Journal of Discrete Algorithms 31 (2015) 104–112

Contents lists available at ScienceDirect

Journal of Discrete Algorithms

www.elsevier.com/locate/jda

On the maximum independent set problem in subclasses of subcubic graphs ☆

Vadim Lozin [a,∗,1], Jérôme Monnot [b,c], Bernard Ries [c,b]

[a] DIMAP and Mathematics Institute, University of Warwick, Coventry, CV4 7AL, UK
[b] CNRS, LAMSADE UMR 7243, France
[c] PSL, Université Paris-Dauphine, 75775 Paris Cedex 16, France

ARTICLE INFO

Article history:
Available online 16 September 2014

Keywords:
Independent set
Polynomial-time algorithm
Subcubic graph
APX-completeness

ABSTRACT

It is known that the maximum independent set problem is NP-complete for subcubic graphs, i.e. graphs of vertex degree at most 3. Moreover, the problem is NP-complete for 3-regular Hamiltonian graphs and for H-free subcubic graphs whenever H contains a connected component which is not a tree with at most 3 leaves. We show that if every connected component of H is a tree with at most 3 leaves and at most 7 vertices, then the problem can be solved for H-free subcubic graphs in polynomial time. We also strengthen the NP-completeness of the problem on 3-regular Hamiltonian graphs by showing that the problem is APX-complete in this class.

1. Introduction

In a graph, an *independent set* is a subset of vertices no two of which are adjacent. The maximum independent set problem (MaxIS for short) consists in finding in a graph an independent set of maximum cardinality. This problem is generally NP-complete [6]. Moreover, it remains NP-complete even under substantial restrictions, for instance, for planar graphs or graphs of large girth [12]. On the other hand, for graphs in some particular classes, such as perfect graphs or claw-free graphs [11], the problem can be solved in polynomial time. In order to better understand the boundary between the NP-complete and polynomially-solvable cases of the problem, in the present paper we study MaxIS restricted to graphs of vertex degree at most 3 (also known as subcubic graphs), which is the best possible restriction expressed in terms of vertex degree under which the problem remains NP-complete. This restriction can also be expressed in terms of forbidden induced subgraphs, in which case the set of excluded graphs consists of 11 minimal graphs containing a vertex of degree 4. However, in terms of forbidden induced subgraphs the restriction to subcubic graphs is not best possible, because the problem is NP-complete in the class of $(K_{1,4}, K_3)$-free graphs, which is a proper subclass of subcubic graphs. This follows, in particular, from the result in [1] that can be stated as follows: if Z is a *finite* set containing no graph every connected component of which is a tree with at most three leaves, then MaxIS is NP-complete in the class of Z-free graphs. In other words, for polynomial-time solvability of the problem in a class defined by *finitely many* forbidden induced subgraphs, we must exclude a graph every connected component of which has the form $S_{i,j,k}$ represented in Fig. 1. Whether this condition is sufficient for polynomial-time solvability of the problem is a big open question.

☆ A preliminary version of this paper has appeared in the proceedings of IWOCA 2013.

∗ Corresponding author.
E-mail addresses: v.lozin@warwick.ac.uk (V. Lozin), monnot@lamsade.dauphine.fr (J. Monnot), ries@lamsade.dauphine.fr (B. Ries).

1 The first author gratefully acknowledges support from DIMAP – the Center for Discrete Mathematics and Its Applications at the University of Warwick, grant EP/D063191/1 (from EPSRC), and from EPSRC, grant EP/L020408/1.

http://dx.doi.org/10.1016/j.jda.2014.08.005
1570-8667/

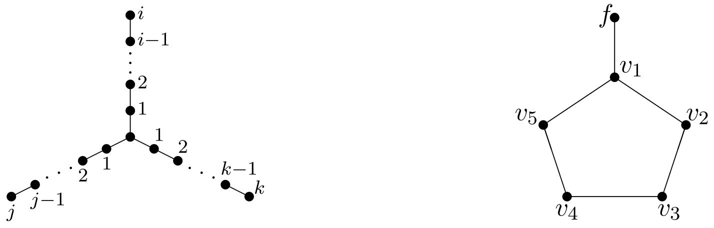

Fig. 1. Graphs $S_{i,j,k}$ (left) and A_5 (right).

Without the restriction on vertex degree, polynomial-time solvability of the problem in classes of $S_{i,j,k}$-free graphs was shown only for very small values of i, j, k. In particular, the problem can be solved for $S_{1,1,1}$-free (claw-free) graphs [11], $S_{1,1,2}$-free (fork-free) graphs [8], and $S_{0,1,1} + S_{0,1,1}$-free ($2P_3$-free) graphs [10]. Recently, Lokshtanov, Vatshelle, and Villanger [7] proved that the independence number of an $S_{0,2,2}$-free (P_5-free) graph can be computed in polynomial time (thereby solving a long-standing open problem).

With the restriction on vertex degree, we can do much better. In particular, we can solve the problem for P_k-free graphs of degree at most d for any k and d, because under this restriction the number of vertices in *connected* graphs is bounded by a function of k and d. More generally, we can solve the problem for $S_{1,j,k}$-free graphs of bounded degree for any j and k, because by excluding $S_{1,j,k}$ we exclude large apples (see definition in the end of the introduction), and for graphs of bounded degree containing no large apples the problem can be solved in polynomial time, which was recently shown in [9]. However, nothing is known about the complexity of the problem in classes of $S_{i,j,k}$-free graphs of bounded degree where all three indices i, j, k are at least 2. To make a progress in this direction, we consider best possible restrictions of this type, i.e. we study $S_{2,2,2}$-free graphs of vertex degree at most 3, and show that the problem is solvable in polynomial time in this class. More generally, we show that the problem is polynomial-time solvable in the class of H-free subcubic graphs whenever H is a graph every connected component of which is isomorphic to $S_{2,2,2}$ or to $S_{1,j,k}$. This result is presented in Section 2.

In Section 3, we switch to the classes where the problem is difficult and prove a new result in this area. In particular, we show that MaxIS is APX-complete in the class of 3-regular Hamiltonian graphs, which strengthens the NP-completeness of the problem in this class.

Section 4 concludes the paper with a number of open problems.

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 are denoted by $V(G)$ and $E(G)$, respectively. For a vertex $v \in V(G)$, we denote by $N(v)$ the neighborhood of v, i.e., the set of vertices adjacent to v, and by $N[v]$ the closed neighborhood of v, i.e. $N[v] = N(v) \cup \{v\}$. For $v, w \in V(G)$, we set $N[v, w] = N[v] \cup N[w]$. The *degree* of v is the number of its neighbors, i.e., $d(v) = |N(v)|$. The subgraph of G induced by a set $U \subseteq V(G)$ is obtained from G by deleting the vertices outside of U and is denoted $G[U]$. If no induced subgraph of G is isomorphic to a graph H, then we say that G is H-free. Otherwise we say that G contains H. If G contains H, we denote by $[H]$ the subgraph of G induced by the vertices of H and all their neighbors. As usual, by C_p we denote a chordless cycle of length p. Also, an *apple* A_p, $p \geq 4$, is a graph consisting of a cycle C_p and a vertex f which has exactly one neighbor on the cycle. We call vertex f the *stem* of the apple. See Fig. 1 for the apple A_5. The size of a maximum independent set in G is called the *independence number* of G and is denoted $\alpha(G)$.

2. Polynomial-time results

In this section, we show that the problem is polynomial-time solvable in the class of H-free subcubic graphs whenever H is a graph every connected component of which is isomorphic to $S_{2,2,2}$ or to $S_{1,j,k}$. We start by solving the problem for $S_{2,2,2}$-free subcubic graphs. To this end, we quote the following result from [9].

Theorem 2.1. *For any positive integers d and p,* MaxIS *is polynomial-time solvable in the class of $(A_p, A_{p+1}, \ldots)$-free graphs with maximum vertex degree at most d.*

We solve MaxIS for $S_{2,2,2}$-free subcubic graphs by reducing it to subcubic graphs without large apples.

Throughout the section we let G be an $S_{2,2,2}$-free subcubic graph and $K \geq 1$ a large fixed integer. If G contains no apple A_p with $p \geq K$, then the problem can be solved for G by Theorem 2.1. Therefore, from now on we assume that G contains an induced apple A_p with $p \geq K$ formed by a chordless cycle $C = C_p$ of length p and a stem f. We denote the vertices of C by $v_1, \ldots, v_p$ (listed along the cycle) and assume without loss of generality that the only neighbor of f on C is v_1 (see Fig. 1 for an illustration).

If v_1 is the only neighbor of f in G, then the deletion of v_1 together with f reduces the independence number of G by exactly 1. This can be easily seen and also is a special case of a more general reduction described in Section 2.1. The deletion of f and v_1 destroys the apple A_p. The idea of our algorithm is to destroy all large apples by means of other

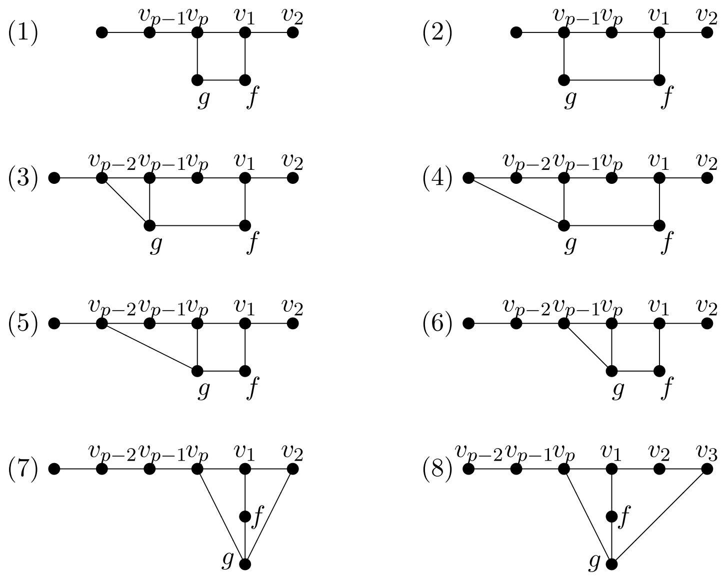

Fig. 2. $A_p + g$.

simple reductions that change the independence number by a constant. Before we describe the reductions in Section 2.1, let us first characterize the local structure of G in the case when the stem f has a neighbor different from v_1.

Lemma 2.2. *If f has a neighbor g different from v_1, then g has at least one neighbor on C and the neighborhood of g on C is of one of the 8 types represented in Fig. 2.*

Proof. First observe that g must have a neighbor among $\{v_{p-1}, v_p, v_2, v_3\}$, since otherwise we obtain an induced $S_{2,2,2}$. If g has only 1 neighbor on C, then clearly we obtain configuration (1) or (2).

Now assume that g has two neighbors on C. Suppose first that g is adjacent neither to v_2 nor to v_p. Then g must be adjacent to at least one of v_{p-1}, v_3. Without loss of generality, we may assume that g is adjacent to v_{p-1} and denote the third neighbor of g by v_j. If $2 < j < p-3$, then we clearly obtain an induced $S_{2,2,2}$ centered at g. Otherwise, we obtain configuration (3) or (4).

Now assume g is adjacent to one of v_2, v_p, say to v_p, and again denote the third neighbor of g by v_j. If $j \in \{p-2, p-1\}$, then we obtain configuration (5) or (6). If $j \in \{2, 3\}$, then we obtain configuration (7) or (8). If $3 < j < p-2$, then G contains an $S_{2,2,2}$ induced by $\{v_{j-2}, v_{j-1}, v_j, v_{j+1}, v_{j+2}, g, f\}$. □

2.1. Graph reductions

As we mentioned earlier, the idea of our algorithm is to destroy all large apples by means of reductions that change the independence number by a constant. In the present section we describe the main reductions used in our solution.

2.1.1. H-subgraph reduction

Let H be an induced subgraph of G.

Lemma 2.3. *If $\alpha(H) = \alpha([H])$, then $\alpha(G - [H]) = \alpha(G) - \alpha(H)$.*

Proof. Since any independent set of G contains at most $\alpha([H])$ vertices in $[H]$, we know that $\alpha(G - [H]) \geq \alpha(G) - \alpha([H])$. Now let S be an independent set in $G - [H]$ and A an independent set of size $\alpha(H)$ in H. Then $S \cup A$ is an independent set in G and hence $\alpha(G) \geq \alpha(G - [H]) + \alpha(H)$. Combining the two inequalities together with $\alpha(H) = \alpha([H])$, we conclude that $\alpha(G - [H]) = \alpha(G) - \alpha(H)$. □

The deletion of $[H]$ in the case when $\alpha(H) = \alpha([H])$ will be called the *H-subgraph reduction*. For instance, if a vertex v has degree 1, then the deletion of v together with its only neighbor is the H-subgraph reduction with $H = \{v\}$.

2.1.2. Φ-reduction

Let us denote by Φ the graph represented on the left of Fig. 3. The transformation replacing Φ by Φ' as shown in Fig. 3 will be called *Φ-reduction*.

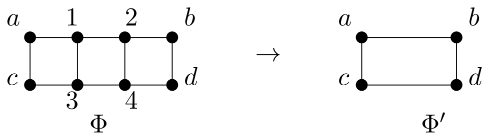

Fig. 3. Φ-reduction.

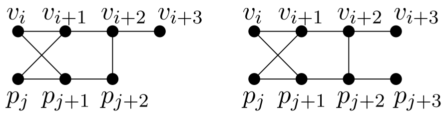

Fig. 4. Induced subgraphs A (left) and B (right).

Lemma 2.4. *By applying the Φ-reduction to an $S_{2,2,2}$-free subcubic graph G, we obtain an $S_{2,2,2}$-free subcubic graph G' such that $\alpha(G') = \alpha(G) - 2$.*

Proof. Let S be an independent set in G. Clearly it contains at most two vertices in $\{a, b, c, d\}$ and at most two vertices in $\{1, 2, 3, 4\}$. Denote $X = S \cap \{1, 2, 3, 4\}$. If the intersection $S \cap \{a, b, c, d\}$ contains at most one vertex or one of the pairs $\{a, d\}$, $\{b, c\}$, then $S - X$ is an independent set in G' of size at least $\alpha(G) - 2$. If $S \cap \{a, b, c, d\} = \{a, b\}$, then X contains at most one vertex and hence $S - (X \cup \{b\})$ is an independent set in G' of size at least $\alpha(G) - 2$. Therefore, $\alpha(G') \geq \alpha(G) - 2$.

Now let S' be an independent set in G'. Then the intersection $S' \cap \{a, b, c, d\}$ contains at most two vertices. If $S' \cap \{a, b, c, d\} = \{a, d\}$, then $S' \cup \{2, 3\}$ is an independent set of size $\alpha(G') + 2$ in G. Similarly, if $S' \cap \{a, b, c, d\}$ contains at most one vertex, then G contains an independent set of size at least $\alpha(G') + 2$. Therefore, $\alpha(G) \geq \alpha(G') + 2$. Combining the two inequalities, we conclude that $\alpha(G') = \alpha(G) - 2$.

Now let us show that G' is an $S_{2,2,2}$-free subcubic graph. The fact that G' is subcubic is obvious. Assume to the contrary that it contains an induced subgraph H isomorphic to $S_{2,2,2}$. If H contains none of the edges ab and cd, then clearly H is also an induced $S_{2,2,2}$ in G, which is impossible. If S contains both edges ab and cd, then it contains $C_4 = (a, b, c, d)$, which is impossible either. Therefore, H has exactly one of the two edges, say ab. If vertex b has degree 1 in H, then by replacing b by vertex 1 we obtain an induced $S_{2,2,2}$ in G. By symmetry, a also is not a vertex of degree 1 in H. Therefore, we may assume, without loss of generality, that a has degree 3 and b has degree 2 in H. Let us denote by x the only neighbor of b in H. Then $(H - \{b, x\}) \cup \{1, 2\}$ is an induced $S_{2,2,2}$ in G. This contradiction completes the proof. □

2.1.3. AB-reduction

The AB-reduction deals with two graphs A and B represented in Fig. 4. We assume that the vertices v_i belong to the cycle $C = C_p$, and the vertices p_j are outside of C.

Lemma 2.5. *If G contains an induced subgraph isomorphic to A, then*

- *either A can be extended to an induced subgraph of G isomorphic to B in which case p_{j+2} can be deleted without changing $\alpha(G)$*
- *or the deletion of $N[v_i] \cup N[p_j]$ reduces the independence number by 2.*

Proof. Assume first that A can be extended to an induced B (by adding vertex p_{j+3}). Consider an independent set S containing vertex p_{j+2}. Then S contains neither p_{j+1} nor p_{j+3} nor v_{i+2}. If neither p_j nor v_i belongs to S, then p_{j+2} can be replaced by p_{j+1} in S. Now assume, without loss of generality, that v_i belongs to S. Then $v_{i+1} \notin S$ and therefore we may assume that $v_{i+3} \in S$, since otherwise p_{j+2} can be replaced by v_{i+2} in S. If p_{j+3} has one more neighbor x in S (different from p_{j+2}), then vertices v_i, v_{i+2}, v_{i+3}, p_{j+1}, p_{j+2}, p_{j+3} and x induce an $S_{2,2,2}$ in G (because the 3 endpoints are in S and the internal vertices have degree 3 in A). Therefore, we conclude that p_{j+2} is the only neighbor of p_{j+3} in S, in which case p_{j+2} can be replaced by p_{j+3} in S. Thus, for any independent S in G containing vertex p_{j+2}, there is an independent set of size $|S|$ which does not contain p_{j+2}. Therefore, the deletion of p_{j+2} does not change the independence number of G.

Now let us assume that A cannot be extended to B. Clearly, every independent set S in $G - N[v_i, p_j]$ can be extended to an independent set of size $|S| + 2$ in G by adding to S vertices v_i and p_j. Therefore, $\alpha(G) \geq \alpha(G - N[v_i, p_j]) + 2$.

Conversely, consider an independent set S in G. If it contains at most 2 vertices in $N[v_i, p_j]$, then by deleting these vertices from S we obtain an independent set of size at least $|S| - 2$ in $G - N[c_i, p_j]$.

Suppose now that S contains more than 2 vertices in $N[v_i, p_j]$. Let us show that in this case it must contain exactly three vertices in $N[v_i, p_j]$, two of which are v_{i+1} and p_{j+1}. Indeed, $N[v_i, p_j]$ contains at most 6 vertices: v_{i-1}, v_i, v_{i+1}, p_j, p_{j+1} and possibly some vertex x. Moreover, if x exists, then it is adjacent to v_{i-1}, since otherwise an $S_{2,2,2}$ arises induced either by vertices x, p_j, p_{j+1}, p_{j+2}, v_{i+2}, v_{i-1}, v_i (if p_{j+2} is not adjacent to v_{i-1}) or by vertices p_j, v_{i+1}, v_{i+2}, v_{i+3}, v_{i+4}, v_{i-1}, p_{j+2} (if p_{j+2} is adjacent to v_{i-1}). Therefore, S cannot contain more than three vertices in $N[v_i, p_j]$, and if it contains three vertices, then two of them are v_{i+1} and p_{j+1}. As a result, S contains neither v_{i+2} nor p_{j+2}. If each of v_{i+2}

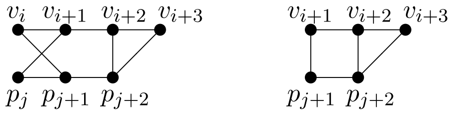

Fig. 5. Induced subgraphs A^* (left) and *House* (right).

and p_{j+2} has one more neighbor in S (different from v_{i+1} and p_{j+1}), then A can be extended to B, which contradicts our assumption. Therefore, we may assume without loss of generality that p_{j+1} is the only neighbor of p_{j+2} in S. In this case, the deletion from $N[v_i, p_j]$ of the three vertices of S and adding to it vertex p_{j+2} results in an independent set of size $|S| - 2$ in $G - N[v_i, p_j]$. Therefore, $\alpha(G - N[v_i, p_j]) \geq \alpha(G) - 2$. Combining with the inverse inequality, we conclude that $\alpha(G - N[v_i, p_j]) = \alpha(G) - 2$. □

2.1.4. Other reductions

Two other reductions that will be helpful in the proof are the following.

- The *A^*-reduction* applies to an induced A^* (Fig. 5) and consists in deleting vertex p_{j+2}.
- The *House-reduction* applies to an induced *House* (Fig. 5) and consists in deleting the vertices of the triangle v_{i+2}, v_{i+3}, p_{j+2}.

Lemma 2.6. *The A^*-reduction does not change the independence number, and the House-reduction reduces the independence number by exactly 1.*

Proof. Assume G contains an induced A^* and let S be an independent set containing p_{j+2}. If S does not contain v_{i+1}, then p_{j+2} can be replaced by v_{i+2}, and if S contains v_{i+1}, then p_{j+2} can be replaced by p_{j+1}. Therefore, G has an independent set of size $|S|$ which does not contain p_{j+2} and hence the deletion of p_{j+2} does not change the independence number.

Assume G contains an induced *House* and let S be a maximum independent set in G. Then obviously at most one vertex of the triangle v_{i+2}, v_{i+3}, p_{j+2} belongs to S. On the other hand, S must contain at least one vertex of this triangle. Indeed, if none of the three vertices belong to S, then each of them must have a neighbor in S (else S is not maximum), but then both v_{i+1} and p_{j+1} belong to S, which is impossible. Therefore, every maximum independent set contains exactly one vertex of the triangle, and hence the deletion of the triangle reduces the independence number by exactly 1. □

2.2. Solving the problem

In the subgraph of G induced by the vertices outside of $C = C_p$ that have at least one neighbor on C, every vertex has degree at most 2 and hence every connected component in this subgraph is either a path or a cycle. Let F be the component of this subgraph containing the stem f. In what follows we analyze all possible cases for F and show that in each case the apple A_p can be destroyed by means of graph reductions described in Section 2.1 or by some other simple reductions.

Lemma 2.7. *If F is a cycle, then A_p can be destroyed by graph reductions that change the independence number by a constant.*

Proof. If F is a triangle, then, according to Lemma 2.2, the neighbors of F in C are three consecutive vertices of C. In this case, F together with two consecutive vertices of C form a House and hence the deletion of F reduces the independence number of G by exactly one.

Assume F is a cycle of length 4 induced by vertices f_1, f_2, f_3, f_4. With the help of Lemma 2.2 it is not difficult to see that the neighbors of F in C must be consecutive vertices, say $v_i, \dots, v_{i+3}$, and the only possible configuration, up to symmetry, is this: v_i is a neighbor of f_1, v_{i+1} is a neighbor of f_2, v_{i+2} is a neighbor of f_4, v_{i+3} is a neighbor of f_3. In this case, the deletion of vertex v_{i+1} does not change the independence number of G. To show this, consider an independent set S containing vertex v_{i+1}. Then S does not contain f_2, v_i, v_{i+2}. If $f_4 \in S$, then $f_1, f_3 \notin S$, in which case v_{i+1} can be replaced by f_2 in S. So, assume $f_4 \notin S$. If $f_3 \notin S$, then we can assume that $v_{i+3} \in S$ (else v_{i+1} can be replaced by f_3 in S), in which case v_{i+1}, v_{i+3} can be replaced by v_{i+2}, f_3. So, assume $f_3 \in S$, and hence $v_{i+3} \notin S$. But now v_{i+1} can be replaced by v_{i+2} in S. This proves that for every independent set S containing v_{i+1}, there is an independent set of the same size that does not contain v_{i+1}. Therefore, the deletion of v_{i+1} does not change the independence number of G.

Assume F is a cycle of length 5 induced by vertices f_1, f_2, f_3, f_4, f_5. With the help of Lemma 2.2 it is not difficult to verify that the neighbors of F in C must be consecutive vertices, say $v_i, \dots, v_{i+4}$, and the only possible configuration, up to symmetry, is this: f_1 is adjacent to v_i, f_2 is adjacent to v_{i+1}, f_3 is adjacent to v_{i+3}, f_4 is adjacent to v_{i+4}, f_5 is adjacent to v_{i+2}. But then the vertices f_2, f_3, f_4, f_5, v_{i+2}, v_{i+4}, v_{i+5} induce an $S_{2,2,2}$.

If F is a cycle of length more than 5, then an induced $S_{2,2,2}$ can be easily found. □

Lemma 2.8. *If F is a path with at least 5 vertices, then A_p can be destroyed by graph reductions that change the independence number by a constant.*

Proof. Assume F has at least 5 vertices $f_1, \ldots, f_5$. Denote the neighbor of f_3 on C by v_i. Assume v_{i-1} has a neighbor in $\{f_1, f_5\}$, say f_1 (up to symmetry). By Lemma 2.2, f_2 is adjacent either to v_{i-2} or v_{i+1}.

Let first f_2 be adjacent to v_{i+1}. Then either f_1 is not adjacent to v_{i-2}, in which case the vertices $v_{i-2}, \ldots, v_{i+1}, f_1, f_2, f_3$ induce an A, or f_1 is adjacent to v_{i-2}, in which case f_4 is adjacent to v_{i+2} (by Lemma 2.2) and hence the vertices $v_i, \ldots, v_{i+3}, f_2, f_3, f_4$ induce an A. In either case, we can apply Lemma 2.5.

Suppose now that f_2 is adjacent to v_{i-2}. Then f_1 is not adjacent to v_{i+1}, since otherwise f_4 is adjacent to v_{i+2} (by Lemma 2.2), in which case the vertices $v_{i+1}, \ldots, v_{i+4}, f_1, f_3, f_4$ induce an $S_{2,2,2}$. As a result, vertices $v_{i-2}, \ldots, v_{i+1}, f_1, f_2, f_3$ induce an A and we can apply Lemma 2.5.

The above discussion shows that v_{i-1} has no neighbor in $\{f_1, f_5\}$. By symmetry, v_{i+1} has no neighbor in $\{f_1, f_5\}$. Then each of v_{i-1} and v_{i+1} has a neighbor in $\{f_2, f_4\}$, since otherwise $f_1, \ldots, f_5, v_i$ together with v_{i-1} or with v_{i+1} induce an $S_{2,2,2}$. Up to symmetry, we may assume that v_{i-1} is adjacent to f_2, while v_{i+1} is adjacent to f_4.

If f_1 is adjacent to v_{i-2} or f_5 is adjacent to v_{i+2}, then an induced Φ arises, in which case we can apply the Φ-reduction. Therefore, we can assume that f_1 is adjacent to v_{i-3}, while f_5 is adjacent to v_{i+3}.

We may assume that vertex v_{i-2} has no neighbor x different from v_{i-3}, v_{i-1}, since otherwise x must be adjacent to f_1 (else vertices $x, v_{i-2}, v_{i-1}, v_i, v_{i+1}, f_1, f_2$ induce an $S_{2,2,2}$), in which case $v_{i-3}, \ldots, v_i, x, f_1, f_2$ induce an A and we can apply the AB-reduction. Similarly, we may assume that vertex f_1 has no neighbor x different from v_{i-3}, f_2. But then $d(f_1) = d(v_{i-2}) = 2$ and we can apply the H-subgraph reduction with $H = \{v_{i-2}, f_1\}$. □

Lemma 2.9. *If F is a path with 4 vertices, then A_p can be destroyed by graph reductions that change the independence number by a constant.*

Proof. Let F be a path (f_1, f_2, f_3, f_4). Without loss of generality we assume that f_2 is adjacent to v_i and f_3 to v_j with $j > i$. By Lemma 2.2, $j = i + 1$ or $j = i + 2$.

Case (a): $j = i + 1$. Assume f_1 is adjacent to v_{i+2}. Then vertices $v_i, v_{i+1}, v_{i+2}, v_{i+3}, f_1, f_2, f_3$ induce either the graph A (if f_1 is not adjacent to v_{i+3}) or the graph A^* (if f_1 is adjacent to v_{i+3}), in which case we can apply either Lemma 2.5 or Lemma 2.6. Therefore, we may assume that f_1 is not adjacent to v_{i+2}, and by symmetry, f_4 is not adjacent to v_{i-1}. Then by Lemma 2.2, f_1 must have a neighbor in $\{v_{i-2}, v_{i-1}\}$ and f_4 must have a neighbor in $\{v_{i+2}, v_{i+3}\}$.

Assume that f_4 is adjacent to v_{i+3}. If v_{i+2} has a neighbor x outside of the cycle C, then x is not adjacent to f_4 (else F has more than 4 vertices) and hence $v_{i-1}, v_i, v_{i+1}, v_{i+2}, x, f_3, f_4$ induce an $S_{2,2,2}$. Therefore, the degree of v_{i+2} in G is 2. Similarly, the degree of f_4 in G is two. But now we can apply the H-subgraph reduction with $H = \{v_{i+2}, f_4\}$. This allows us to assume that f_4 is not adjacent to v_{i+3}, and by symmetry, f_1 is not adjacent to v_{i-2}. But then f_1 is adjacent to v_{i-1} and f_4 is adjacent to v_{i+2}, in which case we can apply the Φ-reduction to the subgraph of G induced by $v_{i-1}, v_i, v_{i+1}, v_{i+2}, f_1, f_2, f_3, f_4$.

Case (b): $j = i + 2$. If f_1 or f_4 is adjacent to v_{i+1}, then an induced graph A arises, in which case we can apply Lemma 2.5. Then f_1 must be adjacent to v_{i-1}, since otherwise it adjacent to v_{i-2} (by Lemma 2.2), in which case vertices $v_{i-2}, f_1, f_2, f_3, f_4, v_i, v_{i+1}$ induce an $S_{2,2,2}$. By symmetry, f_4 is adjacent to v_{i+3}.

If f_1 is adjacent to v_{i-2}, then we can apply the *House*-reduction to the subgraph of G induced by $v_{i-2}, v_{i-1}, v_i, f_1, f_2$, and if f_1 is adjacent to v_{i-3}, then vertices $v_{i-3}, f_1, f_2, f_3, f_4, v_i, v_{i+1}$ induce an $S_{2,2,2}$. Therefore, we may assume by Lemma 2.2 that f_1 has degree 2 in G. By symmetry, f_4 has degree 2. Also, to avoid an induced $S_{2,2,2}$, we conclude that v_{i+1} has degree 2. But now we apply the H-subgraph reduction with $H = \{f_1, v_i, v_{i+2}, f_4\}$, which reduces the independence number of G by 4. □

Lemma 2.10. *If F is a path with 3 vertices, then A_p can be destroyed by graph reductions that change the independence number by a constant.*

Proof. Assume F is a path (f_1, f_2, f_3). Without loss of generality let f_2 be adjacent to v_1. Since G is $S_{2,2,2}$-free, each of f_1 and f_3 must have at least one neighbor in $\{v_{p-1}, v_p, v_2, v_3\}$. Denote $L = \{v_{p-1}, v_p\}$ and $R = \{v_2, v_3\}$.

Case (a): f_1 and f_3 have both a neighbor in R. Due to the symmetry, we assume without loss of generality that f_1 is adjacent to v_2, while f_3 is adjacent to v_3. Then we may further assume that f_1 is adjacent to v_4, since otherwise vertices $v_1, v_2, v_3, v_4, f_1, f_2, f_3$ induced either an A (if f_3 is not adjacent to v_4) or an A^* (if f_3 is adjacent to v_4), in which case we can apply either Lemma 2.5 or Lemma 2.6. But now the deletion of f_3 does not change the independence number of G. Indeed, let S be an independent set containing f_3. If $f_1 \in S$, then f_3 can be replaced by v_3. If $f_1 \notin S$, then we can assume that $v_1 \in S$ (else f_3 can be replaced by f_2), in which case f_3, v_1 can be replaced by f_2, v_2.

The above discussion allows us to assume, without loss of generality, that f_1 has no neighbor in R, while f_3 has no neighbor in L.

Case (b): f_3 is adjacent to v_3. Then we may assume that f_3 is not adjacent to v_2, since otherwise we can apply the *House*-reduction to the subgraph of G induced by v_1, v_2, v_3, f_3, f_2. Let us show that in this case

- *the degree of v_2 is 2.* Assume to the contrary v_2 has a third neighbor x. Then x is not adjacent to v_{p-1}, since otherwise G contains an $S_{2,2,2}$ induced either by $v_{p-1}, x, v_2, v_1, f_2, v_3, v_4$ (if x is not adjacent to v_4) or by $v_{p-2}, v_{p-1}, x, v_2, v_1, v_4, v_5$ (if x is adjacent to v_4). This implies that x is adjacent to v_p, since otherwise $x, v_2, v_1, f_2, f_3, v_p, v_{p-1}$ induce an $S_{2,2,2}$. As a result, f_1 is adjacent to v_{p-1}. Due to the degree restriction, x may have at most one neighbor in $\{v_{p-3}, v_{p-2}, v_4, v_5\}$. By symmetry, we may assume without loss of generality that x has no neighbor in $\{v_4, v_5\}$. Also, f_3 has no neighbor in $\{v_4, v_5\}$, since otherwise this neighbor together with v_{p-1}, f_1, f_2, f_3, v_1, v_2 would induce an $S_{2,2,2}$. But now x, v_2, v_3, v_4, v_5, f_3, f_2 induce an $S_{2,2,2}$. This contradiction completes the proof of the claim.

If f_3 also has degree two, then we can apply the H-subgraph reduction with $H = \{v_3, f_3\}$. Therefore, may assume that f_3 has one more neighbor, which must be, by Lemma 2.2, either v_4 or v_5. If f_3 is adjacent to f_5, then f_1, f_2, f_3, v_5, v_6, v_3, v_2 induce an $S_{2,2,2}$. Therefore, f_3 is adjacent to v_4. But now v_3 can be deleted without changing the independence number. Indeed, let S be an independent set containing v_3. If S does not contain v_1, then v_3 can be replaced by v_2, and if S contains v_1, then v_1, v_3 can be replaced by v_2, f_3.

Cases (a) and (b) reduce the analysis to the situation when f_1 is adjacent to v_p and non-adjacent to v_{p-1}, while f_3 is adjacent to v_2 and non-adjacent to v_3. If f_3 is adjacent to v_4, then vertices v_p, v_1, v_2, v_3, v_4, f_1, f_2, f_3 induce the graph Φ, in which case we can apply Lemma 2.4. Therefore, we can assume by Lemma 2.2 that the degree of f_3 is 2, and similarly the degree of f_1 is 2. But now we can apply the H-subgraph reduction with $H = \{f_1, v_1, f_3\}$, which reduces the independence number of G by 3. □

Lemma 2.11. *If F is a path with 2 vertices, then A_p can be destroyed by graph reductions that change the independence number by a constant.*

Proof. If F is a path with 2 vertices, we deal with the eight cases represented in Fig. 2. It is easy to see that in cases (1) and (7), every maximum independent set must contain exactly one of f, g and thus by deleting f, g we reduce the independence number by exactly 1.

In case (5), the deletion of f, g also reduces the independence number by exactly 1. Indeed, let S be a maximum independent set containing neither f nor g. Since S is maximum it must contain v_1, v_{p-2} and hence it does not contain v_p, v_{p-1}. But then $(S \setminus \{v_1\}) \cup \{v_p, f\}$ is an independent set larger than S, contradicting the choice of S. Therefore, every maximum independent set contains exactly one of f and g and hence $\alpha(G - \{f, g\}) = \alpha(G) - 1$.

In case (2), the deletion of the set $X = \{v_{p-1}, v_p, v_1, f, g\}$ reduces the independence number of the graph by exactly 2. Indeed, any independent set of G contains at most two vertices in X, and hence $\alpha(G - X) \geq \alpha(G) - 2$. Assume now that S is a maximum independent set in $G - X$. If $v_2 \notin S$, then $S \cup \{v_1, g\}$ is an independent set in G of size $\alpha(G - X) + 2$. Now assume $v_2 \in S$. By symmetry, $v_{p-2} \in S$. Assume v_p has a neighbor x in S. Then x is adjacent neither to v_{p-2} nor to v_2, as all three vertices belong to S. Also, x cannot be adjacent to both v_{p-3} and v_3, since otherwise an induced $S_{2,2,2}$ can be easily found. But if x is not adjacent, say, to v_3, then $x, v_p, v_1, v_2, v_3, f, g$ induce an $S_{2,2,2}$. This contradiction shows that v_p has no neighbors in S. Therefore, $S \cup \{v_p, f\}$ is an independent set in G of size $\alpha(G - X) + 2$, and hence $\alpha(G) \geq \alpha(G - X) + 2$. Combining the two inequalities, we conclude that $\alpha(G - X) = \alpha(G) - 2$.

In case (3), we may delete g without changing the independence number, because in any independent set S containing g, vertex g can be replaced either by v_{p-1} (if S does not contain v_p) or by f (if S contains v_p). In case (6), we apply the *House*-reduction.

In cases (4) and (8), we find another large apple A' whose stem f' belongs to a path F' with at least 3 vertices. In case (4), A' is induced by the cycle $v_1, \dots, v_{p-3}, g, f$ with stem $f' = v_{p-1}$, and in case (8) the apple is induced by the cycle $v_3, \dots, v_p, g$ with stem $f' = v_1$. In both cases, the situation can be handled by one of the previous lemmas. □

Theorem 2.12. *Let H be a graph every connected component of which is isomorphic either to $S_{2,2,2}$ or to $S_{1,j,k}$.* MaxIS *can be solved for H-free graphs of maximum vertex degree at most* 3 *in polynomial time.*

Proof. First, we show how to solve the problem in the case when $H = S_{2,2,2}$. Let $G = (V, E)$ be an $S_{2,2,2}$-free subcubic graph and let K be a large fixed constant. We start by checking if G contains an apple A_p with $p \geq K$. To this end, we detect every induced $S_{1,k,k}$ with $k = K/2$, which can be done in time n^K. If G is $S_{1,k,k}$-free, then it is obviously A_p-free for each $p \geq K$. Assume a copy of $S_{1,k,k}$ has been detected and let x, y be the two vertices of this copy at distance k from the center of $S_{1,k,k}$. We delete from G all vertices of $V(S_{1,k,k}) - \{x, y\}$ and all their neighbors, except x and y, and determine if in the resulting graph there is a path connecting x to y. It is not difficult to see that this procedure can be implemented in polynomial time.

Assume G contains an induced apple A_p with $p \geq K$. If the stem of the apple has degree 1 in G, we delete it together with its only neighbor, which destroys the apple and reduces the independence number of G by exactly one. If the stem has degree more than 1, we apply one of the lemmas of Section 2.2 to destroy A_p and reduce the independence number of G. It is not difficult to see that all the reductions used in the lemmas can be implemented in polynomial time.

Thus in polynomial time we reduce the problem to a graph G' which does not contain any apple A_p with $p \geq K$, and then we find a maximum independent set in G' with the help of Theorem 2.1. This also shows that in polynomial time we can compute $\alpha(G)$, since we know the difference between $\alpha(G)$ and $\alpha(G')$. To find a maximum independent set in G, we take an arbitrary vertex $v \in V(G)$. If $\alpha(G-v)=\alpha(G)$, then there is a maximum independent set in G that does not contain v and hence v ignored (deleted). Otherwise, v belongs to every maximum independent set in G and hence it must be included in the solution. Therefore, in polynomial time we can find a maximum independent set in G. This completes the proof of the theorem in the case when $H=S_{2,2,2}$.

By Theorem 2.1 we also know how to solve the problem in the case when $H=S_{1,j,k}$. Now we assume that H contains $s>1$ connected components. Denote by S any of the components of H and let H' be the graph obtained from H by deleting S. Consider an H-free graph G. If G does not contain a copy of S, the problem can be solved for G by the first part of the proof. So, assume G contains a copy of S. By deleting from G the vertices of $[S]$ we obtain a graph G' which is H'-free and hence the problem can be solved for G' by induction on s. The number of vertices in $[S]$ is bounded by a constant independent of $|V(G)|$ (since $|V(S)|<|V(H)|$ and every vertex of S has at most three neighbors in G), and hence the problem can be solved for G in polynomial time as well, which can be easily seen by induction on the number of vertices in $[S]$. □

3. APX-completeness of MaxIS in 3-regular Hamiltonian graphs

In [5], it was shown that MaxIS is NP-complete in 3-regular Hamiltonian graphs, even if the graph is planar. On the other hand, the problem admits a polynomial-time approximation scheme (PTAS) in this class, because it admits a PTAS for general planar graphs (see [2]). In the present section, we prove that without the planarity condition, MaxIS does not admit a PTAS for 3-regular Hamiltonian graphs, i.e. the problem is APX-complete in this class.

Theorem 3.1. MaxIS *is APX-complete in 3-regular Hamiltonian graphs.*

Proof. Our proof will be done using an approximation preserving reduction from the maximum 2-satisfiability problem with variables appearing each exactly 3 times (Max2Sat-3 for short). An instance $I=(\mathcal{C},X)$ of Max2Sat-3 consists of a collection $\mathcal{C}=(C_1,\dots,C_m)$ of clauses over the set $X=\{x_1,\dots,x_n\}$ of Boolean variables, such that each clause C_j contains exactly 2 literals and each variable appears exactly 3 times (without loss of generality, we may assume that each variable appears one time positively and twice negatively). The goal is to find a truth assignment f satisfying a maximum number of clauses. Max2Sat-3 has been shown to be APX-complete in [3,4].

From an instance $I=(\mathcal{C},X)$ of Max2Sat-3, we build a 3-regular Hamiltonian graph G, instance of MaxIS, as follows:

- For every variable x_i, $i=1,\dots,n$, appearing positively in clause C_{i_2} and negatively in clauses C_{i_1}, C_{i_3}, we construct a path $T_i=(V_i,E_i)$ with $V_i=\{v_i(i_1), v_i(i_2), v_i(i_3)\}$ and $E_i=\{v_i(i_1)v_i(i_2), v_i(i_2)v_i(i_3)\}$.
- For $i=1,\dots,n$, T_i is connected to T_{i+1} (modulo n) using the graph F_i (see Fig. 6) by adding the edges $v_i(i_3)a(i)$ and $b(i)v_{i+1}((i+1)_1)$ (see Fig. 7).
- Finally, if a clause C_k contains variables x_i and x_j, then we add an edge $v_i(c)v_j(d)$ where $c \in \{i_1,i_2,i_3\}$ and $d \in \{j_1,j_2,j_3\}$ according to their appearance in I.

The resulting graph G is clearly 3-regular and can be constructed in polynomial time. Furthermore, G contains a Hamiltonian cycle C. Indeed, C is obtained by starting at any vertex $v_i(i_1)$, visiting the vertices of T_i, then traversing F_i visiting in order $a(i), a_1(i), a_2(i), b_2(i), b_1(i), b(i)$ and finally going to $v_{i+1}((i+1)_1)$ and repeating the same procedure until we end up at vertex $v_i(i_1)$ again.

Let f^* be a truth assignment of $I=(\mathcal{C},X)$ satisfying $p=opt_{Max2Sat\text{-}3}(I)$ clauses. Let $C_{k_1},\dots,C_{k_p}$ be the clauses that are satisfied by f^*. For each satisfied clause C_{k_ℓ}, let $v_{i_\ell}(g(k_\ell))$ be a vertex corresponding to the variable x_{i_ℓ} satisfying C_{k_ℓ}.

We set $S=\{a_1(i), b_2(i): i=1,\dots,n\} \cup \{v_{i_\ell}(g(k_\ell)): \ell=1,\dots,p\}$. Clearly, S is an independent set of G because f^* is a truth assignment and we select exactly one vertex per satisfied clause. Furthermore, $|S|=p+2n$. Thus,

$$\alpha(G) \geq opt_{Max2Sat\text{-}3}(I)+2n. \tag{1}$$

Conversely, let S be an independent set in G. Since any independent set contains at most 2 vertices in each F_i, we conclude that the set $S'=S \setminus (\bigcup_{i=1}^n F_i)$ has size at least $|S|-2n$. By construction of G, for every vertex $v_i(c) \in S'$, there exists an edge $[v_i(c), v_j(d)]$ in G corresponding to the unique clause C_k. Hence, the mapping f defined by $f(x_i)=\texttt{true}$ if $v_i(i_2) \in S'$ and $f(x_i)=\texttt{false}$ otherwise, is a truth assignment of $I=(\mathcal{C},X)$ satisfying at least $|S'|$ clauses (because S' is an independent set). Thus,

$$val(f) \geq |S|-2n. \tag{2}$$

Using (1) and (2), we deduce that $\alpha(G)=opt_{Max2Sat\text{-}3}(I)+2n$. It is well-known that any optimal assignment satisfies at least half of the clauses (since in each pair containing an assignment and its complement the better assignment must satisfy at least half of the clauses), therefore $opt_{Max2Sat\text{-}3}(I) \geq \frac{m}{2}$. Also, since we deal with Max2Sat-3, $2m=3n$. As a result, we conclude that

$$\alpha(G) \leq \frac{11}{3} opt_{Max2Sat\text{-}3}(I).$$

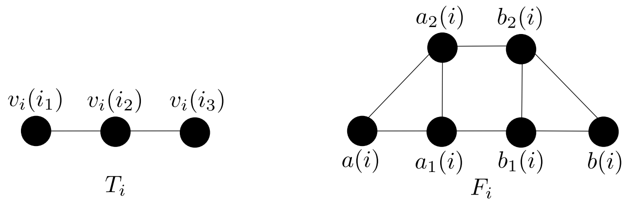

Fig. 6. The gadget variable T_i and the dummy gadget F_i.

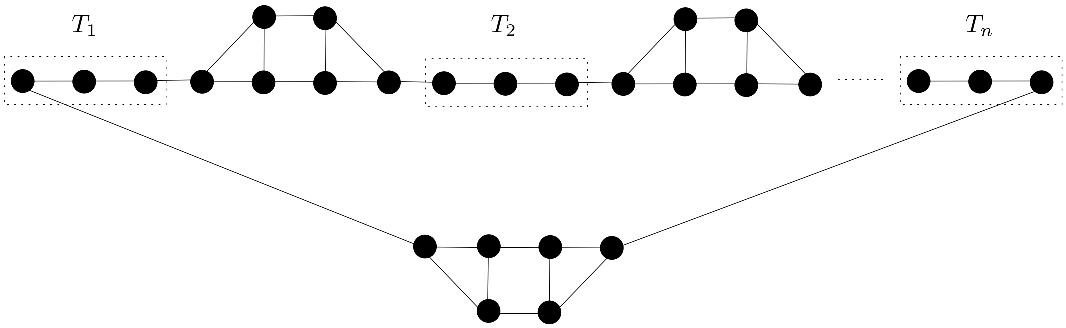

Fig. 7. How the paths T_i are connected using the graphs F_i.

Finally,

$$opt_{Max2Sat\text{-}3}(I) - val(f) \le \alpha(G) - |S|.$$

The last two inequalities show that our reduction is an L-reduction and hence it preserves approximation schemes [13]. □

4. Conclusion

Unless $P = NP$, MaxIS can be solved in polynomial time for H-free subcubic graphs *only if* every connected component of H has the form $S_{i,j,k}$ represented in Fig. 1. Whether this condition is sufficient for polynomial-time solvability of the problem is a challenging open question. In this paper, we contributed to this topic by solving the problem in the case when every connected component of H is isomorphic either to $S_{2,2,2}$ or to $S_{1,j,k}$. Our proof also shows that, in order to answer the above question, one can restrict to H-free subcubic graphs where H is connected. In other words, one can consider $S_{i,j,k}$-free, or more generally, $S_{k,k,k}$-free subcubic graphs. We believe that the answer is positive for all values of k and hope that our solution for $k = 2$ can base a foundation for algorithms for larger values of k.

References

[1] V.E. Alekseev, On the local restrictions effect on the complexity of finding the graph independence number, in: Combinatorial–Algebraic Methods in Applied Mathematics, Gorky University Press, 1983, pp. 3–13 (in Russian).
[2] B.S. Baker, Approximation algorithms for NP-complete problems on planar graphs, J. ACM 41 (1) (1994) 153–180.
[3] P. Berman, M. Karpinski, Efficient amplifiers and bounded degree optimization, Electron. Colloq. Comput. Complex. 8 (53) (2001).
[4] P. Berman, M. Karpinski, Improved approximation lower bounds on small occurrence optimization, Electron. Colloq. Comput. Complex. 10 (8) (2003).
[5] H. Fleischner, G. Sabidussi, V.I. Sarvanov, Maximum independent sets in 3- and 4-regular Hamiltonian graphs, Discrete Math. 310 (2010) 2742–2749.
[6] M.R. Garey, D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W.H. Freeman, ISBN 0-7167-1045-5, 1979.
[7] D. Lokshtanov, M. Vatshelle, Y. Villanger, Independent set in P_5-free graphs in polynomial time, in: Proc. SODA, 2014, pp. 570–581.
[8] V. Lozin, M. Milanič, A polynomial algorithm to find an independent set of maximum weight in a fork-free graph, J. Discrete Algorithms 6 (2008) 595–604.
[9] V.V. Lozin, M. Milanic, C. Purcell, Graphs without large apples and the maximum weight independent set problem, Graphs Comb. 30 (2014) 395–410.
[10] V.V. Lozin, R. Mosca, Maximum regular induced subgraphs in $2P_3$-free graphs, Theor. Comput. Sci. 460 (2012) 26–33.
[11] G.J. Minty, On maximal independent sets of vertices in claw-free graphs, J. Comb. Theory, Ser. B 28 (1980) 284–304.
[12] O.J. Murphy, Computing independent sets in graphs with large girth, Discrete Appl. Math. 35 (1992) 167–170.
[13] C.H. Papadimitriou, M. Yannakakis, Optimization, approximation, and complexity classes, J. Comput. Syst. Sci. 43 (3) (1991) 425–440.