



Maximum independent sets in subcubic graphs: New results [☆]

A. Harutyunyan ^a, M. Lampis ^a, V. Lozin ^{b,*}, J. Monnot ^a

^a Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016 Paris, France

^b Mathematics Institute, University of Warwick, Coventry, CV4 7AL, UK

ARTICLE INFO

Article history:

Received 1 April 2020

Received in revised form 28 August 2020

Accepted 5 September 2020

Available online 11 September 2020

Keywords:

Maximum independent set

Subcubic graph

Polynomial algorithm

ABSTRACT

The maximum independent set problem is known to be NP-hard in the class of subcubic graphs, i.e. graphs of vertex degree at most 3. We study complexity of the problem on hereditary subclasses of subcubic graphs. Each such subclass can be described by means of forbidden induced subgraphs. In case of finitely many forbidden induced subgraphs a necessary condition for polynomial-time solvability of the problem in subcubic graphs (unless $P = NP$) is the exclusion of the graph $S_{i,j,k}$, which is a tree with three leaves of distance i, j, k from the only vertex of degree 3. Whether this condition is also sufficient is an open question, which was previously answered only for $S_{1,k,k}$ -free subcubic graphs and $S_{2,2,2}$ -free subcubic graphs. Combining various algorithmic techniques, in the present paper we generalize both results and show that the problem can be solved in polynomial time for $S_{2,k,k}$ -free subcubic graphs, for any fixed value of k .

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

In a graph, an *independent set* is a subset of vertices no two of which are adjacent. The maximum independent set problem asks to find in a graph G an independent set of maximum size. The size of a maximum independent set in G is called the *independence number* of G and is denoted $\alpha(G)$.

The maximum independent set problem is one of the first problems that has been shown to be NP-hard. Moreover, the problem remains NP-hard under substantial restrictions. In particular, it is NP-hard for graphs of vertex degree at most 3, also known as *subcubic* graphs. In terms of vertex degree, this is the strongest possible restriction under which the problem remains NP-hard, since for graphs of vertex degree at most 2 the problem is solvable in polynomial time. However, with respect to other parameters the restriction to subcubic graphs is not best possible, as the problem remains NP-hard for subcubic graphs of girth at least k for any fixed value of k [11], where the girth of a graph is the size of a smallest cycle. In other words, the problem is NP-hard for (C_3, \dots, C_k) -free subcubic graphs for each value of k , where C_k is a chordless cycle of length k . The idea behind this conclusion is quite simple: it is not difficult to see that a double subdivision of an edge increases the independence number of the graph by exactly one. Applying edge subdivisions repeatedly to a subcubic graph G , we can destroy all small cycles in G without increasing the maximum vertex degree of the graph, i.e. we can transform G in polynomial time into a subcubic graph of large girth.

[☆] Some results presented in this paper appeared in the extended abstract [5] published in the proceedings of the 45th International Workshop on Graph-Theoretic Concepts in Computer Science, WG 2019.

* Corresponding author.

E-mail addresses: ararat.harutyunyan@dauphine.fr (A. Harutyunyan), michail.lampis@dauphine.fr (M. Lampis), V.Loizin@warwick.ac.uk (V. Lozin), Jerome.Monnot@dauphine.fr (J. Monnot).




Fig. 1. The graphs H_k (left) and $S_{i,j,k}$ (right).

Let us observe that by means of edge subdivisions we can also destroy small copies of some other graphs, in particular, graphs of the form H_k represented in Fig. 1 (left). Therefore, the maximum independent set problem remains NP-hard for $(C_3, \dots, C_k, H_1, \dots, H_k)$ -free subcubic graphs for each value of k .

Let us denote by S_k the class of $(C_3, \dots, C_k, H_1, \dots, H_k)$ -free subcubic graphs and by $\kappa(G)$ the maximum k such that $G \in S_k$. If G belongs to no class S_k , then $\kappa(G)$ is defined to be 0, and if G belongs to all classes S_k , then $\kappa(G)$ is defined to be ∞ . Also, for a set of graphs M , $\kappa(M)$ is defined as $\kappa(M) = \sup\{\kappa(G) : G \in M\}$. With this notation, we can derive the following conclusion from the above discussion (see e.g. [8]).

Theorem 1. *Let M be a set of graphs. If $\kappa(M) < \infty$, then the maximum independent set problem is NP-hard in the class of M -free subcubic graphs.*

This theorem suggests that, unless $P = NP$, the maximum independent set problem is solvable in polynomial time in the class of M -free graphs only if the parameter κ is unbounded in the set M . There are three basic ways to unbind this parameter in M :

1. include in M a graph G with $\kappa(G) = \infty$;
2. include in M graphs with arbitrarily large induced cycles;
3. include in M graphs with arbitrarily large induced subgraphs of the form H_k .

To give an example of a polynomial-time result of the first type, let us observe that $\kappa(G) = \infty$ if and only if every connected component of G has the form $S_{i,j,k}$ represented in Fig. 1 (right). We call any graph of the form $S_{i,j,k}$ a *tripod*.

In other words, if the set M of forbidden induced subgraphs is finite, then the presence in M of a graph every connected component of which is a tripod is a necessary condition for polynomial-time solvability of the maximum independent set problem in the class of M -free subcubic graphs (assuming $P \neq NP$). In [6], it was conjectured that this condition is also sufficient. Moreover, for graphs of bounded vertex degree the problem can be easily reduced to connected forbidden induced graphs, in which case the conjecture can be restated as follows.

Conjecture 1. *The maximum independent set problem is polynomial-time solvable for G -free subcubic graphs if and only if G is a tripod.*

One of the minimal non-trivial tripods is the claw $S_{1,1,1}$. The problem can be solved for the claw-free graphs in polynomial time even without the restriction to bounded degree graphs [10]. In [7], the result for claw-free graphs was extended to $S_{1,1,2}$ -free graphs, also known as fork-free graphs, and again without the restriction to bounded degree graphs. However, any further extension becomes much harder even for bounded degree graphs, and only recently a solution was found for $S_{2,2,2}$ -free subcubic graphs [9]. Currently, this is one of the few maximal subclasses of subcubic graphs with polynomial-time solvable independent set problem.

Now we turn to polynomial-time solutions of the second type, i.e. classes of graphs where forbidden induced subgraphs contain arbitrarily large chordless cycles. Clearly, in this case the set of forbidden induced subgraphs must be infinite. A typical example of this type deals with classes of bounded chordality, i.e. classes excluding all chordless cycle of length at least k for a constant k . Without a restriction to bounded degree graphs a solution of this type is known only for $k = 4$, i.e. for chordal graphs [4], and is unknown for larger values of k . Together with the restriction to bounded degree graphs bounded chordality implies bounded tree-width [2] and hence polynomial-time solvability of the maximum independent set problem. In other words, the problem can be solved for (C_k, C_{k+1}, \dots) -free graphs of bounded vertex degree for each value of $k \geq 3$.

An *apple* A_k , $k \geq 4$, is a graph formed of a chordless cycle C_k and an additional vertex, called the *stem*, which has exactly one neighbour on the cycle C_k . The class of (A_4, A_5, \dots) -free graphs generalizes both chordal graphs and claw-free graphs, and a solution for the maximum independent set problem in this class was presented in [3]. In case of bounded degree

graphs this solution can be extended to graphs without large apples, i.e. to (A_k, A_{k+1}, \dots) -free graphs of bounded vertex degree for any fixed value of k [8].

Generalizing both the subcubic graphs without large apples and $S_{2,2,2}$ -free subcubic graphs, in the present paper we prove polynomial-time solvability of the maximum independent set problem for subcubic graphs excluding large apples with a long stem. An *apple with a long stem* A_k^* is obtained from an apple A_k by adding one more vertex which is adjacent to the stem of A_k only. We show that for any fixed value of k , the maximum independent set problem in the class of $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graphs can be solved in polynomial time. Observe that this class contains all $S_{2,p,p}$ -free subcubic graphs for any fixed $p < k$ and hence our result brings us much closer to the proof of Conjecture 1.

2. Preliminaries

All graphs in this paper are simple, i.e. undirected, without loops and multiple edges. The vertex set and the edge set of a graph G is denoted by $V(G)$ and $E(G)$, respectively. The *neighbourhood* $N(v)$ of a vertex $v \in V(G)$ is the set of vertices of G adjacent to v . The *degree* of $v \in V(G)$ is the number of its neighbours, i.e. $|N(v)|$. As usual, P_n and C_n denote a chordless path and a chordless cycle with n vertices, respectively.

A subgraph of G induced by a subset $U \subseteq V(G)$ is denoted $G[U]$. If G contains no induced subgraphs isomorphic to a graph H , we say that G is H -free.

Outline of the proof. To prove polynomial-time solvability of the maximum independent set problem in the class of $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graphs,

1. We start by checking if the input graph G has an induced copy of $S_{2,2,2}$. If G is $S_{2,2,2}$ -free, then the problem can be solved for G in polynomial time [9]. Otherwise, we proceed to checking whether G has a induced cycle of length at least $p = 300k$. This can be done by inspecting all induced paths P with $p - 1$ vertices, deleting all the vertices of P and all their neighbours, except for the endpoints of P and the neighbours of the endpoints, and then checking if the endpoints belong to the same connected component in the resulting graph. If G does not contain induced cycles of length at least p , then the tree-width of G is bounded by a function of k [2] and hence the problem can be solved in polynomial time for G .
2. If G contains an induced copy of $S_{2,2,2}$ and a large induced cycle C , then in the absence of large induced apples with long stems we prove that it must contain a large extended cycle C^* , which is a graph obtained from C by adding two vertices that create a C_6 together with four consecutive vertices of C (see Fig. 8 in Section 4). This is shown in Section 3. An important ingredient of this proof is the assumption that the input graph G is connected and has no separating cliques, i.e. cliques whose removal disconnects the graph. A polynomial-time reduction of the maximum independent set problem to graphs without separating cliques can be found in [12,13].
3. After the previous two steps we can assume that our graph contains a large extended cycle. In Section 4 we show how to destroy such a large extended cycle by means of various local reductions. Each of them transforms G into a smaller graph G' in the same class with a fixed difference $\alpha(G) - \alpha(G')$. The set of reductions is described in Section 4.1 and their application to a graph G containing a large extended cycle is described in Section 4.2. By destroying the large extended cycle C^* , we destroy either the cycle C or the induced copy of $S_{2,2,2}$ (or both) and return to Step 1 to check if there are other copies of a large induced cycle or an induced $S_{2,2,2}$.

The main difficulties lie in the second step (showing that if the graph has an $S_{2,2,2}$ and a large induced cycle, then it has a long extended cycle), which is handled in Section 3; and in the third step (showing how to deal with a large extended cycle), which is handled in Section 4).

3. From large cycles to extended large cycles

We recall that C^* denotes an *extended cycle*, i.e. the graph obtained from a cycle C by adding two vertices that create a C_6 together with four consecutive vertices of C (see Fig. 8 in Section 4). Also, A_p^* denotes an *apple with a long stem*, where p stands for the size of the cycle in the apple. An apple with a long stem consisting of a cycle C and two stem vertices x, y will be denoted $C_{x,y}$.

The main goal of this section is to show that if G contains a large induced cycle and an induced copy of $S_{2,2,2}$, then it contains either a large induced extended cycle or a large induced apple with a long stem. This will be shown in two steps in Lemmas 1 and 2.

Lemma 1. *Let G be a subcubic graph without separating cliques. If G has an induced cycle C of length p and an induced copy of $S_{2,2,2}$, then G has an induced cycle of length at least $p/12$ containing the center of an induced $S_{2,2,2}$.*

Proof. Denote a copy of an induced $S_{2,2,2}$ by H . We denote the vertices of H by a_0 (the center), a_1, b_1, c_1 (vertices of degree 2), a_2, b_2, c_2 (vertices of degree 1 adjacent to a_1, b_1, c_1 , respectively). If vertex a_0 belongs to C , there is nothing to prove. We split the rest of the proof into cases depending on the distance from vertex a_0 to C .

Case 1. Assume first that vertex a_0 is of distance 1 from C . Without loss of generality we may suppose that $a_1 \in C$. Then a_2 also belongs to C due to the degree constraint. If at least one of b_1 or c_1 belongs to C , then an induced cycle of length at least $p/3$ containing a_0 can be easily found. Therefore, we assume that neither b_1 nor c_1 belongs to C . If at least one of b_1 or c_1 has no neighbours among the two vertices of the cycle to the left of a_1 and the two vertices of the cycle to the right of a_1 , then a_1 is the center of an induced $S_{2,2,2}$ belonging to C and the result holds. Therefore, we assume that b_1 has a neighbour on the cycle, which is of distance at most 2 from a_1 (along the cycle), and c_1 has a neighbour on the cycle, which is of distance at most 2 from a_1 (along the cycle). Since neither b_1 nor c_1 is adjacent to a_2 , we may assume that at least one of them, say b_1 , has a neighbour x on C non-adjacent to a_1 , i.e. of distance exactly 2 from a_1 . Denote the vertex of C between a_1 and x by y . If b_1 has no other neighbours on C , then a_0 belongs to an induced cycle C' of length $p+1$ obtained from C by replacing vertex y with vertices a_0 and b_1 . If b_1 has one more neighbour on C , say z , then z is different from y , since otherwise C' is still a large induced cycle containing a_0 , in which case the result holds. Therefore, z belongs to C' . The edge zb_1 splits C' into two induced cycles, one of which, say C'' , has length at least $p/2$. Vertex b_1 necessarily belongs to C'' and, unless a_0 also belongs to C'' , in which case the result holds, b_1 is the center of an induced $S_{2,2,2}$, with two branches belonging to C'' and one branch being a_0, a_1 . Indeed, the only obstacle for b_1 to be the center of an induced $S_{2,2,2}$ is the case, when z is adjacent to a_1 . But then $z \neq c_1$, $z \neq a_2$, implying $y = a_2$ and $x = c_1$, which is impossible.

Case 2. Assume now that a_0 is of distance 2 from C , and let, without loss of generality, x be a neighbour of a_1 on C . If a_1 has no other neighbour on C except for x , then x is the center of an induced $S_{2,2,2}$ belonging to C . If a_1 has two non-consecutive neighbours on C , then G has an induced cycle of length at least $p/2$ containing a_1 , and hence, according to Case 1, G has an induced cycle of length at least $p/6$ containing a_0 . If a_1 has two consecutive neighbours on C , then these neighbours together with a_1 create a clique. Therefore, there must exist a path connecting a_0 to C and avoiding this clique. But then G has an induced cycle of length at least $p/2$ containing a_0 .

Case 3. Assume that a_0 is of distance more than 2 from C . To prove the result in this case, we use the notion of a quasi-chord defined as follows. A *quasi-chord* for C is a chordless path $P = (p_1, \dots, p_s)$ such that each of p_1 and p_s has two consecutive neighbours on C , while the other vertices of P have no neighbours on C . Note that a quasi-chord P splits C into two parts one of which together with P creates an induced cycle of length at least $p/2$.

- (1) First, let us show that if a_0 is of distance more than 2 from C , then there exists at least one quasi-chord for C with the property that the distance between a_0 and this quasi-chord is strictly less than the distance between a_0 and C . Since G is connected, there must exist a path connecting a_0 to C . Let $P' = (x_1, \dots, x_p)$ be a shortest path between a_0 and C with $x_1 = a_0$ and with x_p having a neighbour in C . If x_p has a unique neighbour on C , then this neighbour is the center of an induced $S_{2,2,2}$, in which case we are done. If x_p has two non-consecutive neighbours on C , then x_p is the center of an induced $S_{2,2,2}$ that belongs to a cycle of length at least $p/2$. Therefore, x_p has two consecutive neighbours on C , say c_1 and c_2 . Then x_p, c_1, c_2 is a clique and hence there must exist a path connecting a_0 to C and avoiding this clique. Let $P'' = (y_1, \dots, y_t)$ be shortest path of this type with $y_1 = a_0$ and with y_t having a neighbour in C . Then by analogy with P' we conclude that y_t has two consecutive neighbours on C . We observe that the two paths P' and P'' may have common vertices different from a_0 . Also, there may exist chords (edges) between vertices of these paths. However, we can always find a chordless path P connecting x_p to y_t , which uses only the vertices of P' and P'' by considering a shortest path in $P' \cup P''$. Obviously, this path P is a quasi-chord for C . Moreover, P is closer to a_0 than C , since P contains x_p and by definition the distance between a_0 and x_p is exactly one less than the distance between a_0 and C .
- (2) Now let us show that there exists a quasi-chord for C which is of distance at most 2 from a_0 . To this end, let us denote by $P = (p_1, \dots, p_s)$ a quasi-chord for C which is as close to a_0 as possible. Assume a_0 is of distance more than 2 from P . Then we consider a shortest path $P' = (x_1, \dots, x_p)$ connecting $a_0 = x_1$ to P with x_p having a neighbour in P . Since P is closer to a_0 than C , no vertex of P' belongs to C or has a neighbour on C . Also, since P' is shortest, no vertex of P' has a neighbour on P , except for x_p . We recall that the quasi-chord P splits C into two parts one of which together with P creates an induced cycle C' of length at least $p/2$. To avoid an induced $S_{2,2,2}$ whose center is of distance at most 1 from C' (in which case we are in conditions of Case 1 with respect to the induced cycle C'), we conclude that x_p has two consecutive neighbours on P , say p_i and p_{i+1} . Using the fact G has no separating clique, we find one more chordless path $P'' = (y_1, \dots, y_t)$ connecting $a_0 = y_1$ to $C \cup P$ and avoiding the clique $\{x_p, p_i, p_{i+1}\}$. As before, we assume that no vertex of P'' except for y_t has a neighbour on $C \cup P$. Let P^* be a chordless path connecting x_p to y_t and consisting of the vertices of P' and P'' only. If y_t has no neighbours on C , then a part of P can be replaced by P^* creating a quasi-chord for C , which is closer to a_0 than P , since it contains x_p . This contradicts the choice of P . Therefore, y_t has a neighbour on C . To avoid an induced $S_{2,2,2}$ whose center is of distance at most 1 from C , we conclude that y_t has two consecutive neighbours on C . But then $\{p_1, \dots, p_i\} \cup P^*$ is a quasi-chord for C , which is closer to a_0 than P , since it contains x_p . This final contradiction shows that the distance from a_0 to P is at most 2.

From Claim (2) we conclude that G has an induced cycle of length at least $p/2$, which is of distance at most 2 from a_0 . Therefore, according to Cases 1 and 2, G has an induced cycle of length at least $p/12$ containing the center of an induced $S_{2,2,2}$. \square

Lemma 2. Let G be a subcubic graph without separating cliques. If G has an induced cycle C of length p containing the center of an induced $S_{2,2,2}$, then G has an induced extended cycle C_t^* or an induced apple with a long stem A_t^* with $t \geq p/8$.

Proof. Denote the induced $S_{2,2,2}$ with the center on C by H . We also denote the vertices of H by 0 (the center), 1, 2, 3 (vertices of degree 2), 4, 5, 6 (vertices of degree 1 adjacent to 1,2,3, respectively) and assume without loss of generality that 1 and 2 belong to the cycle (together with 0), while 3 does not belong to C . Finally, we denote the two vertices of C following vertex 1 by a and b (possibly $a = 4$), and the two vertices of C following vertex 2 by c and d (possibly $c = 5$). Let us call the number of edges of H contained in C the H -value of C . Clearly, this value cannot be large than 4 and due to degree constraint it cannot be smaller than 2.

We prove the lemma through a series of claims. The proof of the first claim is evident.

Claim 1. If a vertex $x \notin C$

- has two neighbours on C , then the smaller of the two cycles formed by x and the two parts of C has size at most 5,
- has three neighbours on C , then the two smaller cycles (of the three formed by x and the three parts of C) have size at most 4,

since otherwise an induced A_t^* with $t \geq p/3$ can be easily found.

Claim 2. Either vertex 6 belongs to C or 6 has a neighbour on C or 6 is adjacent to a vertex x that has a neighbour on C .

Proof. Assume vertex 6 neither belongs to C nor has a neighbour on C . Then vertex 3 must have a neighbour on C , for otherwise $C_{3,6}$ is an induced A_p^* . By Claim 1 (and the fact that $6 \notin C$), vertex 3 cannot have neighbours outside of $\{a, b, c, d\}$. Also, since $6 \notin C$, vertex 3 cannot have neighbours both in $\{a, b\}$ and in $\{c, d\}$. Therefore, we may assume without loss of generality that 3 has a neighbour in $\{a, b\}$ and does not have neighbours in $\{c, d\}$.

Since G has no separating cliques, vertex 6 must have a neighbour x different from 3. Then x has a neighbour on C , for otherwise $C'_{6,x}$ is an induced A_{p-1}^* , where C' is the cycle formed by 3 and the long part of C between 0 and the neighbour of 3 in $\{a, b\}$. \square

Claim 3. If the H -value of C is 4, i.e. $4 = a$ and $5 = c$, then G has an induced C_t^* or an induced A_t^* with $t \geq p/2$.

Proof. We split the proof into the following cases:

- (1.1) If vertex 3 is adjacent neither to b nor to d , then 3 has no other neighbours on C (except for 0) by Claim 1. In particular, $6 \notin C$. But then 6 has a neighbour on C , since otherwise $C \cup \{3, 6\}$ is an induced A_p^* . If 6 has two neighbours on C or if 6 has a single neighbour different from b and d , then G contains an induced A_t^* with $t \geq p/2$. If b (or d) is the only neighbour of 6 on C , then $C \cup \{3, 6\}$ is an induced C_p^* .
- (1.2) Now assume without loss of generality that 3 is adjacent to b . Then 3 has no other neighbours on C (except for 0 and b) by Claim 1. In particular, $6 \notin C$. If 6 has neighbours on C , we denote them by y and z (possibly $y = z$). Otherwise, by Claim 2, 6 is adjacent to a vertex x which has neighbours on C . In this case, we denote the neighbours of x on C by y and z (again, possibly $y = z$). The neighbours of 3 on C split the cycle into two paths: the short one of length 3 and a longer one of length $p - 3$. Let us denote the longer path by P .
 - (1.2.1) Suppose first that neither y nor z belongs to $\{1, 4\}$. Then y and z split P into at most 3 paths giving rise to at most 3 cycles. It is not difficult to see that each of these cycles is part of an apple with a long stem, and one of these cycles has length at least $p/2$ (remember that the distance between y and z along C cannot be larger than 3 by Claim 1), i.e. G contains an induced A_t^* with $t \geq p/2$.
 - (1.2.2) Now assume without loss of generality that $y \in \{1, 4\}$. Then y is a neighbour of x , since 6 has no neighbours among $\{1, 4\}$. Also, z belongs P , since otherwise $P \cup \{3, 6, x\}$ induce an A_{p-2}^* . Vertex z splits P into two internally disjoint subpaths P^1 (from b to z) and P^2 (from z to 0). Then either the cycle $C^1 = P^1 \cup \{3, 6, x\}$ or the cycle $C^2 = P^2 \cup \{3, 6, x\}$ has size at least $p - 1$ (remember that the distance between y and z along C cannot be larger than 3 by Claim 1). If C^1 has size at least $p - 1$, then
 - * either $y = 1$, which implies $C^1 \cup \{1, 4\}$ is an induced C_t^* with $t \geq p - 1$,
 - * or $y = 4$, which implies by Claim 1 that $z = 2$ and hence $(C - \{1\}) \cup \{3, x\}$ is an induced C_t^* with $t \geq p - 1$.
 If C^2 has size at least $p - 1$, then
 - * either $y = 4$, implying $C^2 \cup \{1, 4\}$ is an induced C_t^* with $t \geq p - 1$,
 - * or $y = 1$, implying $z \in P$ is adjacent to b by Claim 1 and consequently $P \cup \{1, 3, b, x\}$ is an induced C_t^* with $t \geq p - 1$. \square

Claim 4. If the H -value of C is 3 with $4 \neq a$, then G contains either an induced cycle of length at least $p/2$ with H -value 4 or an induced A_t^* with $t \geq p/2$ or an induced C_p^* .

Proof. To prove the claim, we observe the following:

- a is adjacent to 3. Indeed, if a is adjacent neither to 3 nor to 6, then by replacing 4 with a we obtain an induce $S_{2,2,2}$ containing 4 edges in C . If a is not adjacent to 3 and is adjacent to 6, then
 - if $6 \in C$, i.e. if $6 = b$, then 3 has no other neighbours on C by Claim 1, in which case we obtain an induced cycle of length $p - 1$ containing four edges $(0, 3), (3, 6), (0, 2), (2, 5)$ of H .
 - if $6 \notin C$, then any other neighbour of 6 on C (if any) must be of distance at most 3 from a by Claim 1, in which case we obtain an induced cycle of length at least $p - 2$ containing four edges $(0, 3), (3, 6), (0, 2), (2, 5)$ of H .
- 6 does not belong to C . Indeed, if 6 belongs to C , then it must be of distance at most 2 from a (Claim 1), in which case we obtain an induced cycle of length at least $p - 2$ containing four edges $(0, 3), (3, 6), (0, 2), (2, 5)$ of H .
- 4 has no neighbours on C different from 1. Indeed, if 4 has more neighbours on C , then by Claim 1 the farthest neighbour must be of distance at most 3 (if $|N(4) \cap C| = 2$) or at most 4 (if $|N(4) \cap C| = 3$) from 1, in which case we obtain an induced cycle of length at least $p - 2$ containing four edges $(0, 1), (1, 4), (0, 2), (2, 5)$ of H .
- 4 is adjacent to a vertex $u \notin C$ that has a neighbour on C , since otherwise either 1 is a separating clique or $C_{4,u} = A_p^*$.
- u is adjacent to 2. Indeed, assume u is not adjacent to 2. By Claim 1 the neighbours of u on C must be of distance at most 3 from each other. Therefore, 4, u together with a large part of C between 1 and a neighbour of u on C create an induced cycle C' of length at least $p/2$ containing either the edge $(0, 1)$ or the edge $(1, a)$. If $(0, 1) \in C'$, then the H -value of C' is 4 (it contains the four edges $(0, 1), (1, 4), (0, 2), (2, 5)$). If $(1, a) \in C'$, then either $C \cup \{4, u\} = C_p^*$ (if 5 is the only neighbour of u on C) or $C'_{0,2} = A_t^*$ with $t \geq p/2$ (if u is adjacent to a vertex of C different from 5).
- the degree of 4 in G is 2, since otherwise any other neighbour of 4 is also adjacent to 2, similarly to the previous claim, which is impossible due to degree constraint.

From Claim 2 and the above discussion we conclude that either 6 has neighbours on C or 6 is adjacent to a vertex x that has neighbours on C . These neighbours (of 6 or of x) must be located on C at distance at most 3 from each other (Claim 1) and must be different from 0, 1, a . Therefore, vertices 3, 6 (and possibly x) together with a long part of C create an induced cycle C' of length at least $p/2$ containing either the edge $(0, 3)$ or the edge $(3, a)$. In both cases, $C'_{1,4} = A_t^*$ with $t \geq p/2$, unless x is adjacent to 4. But in the latter case $x = u$ (where u is the neighbour of 4 adjacent to 2) and hence $C_{x,6} = A_p^*$. \square

Claim 5. If the H -value of C is 2, then G contains either an induced cycle of length at least $p/2$ with H -value 3 or an induced A_t^* with $t \geq p/2$ or an induced C_{p-1}^* .

Proof. Let us first establish that both a and c must have a neighbour in $\{3, 6\}$. Suppose that c does not have a neighbour in $\{3, 6\}$. We now observe that a must have a neighbour in $\{3, 6\}$, for otherwise we exchange $\{4, 5\}$ with $\{a, c\}$ and obtain H -value 4. Then, if c is not adjacent to 4, we can exchange c with 5 in H to increase the H -value of C . Suppose then that c is adjacent to 4. By Claim 1 c is the only neighbour of 4 on C . If 5 has no neighbour in C besides 2, then we have an induced $A_{p-1}^* = C \cup \{4, 5\} \setminus \{0\}$. If 5 is adjacent to a , then by Claim 1 a and 2 are the only neighbours of 5 on C , in which case we obtain an induced $C_{p-1}^* = C \cup \{4, 5\} \setminus \{0\}$. If 5 is not adjacent to a , then any neighbour of 5 on C must be of distance at most 3 from c . We therefore have an induced cycle of length at least $p - 2$ that goes through 0, 2, 5, then continues to the neighbour of 5 on C that is as far as possible from c , and then goes on to 0 by using vertices of C and avoiding c . This induced cycle has a higher H -value, as it contains edges $(0, 1), (0, 2)$ and $(2, 5)$.

Because of the above, both a and c must have a neighbour in $\{3, 6\}$. They cannot both be adjacent to 6 by Claim 1, and they cannot both be adjacent to 3, for otherwise, without loss of generality, $6 = a$ and hence 6 is adjacent to 1, which is impossible. Therefore, we assume, without loss of generality, that a is adjacent to 6 and c is adjacent to 3. We observe that 6 does not belong to C by Claim 1.

If 6 has a neighbour x on C different from a , then by Claim 1 x must be of distance at most 3 from a , in which case we obtain an induced cycle C' of length at least $p - 4$ containing 3, 6 and a long part of C between c and x , and an induced $A_t^* = C'_{0,1}$ with $t \geq p - 4$.

Therefore, we assume a is the only neighbour of 6 on C and look at vertex 4. If 4 has no other neighbours on C different from 1, then we have an induced $A_{p-1}^* = C \cup \{3, 4, 6\} \setminus \{0, 2\}$. If 4 has neighbours on C different from 1, then these neighbours must be close to 1 (by Claim 1), in which case we find an induced cycle of length at least $p - 2$ containing both edges $(0, 1)$ and $(1, 4)$ and hence having H -value at least 3. \square

Summarizing, we conclude that if G has an induced cycle C of length p containing the center of an induced $S_{2,2,2}$, then G contains either an induced C_t^* or an induced A_t^* with $t \geq p/8$. \square

4. Destroying large extended cycles

According to the previous section, if an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G contains a large induced cycle and an induced copy of $S_{2,2,2}$, then it must contain a large extended cycle C^* . The goal of the present section is to show how to destroy large extended cycles by means of various local graph reductions. We describe these reductions in Section 4.1 and apply them to large extended cycles in Section 4.2.




Fig. 2. Φ -reduction.




Fig. 3. Π -reduction.

4.1. Graph reductions

4.1.1. Φ -reduction and house-reduction

We start with the Φ -reduction introduced in [9]. It applies to a graph G containing an induced copy of the graph Φ represented on the left of Fig. 2 and consists in replacing Φ by the graph on the right of Fig. 2.

Lemma 3. *By applying the Φ -reduction to an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G , we obtain an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G' with $\alpha(G') = \alpha(G) - 2$.*

Proof. The equality $\alpha(G') = \alpha(G) - 2$ was proved in [9]. To prove $(A_k^*, A_{k+1}^*, \dots)$ -freeness of G' assume by contradiction that G' contains an induced copy H of a large apple with a long stem A_p^* with $p \geq k$. Then at least one of the four vertices a, b, c, d , say d , does not belong to H , since otherwise H contains a C_4 . But then the vertices inducing H together with vertices 1 and 2 induce a subdivision of A_p^* in G , which is impossible. \square

A *house* is the complement of a P_5 . If a graph G contains an induced *house*, the *house*-reduction consists in removing from G the vertices that form a triangle in the *house*. It was shown in [9] that if G is a subcubic graph, then the *house*-reduction reduces $\alpha(G)$ by exactly 1.

4.1.2. Π -reduction

Now we introduce the Π -reduction illustrated in Fig. 3. In a graph G , an induced Π is the graph represented on the left of Fig. 3. We observe that vertex f can be missing, in which case vertices a and c have no other neighbours in G . Similarly, vertex h can be missing, in which case vertices b and d have no other neighbours in G .

Lemma 4. *By applying the Π -reduction to an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G , we obtain an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G' with $\alpha(G') = \alpha(G) - 2$.*

Proof. Let S be a maximum independent set in G and $X = S \cap \{1, 2, 3, 4\}$, $Y = S \cap \{a, b, c, d\}$. If $|Y| = 4$, then $|X| = 0$ and hence $S - \{a, c\}$ is an independent set in G' of size $\alpha(G) - 2$. If $|Y| = 3$, say $Y = \{a, b, c\}$, then $X = \{4\}$ and hence $S - \{4, b\}$ is an independent set in G' of size $\alpha(G) - 2$.

Let $|Y| = 2$. Then, up to symmetry, $Y = \{a, b\}$ or $Y = \{a, d\}$ or $Y = \{a, c\}$. If $Y = \{a, b\}$ or $Y = \{a, d\}$, then $|X| = 1$ and hence $S - (X \cup \{a\})$ is an independent set in G' of size $\alpha(G) - 2$. If $Y = \{a, c\}$ or $|Y| \leq 1$, then $|X| = 2$ and hence $S - X$ is an independent set in G' of size $\alpha(G) - 2$. Therefore, $\alpha(G') \geq \alpha(G) - 2$.

Conversely, let S be a maximum independent set in G' and $Y = S \cap \{a, b, c, d\}$. Clearly, $|Y| \leq 2$ and if $|Y| = 2$ we can assume without loss of generality that $Y = \{a, d\}$ or $Y = \{a, c\}$.

If $|Y| \leq 1$, then S can be always extended by adding two vertices from $\{1, 2, 3, 4\}$ to an independent set of size $\alpha(G') + 2$ in G . If $Y = \{a, d\}$, then $g, h \notin S$ and hence $S \cup \{b, 3\}$ is an independent set of size $\alpha(G') + 2$ in G . If $Y = \{a, c\}$, then $S \cup \{2, 4\}$ is an independent set of size $\alpha(G') + 2$ in G . Therefore, $\alpha(G') + 2 \leq \alpha(G)$ and hence $\alpha(G') = \alpha(G) - 2$.

To prove $(A_k^*, A_{k+1}^*, \dots)$ -freeness of G' assume by contradiction that G' contains an induced copy H of a large apple with a long stem A_p^* with $p \geq k$. Then H contains at least one of the edges ab and cd , since otherwise the same vertices induced H in G .

If H contains both edges ab and cd , then exactly one of the vertices e, f, g, h belongs to H . Indeed, if two of these vertices belong to H , then H contains a small cycle, and if none of them belongs to H , then H is not connected. Assuming,




Fig. 4. Γ -reduction.




Fig. 5. Θ graph.

without loss of generality, that e belongs to H , we conclude that H contains two vertices b and d , each of which has degree 1 in H , which is impossible.

Now assume the edge ab belongs to H and the edge cd does not. Without loss of generality, $d \notin V(H)$. If $c \in V(H)$, then $e \notin V(H)$ or $f \notin V(H)$, since otherwise H contains a C_4 , say $f \notin V(H)$. Then c has degree 1 in H and hence vertex a must have degree 3 in H . But a has degree 2 in H , a contradiction.

If $c \notin V(H)$, then the vertices inducing H together with vertices 1 and 2 induce a subdivision of A_p^* in G , which is impossible. \square

4.1.3. Γ -reduction

One more reduction is illustrated in Fig. 4. We will refer to it as Γ -reduction. Again, vertex f can be missing, in which case vertices b and d have degree 2 in the graph.

Lemma 5. *By applying the Γ -reduction to an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G , we obtain an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph G' with $\alpha(G') = \alpha(G) - 2$.*

Proof. Let S be a maximum independent set in G and $X = S \cap \{1, 2, 3, 4\}$, $Y = S \cap \{a, b, c, d\}$. If $|Y| = 3$, say $Y = \{a, b, d\}$, then $X = \{3\}$ and hence $S - \{a, 3\}$ is an independent set in G' of size $\alpha(G) - 2$.

Let $|Y| = 2$. Then, up to symmetry, $Y = \{a, b\}$ or $Y = \{a, d\}$ or $Y = \{b, d\}$. If $Y = \{a, b\}$ or $Y = \{a, d\}$ then $|X| = 1$ and hence $S - (X \cup \{a\})$ is an independent set in G' of size $\alpha(G) - 2$. If $Y = \{b, d\}$ or $|Y| \leq 1$, then $|X| = 2$ and hence $S - X$ is an independent set in G' of size $\alpha(G) - 2$. Therefore, $\alpha(G') \geq \alpha(G) - 2$.

Conversely, let S be a maximum independent set in G' and $Y = S \cap \{a, b, c, d\}$. Clearly, $|Y| \leq 2$ and if $|Y| = 2$ we can assume without loss of generality that $Y = \{a, d\}$ or $Y = \{b, d\}$.

If $|Y| \leq 1$ or $Y = \{b, d\}$, then S can be extended by adding two vertices from $\{1, 2, 3, 4\}$ to an independent set of size $\alpha(G') + 2$ in G . If $Y = \{a, d\}$, then $e, f \notin S$ and hence $S \cup \{b, 3\}$ is an independent set of size $\alpha(G') + 2$ in G . Therefore, $\alpha(G') + 2 \leq \alpha(G)$ and hence $\alpha(G') = \alpha(G) - 2$.

To prove $(A_k^*, A_{k+1}^*, \dots)$ -freeness of G' assume by contradiction that G' contains an induced copy H of a large apple with a long stem A_p^* with $p \geq k$. Then H contains at least one of the edges ab and cd , since otherwise the same vertices induced H in G .

If H contains both edges ab and cd , then neither e nor f belongs to H , since otherwise H contains a small cycle. Then both b and d have degree 1 in H , which is impossible.

If the edge ab belongs to H and the edge cd does not, then the vertices inducing H together with vertices 1 and 2 induce a subdivision of A_p^* in G , which is impossible. \square

4.1.4. Θ -reduction

Lemma 6. *If a subcubic graph G contains an induced Θ (see Fig. 5), then the deletion of vertices x, y reduces the independence number of G by exactly 1.*

Proof. Let S be a maximum independent set in G . Clearly, S contains at most one vertex in $\{x, y\}$. Now let us show that S contains at least one vertex in $\{x, y\}$. Assume $x, y \notin S$. Then x contains a neighbour in S and y contains a neighbour in S . Without loss of generality let $2, 5 \in S$. But then S is not maximum, since by replacing 2 with 3, x we obtain a larger independent set. \square




Fig. 6. Total structure of G with $H = G[1, 2, 3]$ and $\ell = 2$.




Fig. 7. Graphs A_1 , A_2 and A_3 .

4.1.5. Total structure and subgraph reduction

The idea of total structure was introduced in [1] and can be described as follows. Assume the vertices of a graph G are labelled by integers $1, 2, \dots, n$, where $n = |V(G)|$. The vertex with maximum label in a subset $A \subseteq V(G)$ is denoted $m(A)$, and A^- is defined to be $A - \{m(A)\}$.

Let H be an induced subgraph of G and $[H]$ the subgraph of G induced by the closed neighbourhood of H , i.e. by the set containing all the vertices of H and all their neighbours. We choose a positive integer number $\ell \leq \alpha(H)$, denote $R := V - V([H])$ and associate with the triple (G, H, ℓ) a graph $S(G, H, \ell)$ as follows:

- the vertex set of $S(G, H, \ell)$ is $R \cup W$, where W is the family of all independent sets of cardinality $\ell + 1$ in $[H]$;
- the edge set of $S(G, H, \ell)$ consists of
 - the edges of the subgraph $G[R]$,
 - the edges linking vertices $A \in W$ and $B \in W$ whenever $A^- \neq B^-$ or $m(A)m(B) \in E(G)$,
 - the edges linking a vertex $A \in W$ to a vertex $v \in R$ whenever v has a neighbour in the subset A in the graph G .

The transformation of G into $S(G, H, \ell)$ was called [1] total structure. An example of total structure is given in Fig. 6. The importance of this transformation is due to the following theorem proved in [1].

Theorem 2. $\alpha(S(G, H, \ell)) = \alpha(G) - \ell$.

Directly from this theorem we derive the following corollary.

Corollary 1. If $\alpha(H) = \alpha([H])$, then $\alpha(G - [H]) = \alpha(G) - \alpha(H)$.

The deletion of $[H]$ in the case when $\alpha(H) = \alpha([H])$ was called in [9] the H -subgraph reduction. It is not difficult to check that if A_1, A_2 , or A_3 (see Fig. 7) is an induced subgraph of a subcubic graph, then

- $\alpha(A_1[2, 3, 5, 6, x, y]) = \alpha(A_1) = 3$,
- $\alpha(A_2[1, 2, 3, 5, 6, x, y]) = \alpha(A_2) = 4$,
- $\alpha(A_3[1, 2, 3, 5, 6, x, y]) = \alpha(A_3) = 4$.

Lemma 7. If A_1, A_2 , or A_3 is an induced subgraph of a subcubic graph G , then $\alpha(G - A_1) = \alpha(G) - 3$, $\alpha(G - A_2) = \alpha(G) - 4$, $\alpha(G - A_3) = \alpha(G) - 4$.




Fig. 8. An extended cycle.

4.2. Applying graph reductions to large extended cycles

Let G be an $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graph. For ease of terminology and notation we will refer to any A_t^* with $t \geq k$ simply as a large apple with a long stem. According to Section 3, we may assume that G contains a large extended cycle C_p^* , i.e. a graph that consists of an induced cycle of length p , plus two extra vertices which form a C_6 together with four consecutive vertices of the cycle and have no other neighbours in C_p^* . We denote the vertices of an extended cycle as shown in Fig. 8, where we have given labels to the vertices of the C_6 , plus some other interesting vertices. In the remainder we use simply C^* to denote the extended cycle and C_6 to denote the set of vertices $\{1, 2, 3, 4, 5, 6\}$. Without loss of generality, we assume that $p \geq 3k$.

We will now go through a sequence of cases that covers all possible ways in which C^* may be connected to the rest of the graph.

Case 0: Vertices 2 and 3 both have degree 2 in G . In this case we delete 2, 3 from the graph and add the edge connecting 1 to 4. This decreases $\alpha(G)$ by exactly 1, which is easy to see and also follows from total struction applied with $H = G[2]$ and $\ell = 1$. Also, it is not difficult to check that this transformation does not create any new forbidden induced subgraphs.

Because of the above we can assume that the set $\{2, 3\}$ has a neighbour outside of C^* . We call this vertex x . Without loss of generality we assume that x is connected to 2. Let us consider how x is connected to the rest of C^* .

Case 1.1: $N(x) \cap C^* = \{2\}$. This case leads to a contradiction, as $C^* \cup \{x\} \setminus \{3\}$ is a large apple with a long stem.

Case 1.2: $N(x) \cap C_6 = \{2\}$ and x has exactly one neighbour in $C^* \setminus C_6$. Let f be that neighbour (which may coincide with one of a, b, c, d). Then the graph contains a large apple with a long stem: the stem is made up of $\{5, 6\}$, and the cycle from 2, x, f , plus either the path from f to 1, or the path from f to 3 in C^* (whichever is longer).

Case 1.3: $N(x) \cap C_6 = \{2\}$ and x has two neighbours in $C^* \setminus C_6$. Let f, g be these neighbours. If the distance from f to g in C^* is at least $p/3$, then we have a large apple with a long stem: the cycle is x, f, g plus the path from f to g in C^* and the stem is 2, 3. Otherwise, one of f, g has a path of length at least $p/3$ to 1 or 4 in C^* which does not contain the other vertex from $\{f, g\}$, so we find a large apple with a long stem as in Case 1.2.

From the above cases we conclude that x has at least two neighbours in C_6 . Since the degrees of 1, 4 are already three in C^* , we conclude that x has at least two neighbours in $\{2, 3, 5, 6\}$. Let us also rule out two further cases.

Case 1.4: $N(x) \cap C_6 = \{2, 3\}$. If the degree of x is 2, then we can apply the H -subgraph reduction (Corollary 1) with $H = G[\{x\}]$, which decrease $\alpha(G)$ by 1. If x has a neighbour outside of C^* , we find a large apple with a long stem: $1654cd \dots ba1$ and 2, x . Finally, if x has a neighbour f in $C^* \setminus C_6$, we find a large apple with a long stem as in Case 1.3, where $\{5, 6\}$ is the stem and the cycle is formed by 2, x, f , plus either the path from f to 1, or the path from f to 3 in C^* (whichever is longer).

Case 1.5: $|N(x) \cap C_6| = 3$. Here we can assume without loss of generality that $N(x) \cap C_6 = \{2, 3, 5\}$, as other cases are isomorphic to this. Then, $\{2, 3, 4, 5, x\}$ induces a house and we can apply the house-reduction.

We are now ready to state the two main cases of our analysis:

Lemma 8. *If one of Cases 1.1-1.5 applies, then the instance can be simplified in polynomial time. If none of Cases 1.1-1.5 applies, then either $N(x) \cap C_6 = \{2, 5\}$ or $N(x) \cap C_6 = \{2, 6\}$.*

We handle these two cases separately in the following subsections.

4.2.1. x is adjacent to 2 and 6

Lemma 9. *Let x be a vertex adjacent to 2 and 6 and assume x has a neighbour y not in C^* . Then G contains an induced Φ or an induced Π or an induced Γ or an induced Θ .*




Fig. 9. Vertex y has neighbours on C .

Proof. If y is adjacent to 3, then by Lemma 8 (and symmetry) y is also adjacent to 5 and hence vertices 1, 2, 3, 4, 5, 6, x , y induce a Θ .

If y is adjacent to c , then vertices 2, 3, 4, x , y , c create a cycle of length 6 which, together with the path $1ab\dots d$ gives a second large extended cycle. Therefore, by Lemma 8 applied to this extended cycle, vertex 5 must be adjacent to y and hence vertices 1, 2, x , 6, y , 5, c , 4 induce a Φ .

If y is adjacent to a , then vertices a , y , 1, 2, x , 6, 3, 4, 5 induce a Γ with a possible missing common neighbour of 3 and 5 (any neighbour of these vertices must be common by Lemma 8).

If y is adjacent to b and not adjacent to a , then vertices a , b , y , 1, 2, x , 6, 3, 4, 5 induce a Π with a possible missing common neighbour of 3 and 5 (any neighbour of these vertices must be common by Lemma 8), and a possible missing common neighbour of a and y (any neighbour of these vertices must be common by Lemma 8 applied to the cycle C_6 formed by b , a , 1, 2, x , y).

From now on, we assume y has no neighbours in $\{3, 5, a, b, c\}$. If y has neighbours on $C^* \setminus C_6$, then we can distinguish at most 3 cycles containing y as shown in Fig. 9 (if y has only 1 neighbour on $C^* \setminus C_6$, the cycle C_2 is missing).

We observe that at least one of the cycles C^1 , C^2 , C^3 is large, i.e. has length at least $p/3$. Then G contains a large apple with a long stem

- $C^* \cup \{x, y\} \setminus \{5, 6\}$ if y has no neighbours on $C^* \setminus C_6$,
- $C^1 \cup \{3, 4\}$ if C^1 is large,
- $C^2 \cup \{x, 2\}$ if C^2 is large,
- $C^3 \cup \{1, a\}$ if C^3 is large.

A contradiction in all cases shows that y has a neighbour in $\{3, 5, a, b, c\}$ and hence G contains an induced Φ or an induced Π or an induced Γ or an induced Θ \square

We therefore find ourselves in the following context: $N(x) \cap C_6 = \{2, 6\}$ and $N(x) \setminus C^* = \emptyset$. Before we proceed to the rest of the analysis, let us identify another relevant vertex. If 3 has a neighbour outside C^* we call that vertex y . By Lemma 8 (and appropriate symmetry) y is also connected to 5. We have also argued that x and y are not adjacent. We will in the remainder assume that the degree of x is at least as large as the degree of y . This is without loss of generality, as the two vertices can be exchanged by an appropriate automorphism of C^* . In what follows, we analyze all possible adjacencies of x and y to the vertices of C^* .

Case 2.1: If x has degree 2 and y does not exist (therefore 3, 5 have degree 2), then we apply the H -subgraph reduction (Corollary 1) with $H = G[\{x, 1, 3, 5\}]$, in which case $\alpha(H) = \alpha([H]) = 4$ and hence the removal of $[H]$ decreases $\alpha(G)$ by 4.

Case 2.2: Assume x has degree 2 and y exists (therefore, y is connected to 3, 5). We have assumed without loss of generality that x has at least as high degree as y , therefore y has no other neighbour. We delete from the graph vertices 2, 3, x , y . If G' is the new graph, we claim that $\alpha(G') = \alpha(G) - 2$. The inequality $\alpha(G') \geq \alpha(G) - 2$ is clear, since no independent set can take more than two of the deleted vertices. To see that $\alpha(G) \geq \alpha(G') + 2$, take a maximum independent set in G' . If it contains vertex 5, then it does not contain 4 or 6. Therefore, we can augment it with $x, 3$. If it contains 6, we can augment it similarly by adding $y, 2$. Finally, if it contains neither 5 nor 6, we augment it with x, y .

Case 2.3: If x is connected to a , then the application of total struction with $H = G[\{1, x\}]$ and $\ell = 2$ replaces $\{x, 1, 2, 6, a\}$ with a new vertex adjacent to $b, 3, 5$ and decreases $\alpha(G)$ by 2. It is not difficult to check that this transformation does not introduce any new large apples with long stems.

Case 2.4: If x is connected to c , then $x61ab\dots cx$ together with 3, 4 form a large apple with a long stem.

Case 2.5: If x is connected to d , then $x21ab\dots dx$ together with 3, 4 form a large apple with a long stem.

Case 2.6: If x is connected to a vertex f of C^* in the path from b to d (but not b or d), then: if f is closer to a than to c , we take the path $xf\dots dc432x$ plus 1, a ; otherwise we take $xf\dots ba12x$ plus 3, 4. In both cases these form a large apple with a long stem.

Case 2.7: If x is connected to b and y does not exist, then we can apply the H -subgraph reduction with $H = G[\{x, 1, 3, 5\}]$. It is not difficult to check that $\alpha(H) = \alpha([H]) = 4$ and hence the removal of $[H]$ decreases $\alpha(G)$ by 4.

Case 2.8: Assume x is connected to b , y exists and it has degree 2 (that is, y is connected only to 3, 5). We then apply total struction with $H = G[\{y, 2, 6\}]$ and $\ell = 3$, which replaces $[H] = G[\{x, y, 1, 2, 3, 5, 6\}]$ with a new vertex z adjacent to $a, b, 4$ and decreases $\alpha(G)$ by 3. The new graph does not have a large apple with a long stem that uses z and both a, b , since that would induce a triangle. If, on the other hand, it has an apple with a long stem that uses z and at most two of its neighbours, then G also has a sub-divided copy of the same subgraph if we replace z with 1, 2, 3.

Case 2.9: Finally, suppose x is connected to b , y exists and y has degree 3. Since x and y have the same degree, we may exchange their roles, and by symmetry and the same case analysis that we did for x we conclude that y must be connected to d (otherwise one of the previous cases applies). We transform the graph as follows: we delete the vertices 1, 2, 3, 4, 5, 6, x, y and add two new vertices z, w such that z, w are connected to each other, z is connected to a, b , and w is connected to c, d . We claim that $\alpha(G') = \alpha(G) - 3$. First, to obtain $\alpha(G') \geq \alpha(G) - 3$, take a maximum independent set of G . If it contains a vertex from a, b and a vertex from c, d , then it contains at most three of the deleted vertices, since the six deleted vertices which are not adjacent to a vertex of the independent set induce a C_6 . In all other cases, the independent set in G contains at most four of the deleted vertices. However, if the set does not contain any of a, b , we can augment it with z in G' , while if it does not contain any of c, d we can add to it w . To see that $\alpha(G) \geq \alpha(G') + 3$, take a maximum independent set in G' . If it is using z , then it does not contain a or b . In G we replace z with 1, $x, 3, 5$. The situation is symmetric if the set contains w . Finally, if it does not contain either z or w , we observe that deleting the neighbours of the set among the removed vertices gives a C_6 , of which we can select three vertices. The transformation does not introduce a new large apple with a long stem, since the closed neighbourhoods of z, w include a triangle, therefore if one or two of these vertices is used in the apple we can replace them with an appropriate induced path through the deleted vertices in G .

4.2.2. x is adjacent to 2 and 5

Lemma 10. *Let x be a vertex adjacent to 2 and 5 and assume x has a neighbour y not in C^* . Then G contains an induced A_1 or an induced A_2 or an induced A_3 (Fig. 7).*

Proof. If y is adjacent to 3 or 6, then y is adjacent to both 3 and 6 (Lemma 8) and hence G contains an induced A_1 .

Assume y is adjacent to a . Then, if all three vertices 3, 6, y have degree 2 in G , then G contains an induced A_2 . If vertex 3 has degree three, it has a common neighbour with 6 (by Lemma 8), call this neighbour z . We claim that z must also be connected to y , which will give an induced A_3 . To see this, consider the set of vertices $(C^* \setminus \{2, 3\}) \cup \{x, y\}$. This set induces an extended cycle, where the C_6 is now formed by $a, 1, 6, 5, x, y$. Since z is connected to 6, it must be connected to one of $\{x, y\}$ (Lemma 8). However, x already has three neighbours (2, 5, y), therefore, z is connected to y .

If y is adjacent to c , this is symmetric to y being adjacent to a . So, we suppose that y is adjacent to none of 3, 6, a, c . The rest of the proof is similar to that of Lemma 9 with the only difference that if y is adjacent only to b this time we can find a large apple with a long stem, where the stem is $\{1, 6\}$ and the cycle goes through $byx234cd \dots b$. \square

Lemma 11. *Let x be a vertex adjacent to 2 and 5 and assume x has a neighbour in $C^* \setminus C_6$. Then this neighbour is one of a and c .*

Proof. Let f be the neighbour of x in $C^* \setminus C_6$ different from a and c (note that f may coincide with b or d). Suppose that the distance in $C^* \setminus C_6$ from f to a is at least as large as the distance from f to c (the other case is symmetric). We take the cycle $xf \dots a12x$ and the stem $\{3, 4\}$ to form a large apple with a long stem. \square

Because of symmetry, we will in the remainder assume without loss of generality that if x has a neighbour in $C^* \setminus C_6$, then that neighbour is a . We recall that, since x is connected to 2, 5, if 3 has a neighbour outside of C^* , this neighbour is common with 6. We will call such a vertex (if it exists) y . By the same reasoning that we applied for x , vertex y cannot have a neighbour outside C^* (therefore x and y are not adjacent), and if it has a neighbour in $C^* \setminus C_6$, this must be c . As in the previous section, we will assume without loss of generality that the degree of x is at least as high as that of y , otherwise we can exchange their roles.

Case 3.1: If x has degree two and y does not exist, then we can apply the H -subgraph reduction (Corollary 1) with $H = G[\{x, 3, 6\}]$, in which case $\alpha(G) = \alpha([H]) = 3$ and hence the removal of $[H]$ decreases $\alpha(G)$ by 3.

Case 3.2: Assume x has degree two and y exists: y is adjacent to 3, 6 and no other vertex, since we assumed that x has degree at least as high as y . We then apply total struction with $H = G[\{x, y, 2, 3, 5, 6\}]$ and $\ell = 3$, which replaces $[H] = G[\{x, y, 1, 2, 3, 4, 5, 6\}]$ with a new vertex z adjacent to a, c and decreases $\alpha(G)$ by 3. This transformation does not introduce a new forbidden induced subgraph, as any path through z in the transformed graph can be mapped to the path $a1234c$ in G .

Case 3.3: Assume x is adjacent to a and y does not exist. In this case we delete x from the graph and claim that the independence number is unchanged. To see this, take a maximum independent set S in G . If $x \notin S$ we are done. Suppose then that $x \in S$, therefore S does not contain any of $a, 2, 5$. As a result, it contains at most two vertices from C_6 . Consider now the set $(S \setminus (C_6 \cup \{x\})) \cup \{1, 3, 5\}$. This is a valid independent set (since S does not contain a) of the same size as S .

Case 3.4: Assume x is adjacent to a , y exists and it has degree 2. In this case we delete from the graph vertex 6. We claim that the independence number stays unchanged. To see this, take a maximum independent set S . If $6 \notin S$ we are done, so suppose that $6 \in S$, therefore $1, 5, y \notin S$. If $3 \notin S$, then $S \setminus \{6\} \cup \{y\}$ is an independent set of the same size in the new graph, and we are done. Suppose then that $3 \in S$, which means that $2, 4 \notin S$. We now observe that the set $S \setminus \{a, x, 3, 6\} \cup \{2, 5, y\}$ is an independent set of size $|S|$ in the new graph. To see that it has the same size, we note that a is adjacent to x . To see that it is independent, we note that $S \setminus \{a, x, 3, 6\}$ does not contain any neighbours of $\{2, 5, y\}$.

Case 3.5: If x is adjacent to a and y exists and is adjacent to c , then we can apply the H -subgraph reduction with $H = G[\{x, y, 1, 4\}]$, in which case $\alpha(H) = \alpha([H]) = 4$ and hence the removal of $[H]$ decreases $\alpha(G)$ by 4.

5. Conclusion

Summarizing the discussion in the previous sections, we make the following conclusion, which extends several previously known results.

Theorem 3. *The maximum independent set problem can be solved in polynomial time in the class of $(A_k^*, A_{k+1}^*, \dots)$ -free subcubic graphs for any fixed value of k .*

Since A_t^* contains $S_{2,k,k}$ for any $t > k$, we derive the following corollary from this theorem.

Corollary 2. *The maximum independent set problem can be solved in polynomial time in the class of $S_{2,k,k}$ -free subcubic graphs for any fixed value of k .*

This result brings us much closer to the dichotomy described in Conjecture 1. However, proving this conjecture in its whole generality remains a challenging open problem.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] G. Alexe, P.L. Hammer, V.V. Lozin, D. de Werra, Struction revisited, *Discrete Appl. Math.* 132 (2003) 27–46.
- [2] H.L. Bodlaender, D.M. Thilikos, Treewidth for graphs with small chordality, *Discrete Appl. Math.* 79 (1–3) (1997) 45–61.
- [3] A. Brandstädt, V.V. Lozin, R. Mosca, Independent sets of maximum weight in apple-free graphs, *SIAM J. Discrete Math.* 24 (1) (2010) 239–254.
- [4] F. Gavril, Algorithms for minimum coloring, maximum clique, minimum covering by cliques, and maximum independent set of a chordal graph, *SIAM J. Comput.* 1 (1972) 180–187.
- [5] A. Harutyunyan, M. Lampis, V. Lozin, J. Monnot, Maximum independent sets in subcubic graphs: new results, in: *Lecture Notes in Computer Science*, vol. 11789, 2019, pp. 40–52.
- [6] V. Lozin, From matchings to independent sets, *Discrete Appl. Math.* 231 (2017) 4–14.
- [7] V. Lozin, M. Milanič, A polynomial algorithm to find an independent set of maximum weight in a fork-free graph, *J. Discret. Algorithms* 6 (4) (2008) 595–604.
- [8] V.V. Lozin, M. Milanič, C. Purcell, Graphs without large apples and the maximum weight independent set problem, *Graphs Comb.* 30 (2) (2014) 395–410.
- [9] V. Lozin, J. Monnot, B. Ries, On the maximum independent set problem in subclasses of subcubic graphs, *J. Discret. Algorithms* 31 (2015) 104–112.
- [10] G.J. Minty, On maximal independent sets of vertices in claw-free graphs, *J. Comb. Theory, Ser. B* 28 (1980) 284–304.
- [11] O.J. Murphy, Computing independent sets in graphs with large girth, *Discrete Appl. Math.* 35 (1992) 167–170.
- [12] R.E. Tarjan, Decomposition by clique separators, *Discrete Math.* 55 (1985) 221–232.
- [13] S.H. Whitesides, An algorithm for finding clique cut-sets, *Inf. Process. Lett.* 12 (1981) 31–32.