Saturated simple and 2-simple topological graphs with few edges

Péter Hajnal[1]   *Alexander Igamberdiev*
Günter Rote[2]   *André Schulz*[3]

[1]Bolyai Institute, University of Szeged.
[2]Freie Universität Berlin.
[3]LG Theoretische Informatik, FernUniversität in Hagen, Hagen, Germany.

Abstract

A *simple topological graph* is a topological graph in which any two edges have at most one common point, which is either their common endpoint or a proper crossing. More generally, in a *k-simple topological graph*, every pair of edges has at most k common points of this kind. We construct *saturated* simple and 2-simple graphs with few edges. These are k-simple graphs in which no further edge can be added. We improve the previous upper bounds of Kynčl, Pach, Radoičić, and Tóth [Comput. Geom., **48**, 2015] and show that there are saturated simple graphs on n vertices with only $7n$ edges and saturated 2-simple graphs on n vertices with 14.5n edges. As a consequence, there is a k-simple graph (for a general k), which can be saturated using 14.5n edges, while previous upper bounds suggested 17.5n edges. We also construct saturated simple and 2-simple graphs that have some vertices with low degree.

1   Introduction

Let $G = (V, E)$ be a finite simple graph. A *drawing* of G is a map $\delta \colon V \cup E \to \mathbb{R}^2$ that is one-to-one on $\delta|_V \colon V \to \mathbb{R}^2$, i.e., δ assigns the vertices of the graph to different points of the plane. Furthermore, we require that $\delta|_E \colon E \to \mathcal{C}$, where $\mathcal{C}$ is a set of "nice" non-self-intersecting curves with two boundary points of the plane. For example we might think of $\mathcal{C}$ as the set of all Jordan curves or, more elementary, of the set of all simple polygonal curves. For simplicity, we will not distinguish between an edge and the curve on which it is embedded, and between a vertex and the point on which it is embedded. We assume that for any $e = xy \in E$ the edge $\delta(e)$ is a curve connecting $\delta(x)$ and $\delta(y)$ and it doesn't go through any other vertex, and also that any two different edges meet at finitely many points and any meeting point—that is not a common endvertex—is a proper crossing of the two curves.

The pair (G, δ), i.e., a graph with a drawing, is called a *topological graph*. A topological graph (G, δ) is *simple* if in δ two edges have at most one common point, which can either be a common endpoint or a proper crossing. More generally, the topological graph is called *k-simple* if in δ two edges have at most k common points. For both simple and k-simple graphs we do not allow self-intersecting edges. A topological graph is a *geometric graph* if all its edges are drawn as straight-line segments. Obviously, every geometric graph is simple, provided that the vertices are placed in general position. Thus, every graph has simple drawings. Simple graphs harness the complexity of topological graphs to some degree. From this perspective simple drawings can be considered as more readable compared to general "topological drawings".

For a graph property $\mathcal{T}$, a graph G is $\mathcal{T}$-*saturated* if G has property $\mathcal{T}$, but the addition of any edge joining two non-adjacent vertices of G violates property $\mathcal{T}$. Often structures with property $\mathcal{T}$ are quite hard to grasp, but $\mathcal{T}$-saturated structures might have a more useful character. We direct the interested reader to applications of the saturation technique [1, 3, 5]. This notion can be naturally extended to hypergraphs. A thorough survey by Faudree, Faudree, and Schmitt [2] discusses the case when property $\mathcal{T}$ is "not having F as a sub(hyper)graph".

In this paper we study *saturated k-simple* topological graphs. These are topological graphs that are k-simple, but no edge can be added without violating the k-simplicity of the drawing. Saturated planar drawings are triangulations and have therefore due to Euler's formula $3n - 6$ edges. Recently, Kynčl, Pach, Radoičić, and G. Tóth [4] started to investigate saturated k-simple graphs. The maximum number of edges a saturated simple topological graph can have is clearly $\binom{n}{2}$, since the geometric graph of K_n with vertices in general position is a simple drawing. The more intriguing questions ask about the minimum number of edges for saturated k-simple topological graph. One of the main results of Kynčl et al. [4] is a construction of sparse saturated simple and k-simple topological graphs. We denote by $s_k(n)$ the minimum number of edges a saturated k-simple graph with n vertices can have. Their upper bound on $s_k(n)$ is a linear function of n, for n being the number of vertices; see Table 1 for the

k	1	2	3	4	5	6,8,10	7	9, ≥ 11
old upper bounds [4]	$17.5n$	$16n$	$14.5n$	$13.5n$	$13n$	$9.5n$	$10n$	$7n$
new upper bounds	**7n**	**14.5n**						

Table 1: Old and new upper bounds for $s_k(n)$, the minimum number of edges in a saturated k-simple graph with n vertices.

bounds obtained by Kynčl et al. [4]. The gap between the best known upper and lower bounds for $s_k(n)$ is quite substantial. We only know that $s_1(n) \geq 1.5n$ and that $s_k(n) \geq n$ [4].

Our contribution. We improve the upper bounds for $s_k(n)$ for $k = 1, 2$. We do this by showing that for any positive integer n there exists a saturated simple topological graph with at most $7n$ edges (in Sect. 2), and a saturated 2-simple graph with at most $14.5n$ edges (in Sect. 3). Sections 2 and 3 are independent. This result also implies that there are saturated k-simple graphs with at most $14.5n$ edges for every k. See also Table 1 for a comparison with the old bounds. Our proofs are constructive, i.e., we can explicitly present the sparse saturated graphs.

We complete our results by studying *local saturation* of topological graphs. Here, local saturation refers to drawings in which one (or several) vertices have a small vertex degree even though the full drawing might not be the sparsest. Such observations might be helpful in further studies, e.g., if we want to investigate techniques for proving lower bounds that are based on the minimum vertex degree in saturated graphs. We show that there are arbitrarily large saturated simple graphs that have a vertex of degree 4, and saturated simple graphs in which 10 percent of the vertices have degree 5. For saturated 2-simple graphs we can prove that there are drawings with minimum degree 12. The current lower bounds for $s_k(n)$ are obtained by bounding the minimum vertex degree in saturated k-simple graphs [4]. Our results show the limits of this approach. These results can be found in Sect. 4.

2 Saturated simple topological graphs with few edges

In this section we give a construction that generates sparse saturated simple graphs. We start with defining a graph G, parametrized by an integer ℓ, with $n = 6\ell$ vertices and $9\ell - 6$ edges. This graph is the backbone of our sparse saturated graph.

The drawing is best visualized on the surface of a long circular cylinder. Fig. 1 shows an unrolling of the cylinder into the plane. The cylinder is obtained by cutting the drawing along the two dotted lines and gluing the top and the bottom together. The vertices of the graph are placed in a $3 \times 2\ell$-grid-like

fashion. We draw the vertices together in pairs, with each vertex X_i^L on the *left* and the corresponding vertex X_i^R on the *right*, for $X = A, B, C$ and $i = 1, \ldots, \ell$. We refer to the vertices whose label have the subscript i as the *i-th layer*. G is the union of

- three vertex-disjoint paths of *blue edges* connecting $A_1^L A_2^L \ldots A_\ell^L, B_1^L B_2^L \ldots B_\ell^L$, and $C_1^L C_2^L \ldots C_\ell^L$,

- three vertex-disjoint paths of *red edges* connecting $A_1^R A_2^R \ldots A_\ell^R$, $B_1^R B_2^R \ldots B_\ell^R$, and $C_1^R C_2^R \ldots C_\ell^R$, and

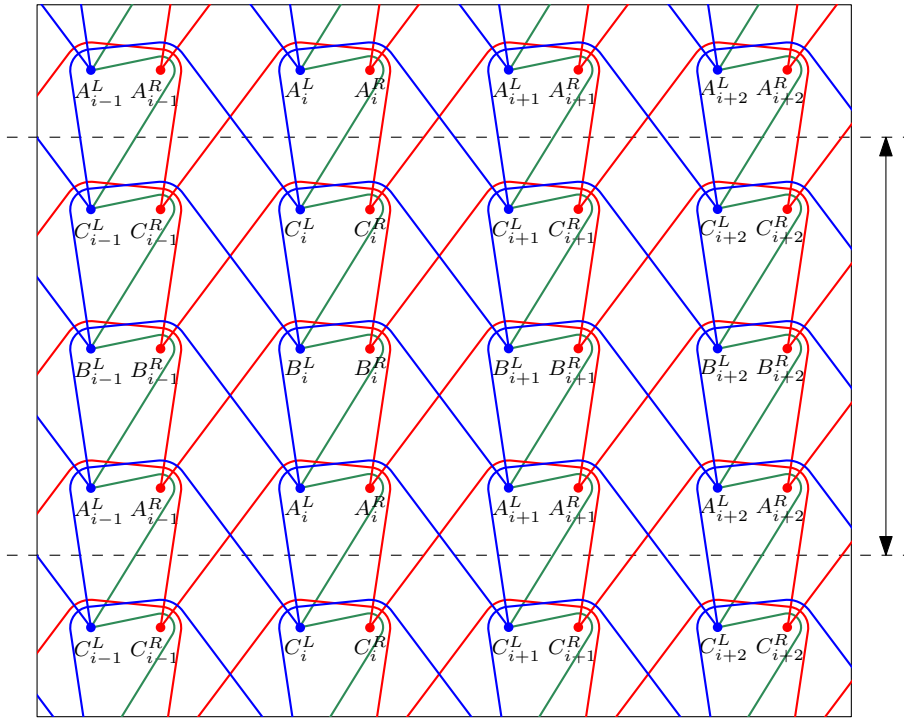- ℓ disjoint cycles of *green edges* connecting $A_i^L B_i^L C_i^L$.

Figure 1: The graph G on an unrolled cylinder.

The cylinder can be homeomorphically mapped into the plane, as shown in Fig. 2 for the red and blue edges only. The horizontal directions turn into radial directions. But the resulting drawings suffer from large distortions, and the left-right symmetry is lost. We therefore prefer the cylindrical drawings, and we extend the cylinder surface periodically beyond the dotted lines (using the plane as a universal cover of the cylinder). One should however be aware that vertices (and edges) that appear as distinct in the figure may denote the same vertex, as indicated by the vertex labels.
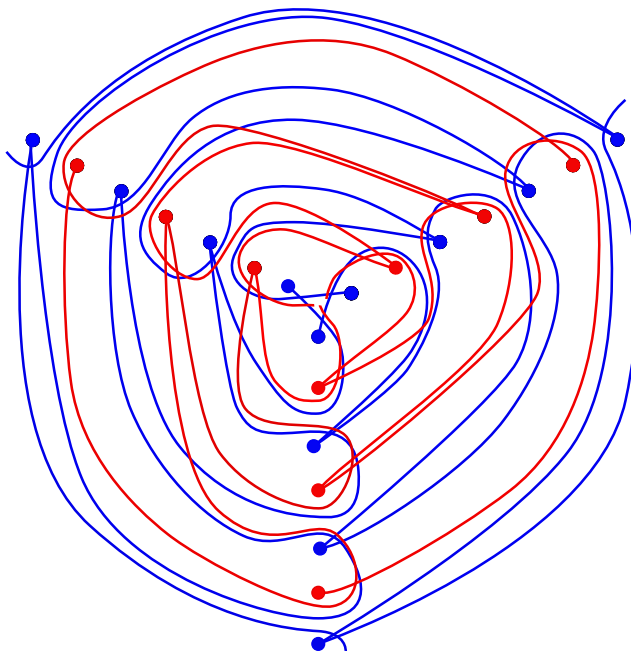
Figure 2: The graph G on the plane.

We will first consider the graph G_{RB} that omits the green edges, because this graph is more symmetric: with the exception of the vertices $X_1^{L/R}$ and $X_\ell^{L/R}$ near the boundary, all vertices look identical. Apart from these boundary effects, the drawing has a rotational symmetry, cyclically shifting the labels $A \to B \to C \to A$, a translational symmetry, shifting indices i up or down, and a mirror symmetry, exchanging left with right and blue with red. The green edges destroy this mirror symmetry: there are then two classes of vertices, the blue vertices X_i^L and the red vertices X_i^R.

We will show that the maximum degree in any saturated drawing which extends G_{RB} is 16. The 16 potential neighbors of a vertex on A_i^L are shown in Fig. 3. This establishes that there are saturated drawings with n vertices and less than $8n$ edges. When the green edges are included, the three dashed edges in Fig. 3 become impossible. Thus, each blue vertex has 13 potential neighbors. The red vertex A_{i+1}^R, which can be taken as a representative of a typical red vertex, loses A_i^L as a potential neighbor. Thus, each red vertex has at most 15 potential neighbors. This improves the upper bound for the smallest number of edges in a saturated drawings with n vertices to $7n$.

Theorem 1 *Let $s(n)$ denote the minimum number of edges that a simple saturated drawing with n vertices can have. Then $s(n) \leq 7n$.*

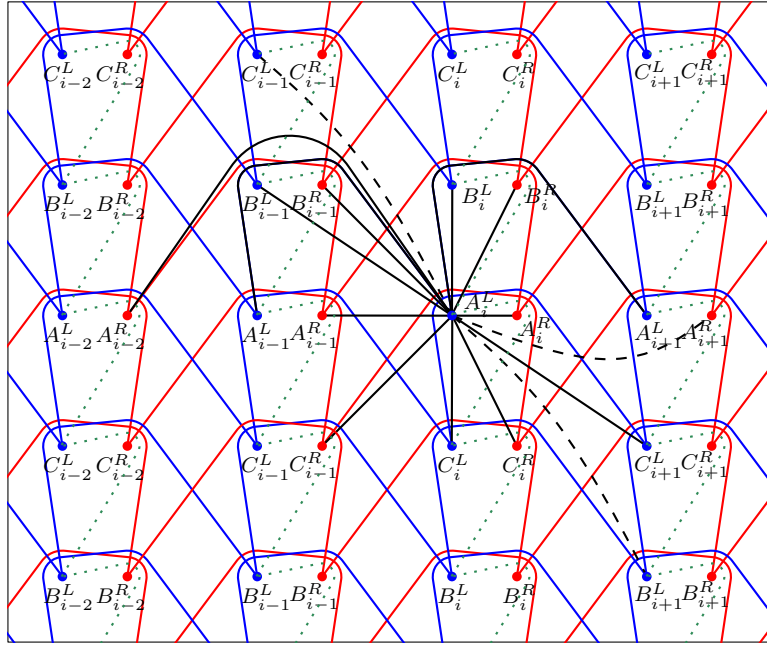The remainder of this section is devoted to proving the above theorem. We

Figure 3: The 16 potential neighbors of a vertex in the graph G_{RB}. The three dashed connections become impossible when the green edges (dotted) are included.

start with the analysis of the graph G_{RB}.

Lemma 1 *The* 16 *potential neighbors of a typical vertex* A_i^L *in* G_{RB} *are all* 11 *vertices of levels* $i-1$ *and* i $(A_{i-1}^L, B_{i-1}^L, C_{i-1}^L;\ A_{i-1}^R, B_{i-1}^R, C_{i-1}^R;\ B_i^L, C_i^L;\ A_i^R, B_i^R, C_i^R)$ *plus the* 5 *vertices* $A_{i-2}^R;\ A_{i+1}^L, B_{i+1}^L, C_{i+1}^L;\ A_{i+1}^R$.

When any of the neighbors listed above does not exist because $i \leq 2$ or $i = \ell$, the lemma still holds in the sense that the remaining vertices form the set of potential neighbors. In the proofs, when we exclude an edge between, say, levels i and j, our arguments will not use edges outside this range.

In the following we will look at the given drawing of G_{RB} (or G) and argue about the additional edges that can be drawn. The implicit assumption is that these edges cannot cross any given edge more than once. Usually, we will regard a new edge as a directed edge, starting at some vertex and trying to reach another vertex.

A *belt* is a substructure of our drawing. It is formed by the 12 vertices of two successive layers with their 6 edges between them, see Fig. 4. This drawing separates a large face on the left from a large face on the right. More precisely, the belt is defined as the part of the plane (or the cylinder) which lies between
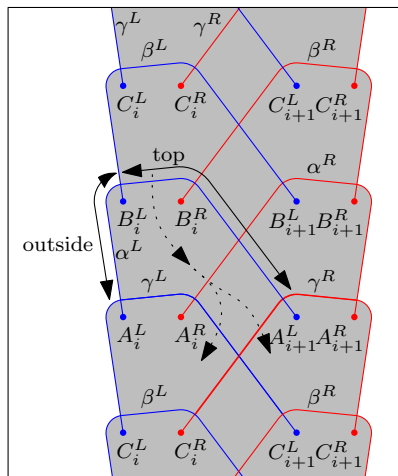
these two large faces.

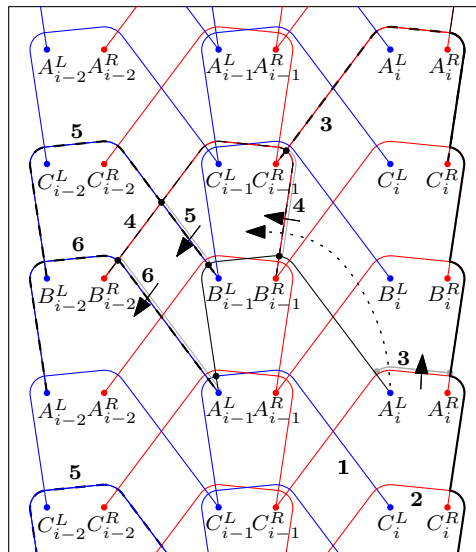Figure 4: Escape from a belt is difficult (Lemma 2).

Figure 5: The situation discussed in the proof of Lemma 1 for left side neighbors.

We denote the six edges of the belt by $\alpha^L = (A_i^L, A_{i+1}^L), \beta^L = (B_i^L, B_{i+1}^L), \gamma^L = (C_i^L, C_{i+1}^L), \alpha^R = (A_i^R, A_{i+1}^R), \beta^R = (B_i^R, B_{i+1}^R), \gamma^R = (C_i^R, C_{i+1}^R)$; as shown in Fig. 4. Each edge is cut into six sections by the intersections with the other edges: Two sections are little "stumps" at the end vertices. One section belongs to the boundary between the belt and the *outside*. The remaining three sections form the *top part* of the edge. We say that a new (directed) edge crosses a belt edge *from the outside* or *from the top* if it crosses the boundary part or the top part in the appropriate direction.

Lemma 2 *In a simple drawing that contains G_{RB}, the following holds:*

1. *If an edge crosses a belt edge from the top or from the outside, it must terminate inside the belt.*

2. *No edge can cross a belt, meaning it cannot pass through the shaded region in Fig. 4.*

Proof: We start with the following observation: If an edge crosses α_L from the outside or from the top, and it does not terminate at B_i^L or at B_i^R, then it must later cross γ^L or γ^R from the top. This observation holds symmetrically for α_R instead of α_L, and cyclically for the other four belt edges. Hence, any edge that "enters" the belt from the outside has to continue by crossing another edge of the belt from the top. There is no way to leave the belt without crossing some edge twice. □

After these preparations, we are ready to prove Lemma 1.

Proof of Lemma 1. Let us first look at the potential neighbors on the left side. A connection from A_i^L to levels $j \leq i-3$ is impossible, because it would have to cross a belt. For the vertices at level $i-2$ we observe the following (see Fig. 5 for the edge numbers we are referring to): When we start from A_i^L we cannot cross the right boundary of the belt formed by levels $i-1$ and i, because then we would have to cross the whole belt to reach level $i-2$. If we cross edge 1 or 2 from the top, then, by Lemma 2, we are restricted to the belt defined by level $i-1$ and i. Thus we can regard edge 1 and 2 as closed from the top. (These edges can later be crossed from the bottom.) We successively conclude that the new edges must cross the shadowed parts of the edges 3, 4, 5, and 6. The endpoints $B_{i-2}^R, B_{i-2}^L, A_{i-2}^L$ of the edges 4, 5, and 6 cannot be taken. C_{i-2}^L and C_{i-2}^R are enclosed in a small face delimited by the edges 4, 5, and 6, and cannot be reached. A_{i-2}^R is thus the only reachable vertex of level $i-2$.

Let us turn to the potential neighbors on the right side. A connection from A_i^L to levels $j \geq i+3$ is impossible, because it would have to cross a belt. Vertices at level $i+2$ cannot be reached either, because (i) if we cross the edge forming the left boundary of the belt spanned by the vertices of level i and $i+1$ we cannot cross this belt anymore and therefore cannot reach level $i+2$, and (ii) if we cross one of the edges in the face that contains A_i^L from the top (edge labeled 1 and 2 in Fig. 6a), then, by Lemma 2, we are also restricted to this belt. Thus we are restricted to the shaded region in Fig. 6a.

The vertices B_{i+1}^R and C_{i+1}^R also cannot be neighbors of A_i^L. We discuss the exclusion of C_{i+1}^R as a potential neighbor – the case for B_{i+1}^R is symmetric. The edges incident to A_i^L and C_{i+1}^R, which we call the *closed edges*, cannot be crossed. The closed edges are depicted as thicker curves in Fig. 6b. Consider the portion of the red edge π_r that runs between A_i^R and A_{i+1}^R above the closed edges (see Fig. 6b). The curve π_r bounds a region below in which the remaining edges bounding this region are parts of the closed edges. Hence, if we enter this region we cannot leave and therefore we cannot cross π_r (see Fig. 6b). Let us now consider the partial edge π_b that runs between B_{i+1}^L and B_i^L above the closed edges and π_r. Again, there is a region whose boundary is part of the closed edges and also π_b. To enter and leave this region we have to cross either one of the closed edges or π_r, or we have to cross π_b twice. Since all these options are invalid, we have to avoid this region, and therefore are not allowed to cross π_b. We observe that the closed edges together with π_b and π_r leave A_i^L and C_{i+1}^R in different faces, which shows that these vertices cannot be neighbors unless we cross one edge twice. □

Now we turn back to G. The additional green edges exclude some of the possible edges from the Lemma 1.

Lemma 3   *1. The* 13 *potential neighbors of a typical vertex A_i^L in G are all 5 vertices of level i (B_i^L, C_i^L; A_i^R, B_i^R, C_i^R), all but one vertex of level $i-1$ (A_{i-1}^L, B_{i-1}^L; $A_{i-1}^R, B_{i-1}^R, C_{i-1}^R$) plus the 3 vertices A_{i-2}^R; A_{i+1}^L, C_{i+1}^L.*

2. The 15 *potential neighbors of a typical vertex A_i^R in G are all* 11
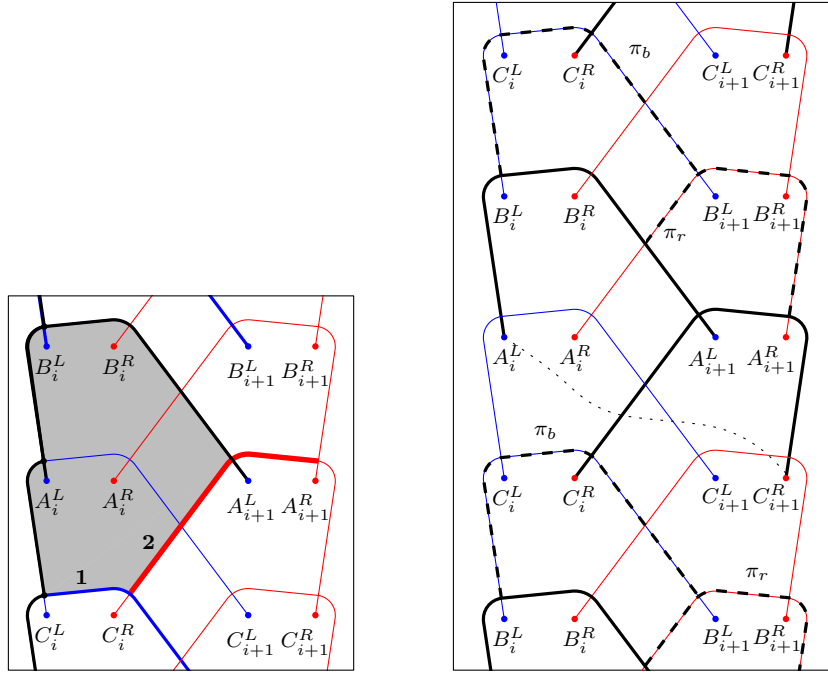
(a) Level $i + 2$ cannot be reached from A_i^L.   (b) C_{i+1}^R cannot be reached from A_i^L.

Figure 6: Restricting the neighbors to the right.

vertices of levels i and $i + 1$ (A_i^L, B_i^L, C_i^L; B_i^R, C_i^R; $A_{i+1}^L, B_{i+1}^L, C_{i+1}^L$; $A_{i+1}^R, B_{i+1}^R, C_{i+1}^R$) plus the 4 vertices $A_{i-1}^R, B_{i-1}^R, C_{i-1}^R$; A_{i+2}^L.

The claim immediately follows from the next two lemmas.

Lemma 4 *In a simple extension of G, A_{i+1}^R cannot be a neighbor of A_i^L.*

Proof: We call the edges incident to A_{i+1}^R and A_i^L the *closed edges*. Let π_b the portion of the edge connecting B_i^L with B_{i+1}^L that runs above the closed edges (see Fig. 7a). The cell "below" π_b is only bounded by π_b and the closed edges. Hence, we cannot leave this cell once we have entered. As a consequence we cannot cross π_b. Since the closed edges together with π_b disconnect A_{i+1}^R and A_i^L, these two vertices cannot be neighbors. □

Lemma 5 *In a simple extension of G, B_{i+1}^L cannot be a neighbor of A_i^L.*

Proof: All edges that are incident to either B_{i+1}^L or A_i^L cannot be crossed. These edges are drawn as black curves in Fig. 7b and are now considered as being the *closed* edges. The only chance to connect A_i^L with B_{i+1}^L is to enter the region that is bounded by the closed edges and the edge π_A from A_i^R to A_{i+1}^R. Thus we have to cross this edge to leave this face. This leads us to a

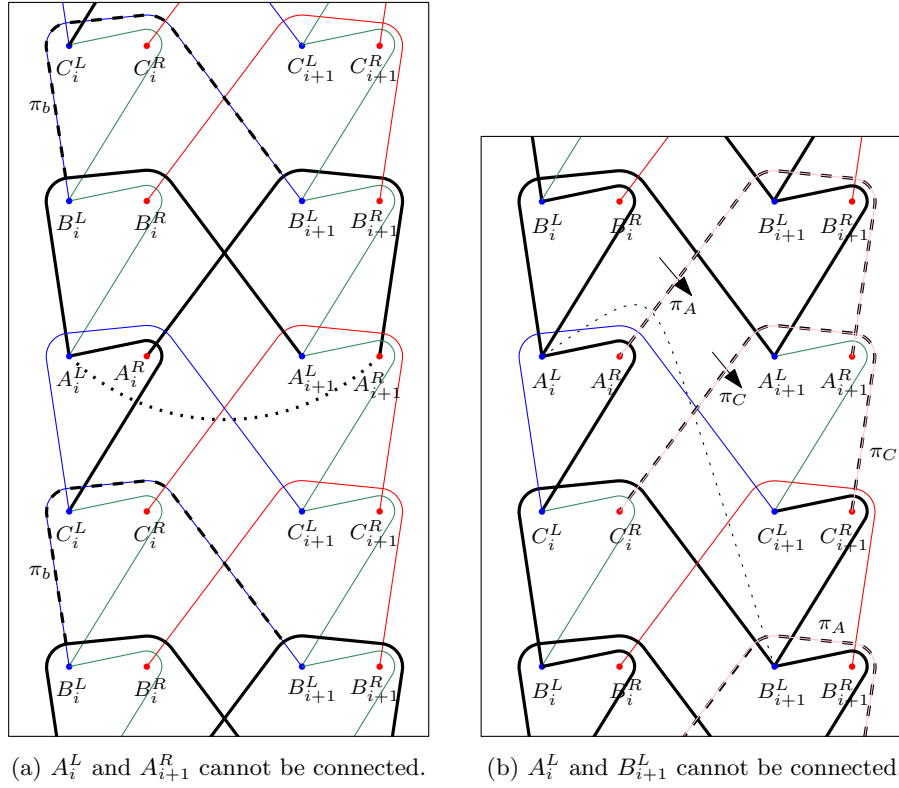(a) A_i^L and A_{i+1}^R cannot be connected.   (b) A_i^L and B_{i+1}^L cannot be connected.

Figure 7: Restricting more neighbors by putting back the green edges.

region that is bounded by the closed edges, π_A and the edge π_C from C_i^R to C_{i+1}^R. Clearly we have to cross π_C to leave this region. Now we have entered a region that is separated from B_{i+1}^L by the closed edges, π_A, and π_C, thus B_{i+1}^L can not be reached.   □

By symmetry, C_{i-1}^L and A_i^L cannot be neighbors, and this concludes the proof of Lemma 3. Moreover, as a consequence of Lemma 3 the average degree in a saturated extension of G is at most 14, which proves Theorem 1 when the number n of vertices is a multiple of 6.

We can determine the vertex degrees more carefully. If $\ell \geq 3$, then

1. the degrees of A_1^L, B_1^L, C_1^L are at most 7,

2. the degrees of A_1^R, B_1^R, C_1^R are at most 12,

3. the degrees of A_2^L, B_2^L, C_2^L are at most 12,

4. the degrees of A_i^R, B_i^R, C_i^R are at most 15, when $1 < i < \ell - 1$,

5. the degrees of A_i^L, B_i^L, C_i^L are at most 13, when $2 < i < \ell$,

6. the degrees of $A_{\ell-1}^R, B_{\ell-1}^R, C_{\ell-1}^R$ are at most 14,

7. the degrees of $A_\ell^L, B_\ell^L, C_\ell^L$ are at most 11,

8. the degrees of $A_\ell^R, B_\ell^R, C_\ell^R$ are at most 8.

A straightforward calculation gives that any saturated extension of G has at most $7n - 30$ edges. For $\ell = 2$, the degrees of $X_1^L, X_1^R, X_2^L, X_2^R$ are bounded by $7, 11, 10, 8$, respectively, for a total of 54 edges, which also agrees with the formula $7n - 30$. Hence, for any $n \geq 12$ that is a multiple of 6, there exists a saturated simple topological graph with n vertices and at most $7n - 30$ edges.

Our construction can be extended to any vertex size by *cloning* some vertices. Take a saturated simple topological graph and any vertex P of it. Next to P we add ρ new copies of P – the clones. Connect the neighbors of P to each clone by edges that are perturbations of the edges incident to P (not intersecting themselves). By this we obtain a simple drawing. A saturation of this drawing can include as additional edges only edges among P and its clones.

For $n \geq 12$, we can write n as $6r + \rho$ where $0 \leq \rho \leq 5$. If $\rho = 0$, we are done. If $\rho \geq 1$, then start with a construction for a saturated simple topological graph with $6r$ vertices. Add ρ clones of its lowest-degree vertex P, and saturate. In our construction, the lowest degree is 7. Cloning such a vertex ρ times adds up to $7\rho + \binom{\rho+1}{2}$ additional edges after saturation. Since $\rho \leq 5$, the number of edges is bounded by

$$7(6r) - 30 + 7\rho + \binom{\rho+1}{2} \leq 7(6r + \rho) - 30 + 15 = 7(6k + \rho) - 15 < 7n$$

The resulting simple topological graph proves Theorem 1 for $n \geq 12$. If $n \leq 11$, then the bound of Theorem 1 holds since even the complete graph has at most $\binom{n}{2} \leq 5n$ edges.

3 Saturated 2-simple topological graphs with few edges

3.1 The grid-block configuration

To begin, we study a drawing of 6 edges (three red edges and three black edges) as depicted in Fig. 8. The drawing consists of three disjoint horizontal segments representing the red edges r_1, r_2, r_3, and three disjoint black edges b_1, b_2, b_3 that are drawn such that one crosses (in order) $r_1, r_2, r_3, r_1, r_2, r_3$, the other $r_2, r_3, r_1, r_2, r_3, r_1$, and the last one $r_3, r_1, r_2, r_3, r_1, r_2$. There are no other crossings in the drawing. Note that the configuration superimposes a grid. We call such an arrangement of edges a *grid-block*. These blocks have been used by Kynčl et al. as building blocks in their saturated graphs [4]. In the terminology of Kynčl et al., our grid-blocks would be called (3,2)-grid-blocks.

As done in the previous section we consider the graph as drawn on the cylinder. Different to the presentation in the previous section we draw no overlap. In

particular, we draw the graph inside a rectangle in which we identify two sides in opposition (*bottom side* and *top side*), while the other sides are named *right side* and *left side*. If an edge uses the transition across the bottom/top edge we say that it *wraps around*. In the following we assume that the grid-blocks are drawn such that only the black edges wrap around. We label every face of the drawing of a grid-block with 2 numbers. These numbers refer to the coordinates of the (dual) superimposed grid, with $(0,0)$ being the label of the face that contains the two bottom most endpoints of the black edges on the left side. All "vertical" coordinates are considered modulo 3.
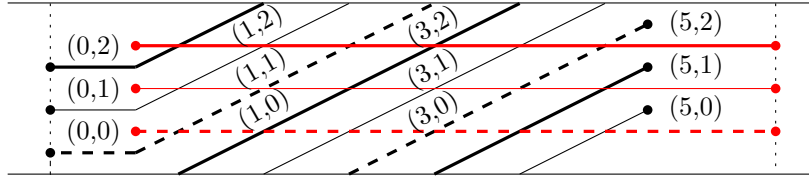
Figure 8: A grid-block with some labeled faces.

Throughout the section we study *paths* connecting the left and the right sides of the cylinder and passing through some blocking configurations. By a *path* in this context we always mean a path in a graph dual to the arrangement of the blocking configuration in question.

Kynčl et al. observed that every path connecting the left with the right side of the cylinder has to intersect the edges of grid-block at least 5 times. For our construction we need a stronger statement which is presented in Lemma 6. The following lemma simplifies the treatment of paths passing through the grid-block.

Lemma 6 *Let γ be a path crossing the grid-block that starts in face $(0, i)$ and ends in face $(5, j)$ and that never visits the faces $(0, \cdot)$, $(5, \cdot)$ again, see Fig. 9. Then γ can be transformed, keeping its endpoints fixed, to a path $\widetilde{\gamma}$ passing through the grid block, such that $\widetilde{\gamma}$:*

1. *crosses (with the same or smaller multiplicity) only the edges of the grid-block crossed by γ,*

2. *first walks between the faces $(0, i)$, $0 \leq i \leq 2$, then crosses some black edges to the right, passing from a face (a, i) to a face $(a+1, i)$, then crosses some red edges upwards, passing from a face (a, i) to a face $(a + 1, i + 1)$.*

Proof: We refer to the transition of the path from one cell of the arrangement to an adjacent cell as a *step*. There are four different types of steps: $\rightarrow$, $\leftarrow$, $\nearrow$ or $\swarrow$, depending on the crossed edge and the direction, see Fig. 9.

We execute the path simplification through a series of local modifications on pairs of two consecutive steps: (1) annihilation of two consecutive steps in opposite directions and (2) changing places of two consecutive steps that are not yet in a desired order.
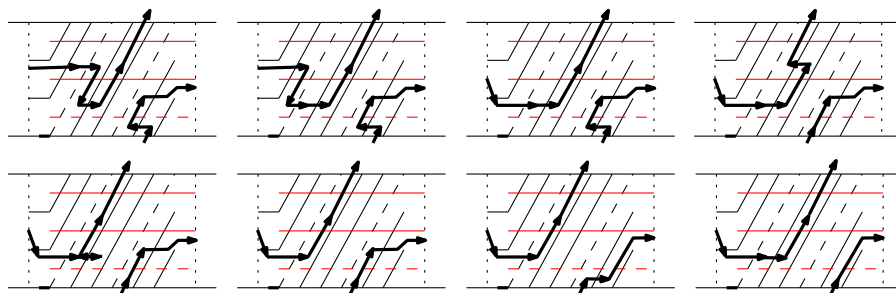
Figure 9: Process of the simplification of a path passing through a grid-block.

The simplification is carried out in two stages. In the first stage (shown in the first 6 pictures in Fig. 9) we remove all "backward steps" $\swarrow$ and $\leftarrow$, while possibly increasing the number of steps the path $\widetilde{\gamma}$ walks between faces $(0, i)$, $0 \leq i \leq 2$. In the second stage we reorder the steps $\nearrow$ and $\rightarrow$ such that no $\nearrow$ precedes any $\rightarrow$.
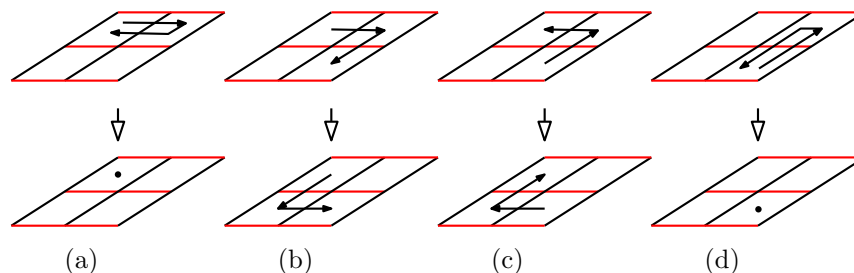

Figure 10: 4 possible 2-step configurations involving "backward steps" as a second step before (first row) and after (second row) the appropriate local modification.

Stage 1: We traverse the path until we meet the first $\leftarrow$ or $\swarrow$ step. Together with its preceding step it forms one of the 4 configurations shown in Fig. 10. In cases (a) and (d) the steps only differ in their orientation, hence we can annihilate two steps. In the remaining cases (b) and (c) we reorder the two steps. This reordering can be safely executed unless it forces the path to leave the grid-block. This, however, may happen only when the backward step ($\leftarrow$ or $\swarrow$) starts from one of the faces labeled $(2, \cdot)$. Since this backward step is the first backward step of the path, we are left with two subcases for each (b) and (c) depending on the preceding step, which might be either $\rightarrow$ or $\nearrow$. The four cases are depicted in Fig. 11. All the cases can be handled by further local simplifications that are shown in the figure.

We finish the proof of the stage 1 using double induction on the number of backward steps and, within it, on the distance from the beginning of the path to the first backward step.
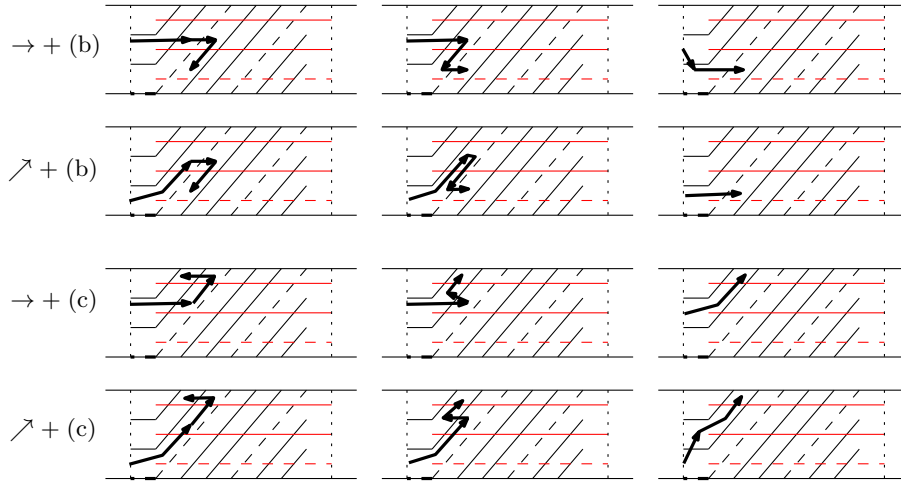
Figure 11: Handling of the 4 possible cases (each row represents one case) when the local modifications (b) or (c) force the path out of the grid-block.

Stage 2: After the stage 1 our path through the grid, leaving aside its first steps between faces $(0, i)$, has only $\rightarrow$ and $\nearrow$ steps. These two types of steps can be reordered without changing the number of times the path crosses any edge of the grid. Moreover, this reordering never leads the path out of the grid-block. □

3.2   A blocking configuration

We call the building blocks of the following constructions *black block* and *red block*, see Fig. 12. We refer to the edges of the red (black) block as *red edges* (*black edges*). Any two red edges, as well as any two black edges, cross exactly twice. Note that up to a reflection the red block is isotopic to the black block.
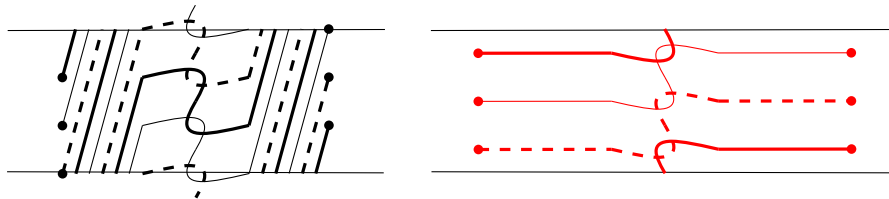

Figure 12: A black (left) and a red (right) blocks.

We combine two black blocks and a red block as shown in Fig. 13 to obtain a drawing that we call a *3-block*. Since the red block differs from the black block only by a reflection, the 3-block built from consecutive black-red-black blocks is a mirror image of the 3-block built from consecutive red-black-red blocks.
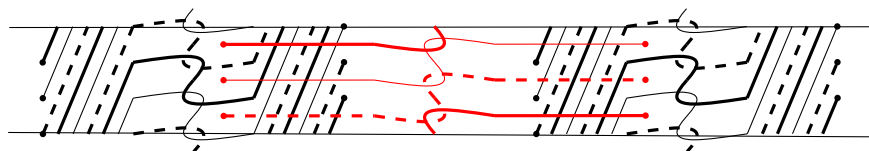
Figure 13: A 3-block, formed by consecutive black, red and again black blocks.

The following theorem is the key observation that we need for the construction of the sparse 2-simple drawing.

Theorem 2 *Any path connecting the left with the right sides of the cylinder while passing through the 3-block crosses one of the edges forming the 3-block at least 3 times.*

Before proving the theorem we provide some helpful lemmas. We label some of the faces of the arrangement as shown in Fig. 14. In particular, for $i = 0, 1, 2$, we denote the faces containing the left endpoint of the red edges r_i as L_i, and the faces containing the right endpoint as R_i. The edges of the left black block are named b_i and the edges of the right black block are named b_i'. Finally, let LM_i be the face that contains the right endpoint of b_i, and let RM_i be the face that contains the left endpoint of b_i'. The region spanned by L_0, L_1 and L_2 is denoted by L. We similarly define regions LM, RM and R.
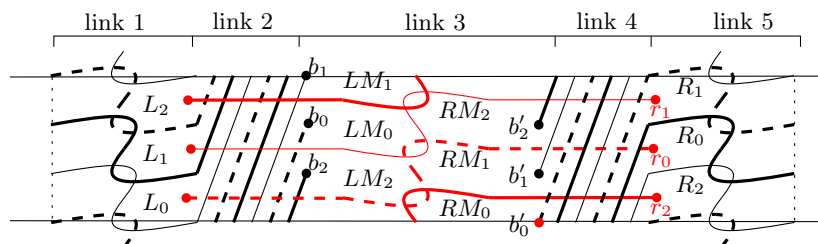

Figure 14: A 3-block with some distinguished faces (capital letters) and edges. The red edges forming the blocks are labeled b_i, b_i' and r_i. The "zones" at which we subdivide the path into links are labeled above the strip.

Let γ be a path that passes through the 3-block. To facilitate the analysis we subdivide the path γ into smaller pieces, which we call *links*. The links are defined as follows:

link 1: from the start point (left) of γ to the last point of γ in L,

link 2: from the last point of γ in L to its first point in LM,

link 3: from the first point of γ in LM to its last point in RM,

link 4: from the last point of γ in RM to its first point in R,

link 5: from the first point of γ in R to its (right) endpoint.

Before we proceed we check that the links are well defined, i.e., that the points defining the links appear in order. For the links 1, 3 and 5 this holds trivially, while to check it for links 2 (and, symmetric, 4), we need to prove that the last point in L precedes the first point in LM:

Lemma 7 *No path can visit the regions $L \to LM \to L \to LM$ in this order without crossing some of the edges forming the 3-block at least 3 times.*

Proof: The faces L and LM are separated by a grid-block. Passing through it requires at least 5 crossings of its edges. Any path visiting $L \to LM \to L \to LM$ would cross the grid-block at least 3 times, and hence it would cross the edges of the grid-block at least $3 \cdot 5 = 15$ times. Since a grid-block is formed by 6 edges, at least one of them will be crossed 3 times or more. □

We continue by analyzing the path through the 3-block following its links.

Lemma 8 *Any path passing through the 3-block from left to right with the last point of link 1 at L_i crosses the edge b_{i+1} at least once or one of the edges b_i and b_{i+2} at least twice at its first link (all indices modulo 3).*

Proof: A path that ends in L_i crosses either b_{i+1} or it crosses b_{i+2} while entering from L_{i+1}. If it came from L_{i+1} it has to enter L_{i+1} before. This can be achieved by either crossing b_{i+2} or by crossing b_i. In the former case we would have crossed b_{i+2} twice. In the later case the path came from L_{i+2}. To enter L_{i+2} it either has to cross b_i a second time or it has to go back to L_i. In the latter case we could as well start form the beginning. So this case is already covered. □

The following lemma summarizes the behavior of the path on the first two links:

Lemma 9 *Any path γ passing the 3-block that does not intersect any edge 3 times or more crosses the red edges r_j, r_{j+1} before it first visits the zone LM at LM_j.*

Proof: We modify the path γ along link 2 following the simplification procedure described in Lemma 6 to get a path $\widetilde{\gamma}$. Lemma 6 also implies that the link 2 of $\widetilde{\gamma}$ consists of exactly 5 "steps": first, $0 \le h \le 5$ steps crossing the black edges $\to$ to the right, followed by $v = 5 - h$ steps crossing red edges $\nearrow$ upward.

Assume that the first point of link 2 of $\widetilde{\gamma}$ lies inside the face L_i. Then h horizontal steps of link 2 cross the b_{i+1}, b_i, b_{i-1}, ..., $b_{i+1-(h-1)}$. Moreover, Lemma 8 guarantees that already link 1 of the path $\widetilde{\gamma}$ crossed either b_{i+1} once or one of b_i or b_{i+2} twice. Since $\widetilde{\gamma}$ does not cross any of the black edges more than twice, it follows that $h \le 3$. This, however, shows that $v \ge 2$, which implies that the path $\widetilde{\gamma}$ crosses the red edges r_{j+1}, r_j before it reaches the last point of its second link in face LM_j. To finish the proof we recall that the path γ crosses every edge of the 3-block at least as many times as $\widetilde{\gamma}$ and that the last points of the link 2 of γ and $\widetilde{\gamma}$ coincide. □

Proof of Theorem 2.   We prove by contradiction, namely, we assume that there is a path γ that passes through the 3-block while crossing every edge of the 3-block at most twice. Let LM_j be the face where link 2 ends, and let RM_ℓ be the face where link 4 starts. By Lemma 9 we know that γ crosses r_j and r_{j+1} in link 1 and link 2. Since the structure of the link 4 and 5 coincides with the structure of link 2 and 1 we can apply Lemma 9 also to the last two links. Thus, γ crosses $r_{\ell-1}$, r_ℓ in link 4 and 5. A short case distinction (ℓ might be either j, $j+1$, or $j+2$) shows that γ cannot connect endpoints of link 2 and 4 via link 3 without crossing at least one of the red edges 3 times; see Fig. 15. The figure depicts all ways of how to possibly route the path γ in link 3. Each of the possible continuations crosses some of the red edges r_j, r_{j+1}, r_{j-1} twice and is blocked within one of the faces before it reaches the face RM_ℓ. As a consequence the path γ cannot exists.   □
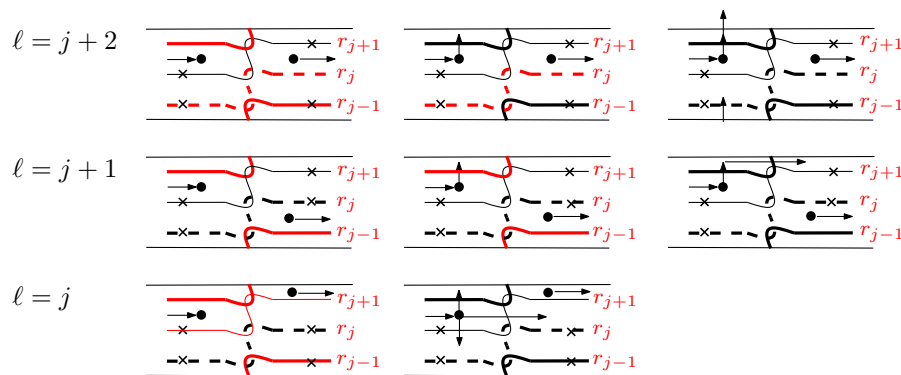

Figure 15: Each row depicts a case. Black dots inside faces mark the faces LM_j (left) and RM_ℓ (right). Black crosses on red edges mark the edges that are, due to Lemma 9, crossed by the path outside link 3. We color red edges black as soon as they are crossed by the path γ twice and no more crossings are allowed. In the case $\ell = j$ the path can be continued in 3 different directions, in each of them the path is blocked after one step.

3.3   A sparse saturated 2-simple drawing

We show now how to combine a sequence of 3-blocks to obtain a 2-simple saturated drawing with few edges.

Theorem 3 *Let $s_2(n)$ denote the minimum number of edges that a 2-simple saturated drawing with n vertices can have. Then $s_2(n) \leq 14.5n$.*

Proof: We consider the drawing that repeats the pattern shown in Fig. 16. The drawing is formed by ℓ consecutive black and red blocks; see Fig. 12. Each block contains 6 vertices, so the total number of vertices is 6ℓ. Clearly, the drawing is 2-simple.
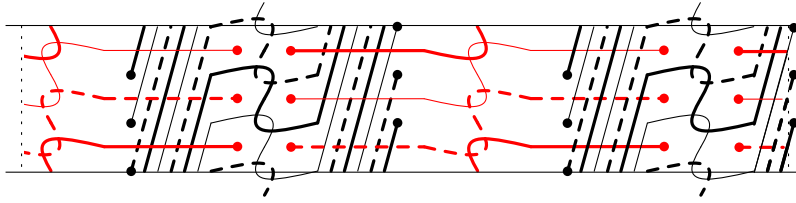
Figure 16: A 2-simple drawing that does not allow too many edges to be added.

Now we add as many edges as possible without violating the 2-simplicity, so that the drawing becomes saturated (this padding procedure is definitely not unique). Theorem 2 implies that without violating the 2-simplicity any vertex can be connected by an edge only to 29 other vertices; see Fig. 17 for "internal" vertices and Fig. 18 for vertices close to the left (right) boundary of the cylinder. This implies that the maximal number of edges in the resulting saturated 2-simple drawing is less than or equal to $14.5n$.
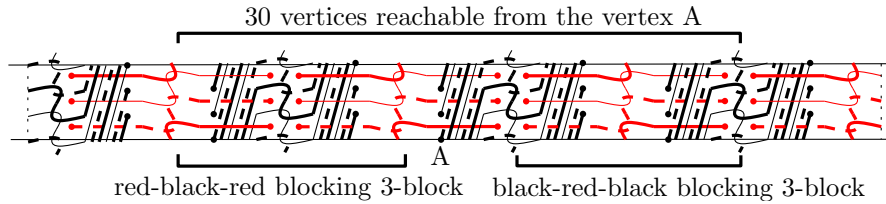

Figure 17: The potential neighbors of a typical vertex A.

For n not divisible by 6 we build the construction above with $\ell = \lfloor n/6 \rfloor$. We split the remaining $n - 6\lfloor n/6 \rfloor$ vertices into two groups of no more than 3 vertices each, and place one group with l_1 vertices to the left and one group with l_2 vertices on to the right of the resulting arrangement.
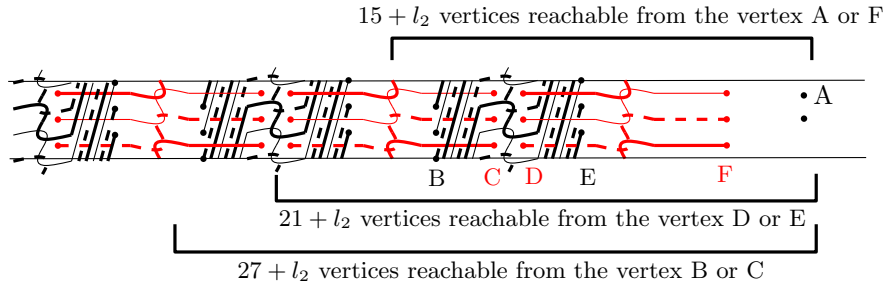

Figure 18: The potential neighbors of vertices close to the boundary.

The possible connections with the newly introduced vertices are illustrated in Fig. 18. Since $l_1, l_2 \leq 3$, no vertex has degree greater than 29. $\qquad\square$

4 Local saturation

4.1 Simple drawings

The lower bound in [4] on the number of edges in a saturated simple topological graph is based on the following lemma.

Lemma 10 ([4]) *Let G be a simple topological graph with at least four vertices, and let A be a vertex of degree at most two. Then G has a simple extension by an edge incident to A.*

This lemma implies that in a simple saturated topological graph with at least four vertices, every vertex must have degree at least three, and hence the number of edges is at least $1.5n$. Can we improve the bound on the edge number by strengthening the lower bound on the degree? The following considerations establish a limit to this approach: There are saturated graphs with minimum degree four.

We say that a vertex S in a simple topological graph is *saturated* if it cannot be connected to a non-adjacent vertex while maintaining simplicity. The above lemma implies that in a simple topological graph with at least four vertices, a saturated vertex must have degree at least three.

Observation 1 *For any positive integer $n \geq 6$, there is a simple topological graph on n vertices with a saturated vertex of degree four.*
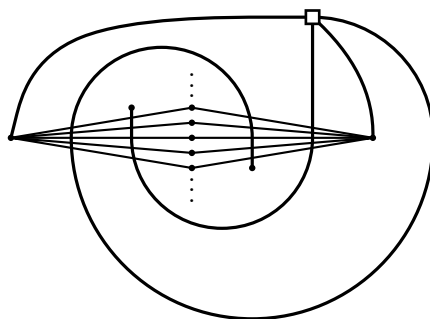

Figure 19: The boxy vertex of degree four is saturated.

The observation is due to the construction presented in Fig. 19. This example is an extension of the case $n = 6$ from [4, Fig. 2]. The topmost vertex is saturated, since only the straight edges are not incident to that vertex. It is easy to see, that in order to connect a vertex p to the degree four vertex, one has to cross an edge incident to p before reaching p.

The following lemma presents a construction that realizes small vertex degrees for many vertices.

Lemma 11 *For any positive integer ℓ, there exists a saturated simple topological graph on 10ℓ vertices with ℓ vertices of degree 5.*
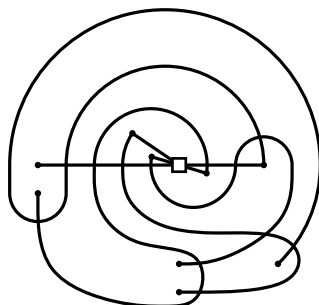
Figure 20: In the simple topological graph above, the central vertex has degree 5, and it cannot be connected by an edge to any point in the unbounded region while keeping simplicity.

Proof: The main idea of our construction is depicted in Fig. 20. A simple case distinction verifies that no edge can connect the central vertex with a point on the outer face without violating the simplicity of drawing.

Now, take ℓ copies of the drawing in Fig. 20, and place them on the plane next to each other such that the interior faces of the copies are non-overlapping. The ℓ copies of the central vertex will remain degree-5 vertices no matter how we saturate the graph. □

4.2   2-simple drawings

To study local saturation in 2-simple case we use a slight modification of the 3-block introduced in Sect. 3; see Fig. 21.
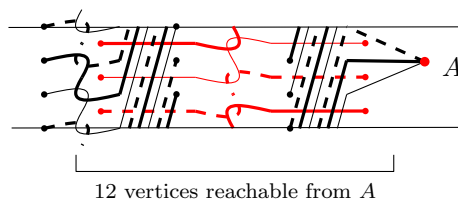

12 vertices reachable from A

Figure 21: The rightmost vertex A cannot be connected to any vertex that belongs to the leftmost (unbounded) face without violating 2-simplicity.

By the arguments given in the proof of Theorem 2 the rightmost vertex can be connected to only 12 other vertices (Fig. 21) and thus it cannot be connected to any vertex that belongs to the leftmost (unbounded) face of the drawing without violating 2-simplicity.

The "unrolling" of this configuration from the cylinder to the plane (with center of the unrolling in the rightmost vertex) is presented in Fig. 22. The central vertex cannot be connected by an edge to any vertex that belongs to the unbounded region without violating 2-simplicity, and so it has degree no larger
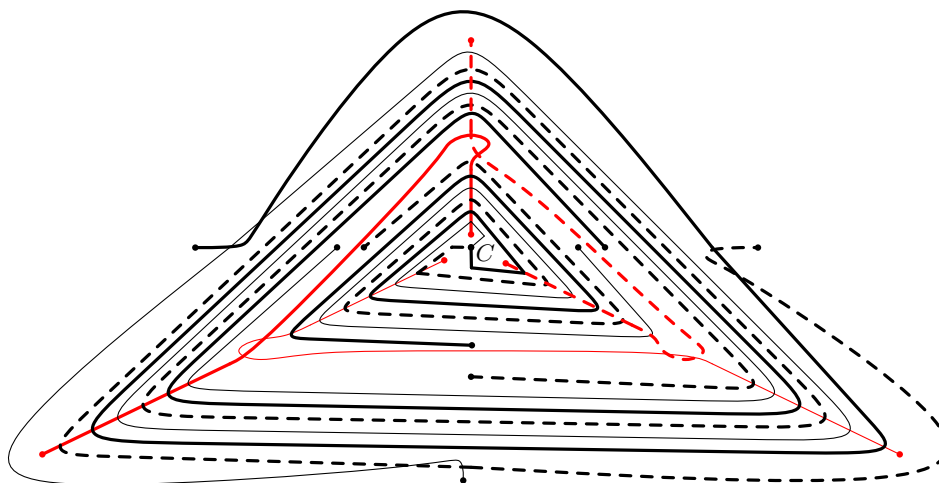
Figure 22: Unrolling of Fig. 21 to the plane. The central vertex C corresponds to the rightmost vertex A of Fig. 21.

than 12 in any saturation. After placing ℓ disjoint copies of this construction to the plane next to each other we obtain the following result:

Lemma 12 *For any positive integer ℓ, there exists a saturated 2-simple topological graph on 16ℓ vertices with ℓ vertices of degree 12.*

5   Open Problems

The main open problem is obviously to close the gaps between the upper and lower bounds. It is somehow intriguing that it is unclear if the function $s_k(n)/n$ is increasing, decreasing or non-monotone in k.

We expect that the lower bounds are not tight and that the true value of $s_k(n)/n$ is closer to the upper bounds. In fact, the lower bound constructions are rather elementary and ignore many facets of the problem by looking only at the minimal vertex degree [4]. Even though we showed in Sect. 4 that this strategy is limited, we think that a more careful execution with a clever case distinction could improve the current lower bounds.

Acknowledgment

References

[1] P. Erdős, A. Hajnal, and J. W. Moon. A problem in graph theory. *The American Mathematical Monthly*, 71(10):1107–1110, 1964. `doi:10.2307/2311408`.

[2] J. R. Faudree, R. J. Faudree, and J. R. Schmitt. A survey of minimum saturated graphs. *Electron. J. Combin*, 18:36, 2011.

[3] L. Kászonyi and Zs. Tuza. Saturated graphs with minimal number of edges. *Journal of Graph Theory*, 10(2):203–210, 1986. `doi:10.1002/jgt.3190100209`.

[4] J. Kynčl, J. Pach, R. Radoičić, and G. Tóth. Saturated simple and k-simple topological graphs. *Computational Geometry*, 48(4):295–310, 2015. `doi:10.1016/j.comgeo.2014.10.008`.

[5] L. Lovász. Three short proofs in graph theory. *Journal of Combinatorial Theory, Series B*, 19(3):269–271, 1975. `doi:10.1016/0095-8956(75)90089-1`.