Considering the three sets of vertices $\mathbb{P}_i \setminus \{U_i^b \cup U_i^g\}$, U_i^b and U_i^g separately, we have:

- (a) $\mu(G, \mathbb{P}_i \setminus \{U_i^b \cup U_i^g\}) \leq \mathcal{M}^\beta(t_\omega) < M$.
- (b) $\mu(G, U_i^b) \leq k(k-1)$, noticing that from Propositions 4.4 and 4.5 two blue vertices in the same fat vein and column have the same neighbourhood in $V \setminus \mathbb{P}_i$.
- (c) $\mu(G, U_i^g) \leq k(k+1)$, noticing that from Propositions 4.4 and 4.5 two green vertices in the same slice and column have the same neighbourhood in $V \setminus \mathbb{P}_i$.

This covers all vertices of \mathbb{P}_i and the result follows. □

4.6 When \mathcal{G}^δ is a minimal class of unbounded clique-width

Our strategy for proving that an arbitrary graph G in a proper hereditary subclass of \mathcal{G}^δ has bounded linear clique-width (and hence bounded clique-width) is to define an algorithm to create a linear clique-width expression that allows us to recycle labels so that we can put a bound on the total number of labels required, however many vertices there are in G . We do this by constructing a linear clique-width expression for each panel P_i in G in a linear sequence,

leaving the labels on each vertex of previously constructed panels \mathbb{P}_{i-1} with an appropriate label to allow edges to be constructed between the current panel/vertex and previous panels. To be able to achieve this we require the following ingredients:

- (a) δ to be recurrent so we can create the panels,
- (b) a bound on the number of labels required to create each new panel,
- (c) a process of relabelling so that we can leave appropriate labels on each vertex of the current panel to enable connecting to previous panels, before moving on to the next panel, and
- (d) a bound on the number of labels required to create edges to previously constructed panels.

We have (a) by assumption and we will deal with (c) and (d) in the proof of Theorem 4.10. The next two lemmas show how we can restrict δ further, using a new concept of ‘gap factors’, to ensure (b) is achieved.

Lemma 4.8. *For any δ and graph $G \in \mathcal{G}^\delta$ and any $j_1, j_2 \in \mathbb{N}$ where $|j_2 - j_1| = \ell - 1$*

$$\text{lcw}(G_{[j_1, j_2]}) \leq 2\ell$$

Proof. We construct $G_{[j_1, j_2]}$ using a row-by-row linear method, starting in the bottom left. For each of the ℓ columns, we create 2 labels: one label c_1, \dots, c_ℓ for the vertex in the *current* row being constructed, and one label e_1, \dots, e_ℓ for the vertices in all *earlier* rows.

For the first row, we insert the (max) ℓ vertices using the labels c_1, \dots, c_ℓ , and since every vertex has its own label we can insert all necessary edges. Now relabel $c_i \rightarrow e_i$ for each i .

Suppose that the first r rows have been constructed, in such a way that every existing vertex in column i has label e_i . We insert the (max) ℓ vertices in row $r+1$ using labels c_1, \dots, c_ℓ . As before, every vertex in this row has its own label, so we can insert all edges between vertices within this row. Next, note that any vertex in this row has the same relationship with all vertices in rows $1, \dots, r$ of any column i . Since these vertices all have label e_i and the vertex in row $r+1$ has its own label, we can add edges as determined by α , β and γ as necessary. Finally, relabel $c_i \rightarrow e_i$ for each i , move to the next row and repeat until all rows have been constructed. \square

As a consequence of Lemma 4.8, if δ is almost periodic, so that for any factor δ^* of δ every factor of δ of length at least $\mathcal{L}(\delta^*)$ contains δ^* , then each panel P_i covers a maximum of $\mathcal{L}(\delta^*) + k$ columns and $G[P_i]$ can be constructed with at most $2\mathcal{L}(\delta^*) + 2k$ labels. So (b) is satisfied for δ almost periodic, but we wish to go further.

We will call the factors between, and including, any consecutive pair of occurrences of a k -factor δ^* the δ^* -gap factors. If δ is recurrent but not almost periodic then there exists a k -factor $\delta^* = \delta_{[j, j+k-1]}$, say, such that the δ^* -gap factors can be arbitrarily long and Lemma 4.8 will not be sufficient to demonstrate a bound on the number of labels required to construct P_i .

Using the same column indices $\{t_i\}$ used for panel construction in Section 4.5, let the i -th δ^* -gap factor be denoted δ_{q_i} where $q_1 = [t_0 + 1, t_1]$ and $q_i = [t_{i-1} - k + 1, t_i]$ for $1 < i < \omega$. Note that for every i , $P_i \subseteq C_{q_i}$. We define a two-row graph $T^\delta(q_i)$ (see Section 3.1) corresponding to each q_i , with vertex partition $R_1(q_i) \cup R_2(q_i)$ and with related parameter $\mathcal{N}^\delta(q_i) = \mu(T^\delta(q_i), R_1(q_i))$.

We will define $\mathcal{N}_*^\delta(n) = \sup_{1 \leq i \leq n} \mathcal{N}^\delta(q_i)$ and will say \mathcal{N}_*^δ is bounded if $\mathcal{N}_*^\delta(n)$ is bounded as $n \rightarrow \infty$.

Lemma 4.9. *Let δ be a recurrent triple with k -factor $\delta^* = \delta_{[j, j+k-1]}$ such that \mathcal{N}_*^δ is bounded. Then there exists a bound on the linear clique-width of any graph $G \in \mathcal{G}^\delta$, with $V[G] \subseteq C_{q_i}$ for some $i \in \mathbb{N}$.*

Proof. As \mathcal{N}_*^δ is bounded then there exists fixed N such that $\mathcal{N}^\delta(q_i) < N$ for any $i \in \mathbb{N}$. Following the proof of Lemma 3.15 it follows that there is a bound, say J , on the number of 2s and 3s in any δ^* -gap factor.

We can use the row-by-row linear method from the proof of Lemma 3.15 to show that for any graph $G \in \mathcal{G}^\delta$, with $V[G] \subseteq C_{q_i}$ for some $i \in \mathbb{N}$ it follows that $\text{lcw}(G) \leq 2J + N + 2$. \square

An example of a recurrent but *not* almost periodic triple $\delta = (\alpha, \emptyset, 0^\infty)$ for which \mathcal{N}^δ is unbounded, but where for any factor δ^* of δ the parameter \mathcal{N}_*^δ is bounded, was given in [3], where α is a recurrent but not almost periodic binary word constructed by a process of substitution.

We are now in a position to define a set of hereditary graph classes \mathcal{G}^δ that are minimal of unbounded clique-width. We will denote $\Delta_{\min} \subseteq \Delta$ as the set of all δ -triples in Δ with the characteristics:

- (a) δ is recurrent,
- (b) for any factor δ^* of δ , the parameter \mathcal{N}_*^δ is bounded (includes all almost periodic δ in Δ), and
- (c) the bond set β has bounded \mathcal{M}^β .

Theorem 4.10. *If $\delta \in \Delta_{\min}$ then \mathcal{G}^δ is a minimal hereditary class of both unbounded linear clique-width and unbounded clique-width.*

Proof. To prove this Theorem we must show that if $\delta \in \Delta_{\min}$ then every proper hereditary subclass $\mathcal{C} \subsetneq \mathcal{G}^\delta$ has bounded linear clique-width. From the introduction to this section we know that for such a subclass \mathcal{C} there must exist some $H_{j,1}^\delta(k, k)$ for some j and $k \in \mathbb{N}$ such that $\mathcal{C} \subseteq \text{Free}(H_{j,1}^\delta(k, k))$.

From Lemma 4.9 we know there exist fixed J and $N \in \mathbb{N}$ such that the number of labels required to construct each panel P_i by the row-by-row linear method for all $i \in \mathbb{N}$ is no more than $2J + N + 2$.

As the bond-set β has bounded \mathcal{M}^β , let $M \in \mathbb{N}$ be a constant such that $\mathcal{M}^\beta(n) < M$ for all $n \in \mathbb{N}$.

Although a single panel P_i can be constructed using at most $2J + N + 2$ labels, we need to be able to recycle labels so that we can construct any number of panels with a bounded number of labels. We will show that any graph $G \in \text{Free}(H_{j,1}^\delta(k, k))$ can be constructed by a linear clique-width expression that only requires a number of labels determined by the constants M, N, J and k .

For our construction of panel P_i , we will use the following set of $4k^2 + MN + M + 2J + 2$ labels:

- 2 *current vertex labels*: a_1 and a_2 ;
- J *current row labels*: $\{c_y : y = 1, \dots, J\}$ for first J columns;
- J *previous row labels*: $\{p_y : y = 1, \dots, J\}$ for first J columns;
- MN *partition labels*: $\{s_{x,y} : x = 1, \dots, M, y = 1, \dots, N\}$, for vertices in $U_{i-1}^y \cup U_{i-1}^r \cup U_i^w$;
- k^2 *blue current panel labels*: $\{bc_{x,y} : x = 1, \dots, k, y = 1, \dots, k\}$, for vertices $\mathcal{V}_{i,x}^f \cap U_i^b \cap C_y$;
- k^2 *blue previous panel labels*: $\{bp_{x,y} : x = 1, \dots, k, y = 1, \dots, k\}$, for vertices $\mathcal{V}_{i-1,x}^f \cap U_{i-1}^b \cap C_y$;
- k^2 *green current panel labels*: $\{gc_{x,y} : x = 0, \dots, k-1, y = 1, \dots, k\}$, for vertices $\mathcal{S}_{i,x} \cap U_i^g \cap C_y$;
- k^2 *green previous panel labels*: $\{gp_{x,y} : x = 0, \dots, k-1, y = 1, \dots, k\}$, for vertices $\mathcal{S}_{i-1,x} \cap U_{i-1}^g \cap C_y$;
- M *bond labels*: $\{m_y : y = 1, \dots, M\}$, for vertices in previous panels for creating the β -bond edges between columns.

We carry out the following iterative process to construct each panel P_i in turn.

Assume $\mathbb{P}_{i-1} = \cup_{s=1}^{i-1} P_s$ has already been constructed such that labels $m_y, bp_{x,y}$ and $gp_{x,y}$ have been assigned to the $M + 2k^2 V \setminus \mathbb{P}_{i-1}$ -similar sets as described in Lemma 4.7.

Using the same column indices $\{t_i\}$ used for panel construction (Section 4.5) we assign a default partition label $s_{x,y}$ to each column of $U_{i-1}^y \cup U_{i-1}^r \cup U_i^w$ as follows:

- (a) Consider the bond-graph $B^\beta([1, t_\omega])$ (Section 4.1). We partition the interval $Q = [t_{i-1} - k + 1, t_i - k]$ into $[t_i - k + 1, t_\omega]$ -similar sets of which there are at most M , and use label index x to identify values in Q in the same $[t_i - k + 1, t_\omega]$ -similar set. Consequently, vertices in two columns of $U_{i-1}^y \cup U_{i-1}^r \cup U_i^w$ that have the same default label x value have the same neighbourhood in $G_{[t_i - k + 1, t_\omega]}$ and hence are in the same $V \setminus \mathbb{P}_i$ -similar set.
- (b) Consider the two-row graph $T^\delta(Q)$ (Section 3.1). We partition vertices in $R_1(Q)$ into $R_2(Q)$ -similar sets of which there are at most N . We create a corresponding partition of the interval Q such that $v_{x,1}$ and $v_{y,1}$ are in the same equivalence class of $R_1(Q)$ if and only if x and y are in the same partition set of Q . We now use label index y to identify values in the same partition set. Consequently, vertices in two columns of $U_{i-1}^y \cup U_{i-1}^r \cup U_i^w$ that have the same default label y value have the same neighbourhood within G_Q .

We construct each panel P_i in the row-by-row linear method used for the graph with a finite number of 2s and 3s with bounded \mathcal{N}^δ constructed in Lemma 3.15. The current vertex always has a unique label. Thus, for each row, we use labels c_1, \dots, c_J for vertices in the first J columns and then alternate a_1 and a_2 for the current and previous vertices for the remainder of the row.

For each new vertex in the current row we add edges as follows:

- (a) Insert required edges to the $\mathcal{M}^\beta + 2k^2 V \setminus \mathbb{P}_{i-1}$ -similar sets – see Lemma 4.7. This is possible because vertices within each of these sets are either all adjacent to the current vertex or none of them are.
- (b) Insert required edges to vertices in the same or lower rows in the current panel. This is possible as these vertices all have labels $p_y, s_{x,y}, bc_{x,y}$ or $gc_{x,y}$ and, from the construction, vertices with the same y value are either all adjacent to the current vertex or none of them are.

Following completion of edges to the current vertex, we relabel the previous vertex as follows:

- from c_y to p_y if it is in the first J columns,
- from a_2 (or a_1) to its default partition label $s_{x,y}$ if it is in $U_{i-1}^y \cup U_{i-1}^r \cup U_i^w$ but not in the first J columns.
- from a_2 (or a_1) to $bc_{x,y}$ if it is in $\mathcal{V}_{i,x}^f \cap U_i^b$, and
- from a_2 (or a_1) to $gc_{x,y}$ if it is in $\mathcal{S}_{i,x} \cap U_i^g$.

We now repeat for the next row of panel P_i .

Once panel P_i is complete, relabel as follows:

Relabel vertices in accordance with their $V \setminus \mathbb{P}_i$ -similar set, of which there are at most M . Note from Proposition 4.1, that two vertices with the same label m_y from the previous \mathbb{P}_{i-1} partition sets will still need the same label in \mathbb{P}_i . Two equivalence classes from the \mathbb{P}_{i-1} partition may merge to form a new equivalence class in the \mathbb{P}_i partition. Hence, it is possible to relabel with the same label the old equivalence classes that merge, and then use the spare m_y labels for any new equivalence classes that appear. We never need more than M such labels.

Also relabel all vertices with labels $bp_{x,y}, gp_{x,y}, p_y$ and $s_{x,y}$ with the relevant bond label m_y of their $V \setminus \mathbb{P}_i$ -similar set. This is possible for the vertices labelled $s_{x,y}$ as the index x signifies their $V \setminus \mathbb{P}_i$ -similar set.

Now relabel $bc_{x,y} \rightarrow bp_{x,y}$ and $gc_{x,y} \rightarrow gp_{x,y}$ ready for the next panel. For the next panel we can reuse labels $a_1, a_2, c_y, p_y, s_{x,y}, bc_{x,y}$ and $gc_{x,y}$ as necessary.

This process repeated for all panels completes the construction of G .

The maximum number of labels required to construct any graph $G \in \text{Free}(H_{j,1}^\delta(k, k))$ is $4k^2 + MN + M + 2J + 2$ and hence \mathcal{C} has bounded linear clique-width. \square

The conditions for δ to be in Δ_{\min} are sufficient for the class \mathcal{G}^δ to be minimal. It is fairly easy to see that the first two conditions (recurrence and bounded clique-width for gap factors) are also necessary. However, there remains a question regarding the necessity of the bond set β to have bounded \mathcal{M}^β . We have been unable to identify any $\delta \notin \Delta_{\min}$ such that \mathcal{G}^δ is a minimal class of unbounded clique-width, hence:

Conjecture 4.11. *The hereditary graph class \mathcal{G}^δ is minimal of unbounded clique-width if and only if $\delta \in \Delta_{\min}$.*

5 Examples of new minimal classes

It has already been shown in [3] that there are uncountably many minimal hereditary classes of graphs of unbounded clique-width. However, armed with the new framework we can now identify many other types of minimal classes. Some examples of $\delta = (\alpha, \beta, \gamma)$ values that yield a minimal class are shown in Table 2.

Example	α	β ($x, y \in \mathbb{N}$)	γ	\mathcal{M}^β bound
1.	0^∞	\emptyset	1^∞	1
2.	1^∞	$(1, x + 2)$	0^∞	2
3.	$(23)^\infty$	$(x, x + 2)$	0^∞	3
4.	0^∞	$(x, y) : x - y \neq 1, x - y \equiv 1 \pmod{2}$	0^∞	3
5.	1^∞	$(x, y) : x \neq y, x - y \equiv 0 \pmod{2}$	1^∞	2
6.	2^∞	$(x, y) : 1 < x - y \leq n$ (fixed n)	0^∞	n

Table 2: New minimal hereditary graph classes of unbounded clique-width

6 Concluding remarks

The ideas of periodicity and recurrence are well established concepts when applied to symbolic sequences (i.e. words). Application to δ -triples and in particular β -bonds is rather different and needs further investigation.

The β -bonds have been defined as generally as possible, allowing a bond between any two non-consecutive columns. The purpose of this was to capture as many minimal classes in the framework as possible. However, it may be observed that the definition is so general that for any finite graph G it is possible to define β so that G is isomorphic to an induced subgraph of $B^\beta(Q)$ and hence \mathcal{G}^δ .

In these \mathcal{G}^δ graph classes we have seen that unboundedness of clique-width is determined by the unboundedness of a parameter measuring the number of distinct neighbourhoods between

two-rows. The minimal classes are those which satisfy defined recurrence characteristics and for which there is a bound on a parameter measuring the number of distinct neighbourhoods between vertices in one row.

Hence, whilst we have created a framework for many types of minimal classes, there may be further classes 'hidden' in the β -bonds. Indeed, we believe other types of minimal hereditary classes of unbounded clique-width exist and this is still an open area for research.

References

- [1] B. Alecu, M. M. Kanté, V. Lozin, and V. Zamaraev. Between clique-width and linear clique-width of bipartite graphs. *Discrete Math.*, 343(8):111926, 14, 2020.
- [2] A. Atminas, R. Brignall, V. Lozin, and J. Stacho. Minimal classes of graphs of unbounded clique-width defined by finitely many forbidden induced subgraphs. *Discrete Appl. Math.*, 295:57–69, 2021.
- [3] R. Brignall and D. Cocks. Uncountably many minimal hereditary classes of graphs of unbounded clique-width. *Electron. J. Combin.*, 29(1):Paper No. 1.63, 27, 2022.
- [4] R. Brignall, N. Korpelainen, and V. Vatter. Linear clique-width for hereditary classes of cographs. *J. Graph Theory*, 84(4):501–511, 2017.
- [5] A. Collins, J. Foniok, N. Korpelainen, V. Lozin, and V. Zamaraev. Infinitely many minimal classes of graphs of unbounded clique-width. *Discrete Appl. Math.*, 248:145–152, 2018.
- [6] B. Courcelle, J. Engelfriet, and G. Rozenberg. Handle-rewriting hypergraph grammars. *J. Comput. System Sci.*, 46(2):218–270, 1993.
- [7] B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.
- [8] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, fifth edition, 2017.
- [9] J. Geelen, O-J. Kwon, R. McCarty, and P. Wollan. The grid theorem for vertex-minors. *J. Combin. Theory Ser. B*, 2020.
- [10] V. Lozin. Minimal classes of graphs of unbounded clique-width. *Ann. Comb.*, 15(4):707–722, 2011.
- [11] F. P. Ramsey. On a Problem of Formal Logic. *Proc. London Math. Soc. (2)*, 30(4):264–286, 1929.
- [12] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Combin. Theory Ser. B*, 41(1):92–114, 1986.