NP-hardness Results for Partitioning Graphs into Disjoint Cliques and a Triangle-free Subgraph

Carl Feghali*[1], Faisal N. Abu-Khzam†[2], and Haiko Müller‡[3]

[1]School of Engineering and Computing Sciences, Durham University, Durham, U.K.
[2]Department of Computer Science and Mathematics, Lebanese American University, Beirut, Lebanon.
[3]School of Computing, University of Leeds, Leeds, U.K.

March 25, 2014

Abstract

This paper investigates the computational complexity of deciding whether the vertices of a graph can be partitioned into a disjoint union of cliques and a triangle-free subgraph. This problem is known to be NP-complete on arbitrary graphs. We show that this problem remains NP-complete even when restricted to planar graphs and perfect graphs.

1 Introduction

According to [3], given graph properties $\mathcal{P}$ and $\mathcal{Q}$, a $(\mathcal{P}, \mathcal{Q})$-colouring of a graph G is a partition of its vertex set into two sets A and B such that A

*carl.feghali@durham.ac.uk
†faisal.abukhzam@lau.edu.lb
‡h.muller@leeds.ac.uk

induces a subgraph that belongs to $\mathcal{P}$ and B induces a subgraph that belongs to $\mathcal{Q}$. A graph G is $(\mathcal{P}, \mathcal{Q})$-colourable if G admits a $(\mathcal{P}, \mathcal{Q})$-colouring.

In this paper, we investigate the computational complexity of deciding whether a graph G is (P_3-free, K_3-free)-colourable, that is, whether G admits a partition of its vertex into two sets A and B such that A induces a P_3-free graph (i.e., a disjoint union of cliques) and B induces a K_3-free graph (i.e., a graph with no triangle). This problem is known to be NP-complete on general graphs [3]. We thus restrict our attention to special classes of graphs. Our hardness results are stated in the following two theorems.

Theorem 1.1. *Deciding whether a planar graph G is (P_3-free, K_3-free)-colourable is* NP*-complete.*

Theorem 1.2. *Deciding whether a short-chorded graph G is (P_3-free, K_3-free)-colourable is* NP*-complete.*

Theorem 1.2 implies the same for perfect graphs (see [1, 6]). Section 2 introduces the terminology that will be used in the rest of this paper. Sections 3 and 4 contain the hardness proofs on planar graphs and perfect graphs, respectively. Finally, we discuss open problems in Section 5.

2 Background

All graphs considered here are finite and have no multiple edges and no loops. For undefined graph terminology we refer the reader to Diestel [4]. Let $G = (V, E)$ be a graph and $V' \subseteq V$. The graph G' induced by deleting the vertices $V \setminus V'$ from G is denoted by $G' = G[V']$. K_n, C_n, P_n denote a complete graph, a cycle, and a path on n vertices respectively. We say that G is H-free if it contains no subgraph isomorphic to some graph H. The graph $G \setminus v$ is obtained from G by deleting the vertex v. We do not distinguish between isomorphic graphs. The union $Q = G \cup H$ of graphs G and H is such that for any $v \in V(G)$ and $u \in V(H)$, $uv \notin E(Q)$.

A graph is said to be *embeddable in the plane*, or *planar*, if it can be drawn in the plane so that its edges intersect only at their ends.

An *odd hole* is an induced cycle of odd length at least 5. A graph G is *short-chorded* (also known as *Raspail*) if every odd cycle C of length at least 5 in G has a short chord, i.e., a chord joining two vertices of distance 2 in C. Short-chorded graphs were introduced in [1]. A graph G is *perfect* if for every induced subgraph H of G, the chromatic number of H equals the size of the largest clique of H. By the strong perfect graph theorem [6], short-chorded graphs are perfect.

Unless otherwise specified, let colouring mean (P_3-free, K_3-free)-colouring, and let colourable mean (P_3-free, K_3-free)-colourable.

3 Proof of Theorem 1.1

In this section, we establish Theorem 1.1. The problem is clearly in NP. To show NP-hardness, we provide a reduction from Planar 3-SAT, which is known to be NP-hard [5], and defined as follows: Given a boolean formula ψ, its *associated graph* $G(\psi)$ has one vertex v_x for each variable x in ψ and one vertex v_C for each clause C in ψ. There is an edge between v_x and v_C iff x or $\neg x$ appears in C. A instance of Planar 3-SAT is a set of variables $X = \{x_1, x_2, \ldots, x_n\}$ and a set of clauses $C = \{C_i \mid i = 1, 2, \ldots, m\}$, such that each $C_i = (l_{i,1} \vee l_{i,2} \vee l_{i,3})$ consists of 3 literals and each literal $l_{i,k}$ is x_p or $\overline{x_p}$ for some $x_p \in X$. Given a boolean formula $\theta = C_1 \wedge C_2 \wedge \ldots \wedge C_m$, the problem is to determine whether there exists a truth assignment to the variables in X such that θ is satisfiable, where $G(\theta)$ is known to be planar. We can safely assume that a literal and its negation do not occur in the same clause.

Gadgets. The *weak negator gadget* with endpoints x, y is presented in Figure 1. The following observations are left as a simple exercise to the reader. Clearly, the gadget admits a colouring. Moreover, x and y cannot both be in the P_3-free part, and if x (or y) is in the P_3-free part, then there exists a colouring such that x (or y respectively) does not have a neighbour in the P_3-free part.

The *strong negator gadget* with endpoints x, y is presented in Figure 2. The following observations are left as a simple exercise to the reader. Obviously, the gadget admits a colouring. Moreover, x and y have different colours, and if x (or y) is in the P_3-free part, then there exists a colouring such that x (or y respectively) does not have a neighbour in the P_3-free part.

The weak and strong negator gadgets are clearly planar.

Given an instance of Planar 3-SAT, we construct the following reduction graph.

Construction. Let m_x be the number of occurrences of variable x. Each variable x is represented by a variable component X (see Figure 3), which is a cycle of length $2m_x$ whose edges are replaced by a strong negator gadget. We number the vertices from 1 to $2m_x$ in a clockwise traversal. Its odd numbered vertices, denoted by *negative literal vertices*, represent the negative occurrences of x, while its even numbered vertices, denoted by *positive literal vertices*, represent the positive occurrences of x. Each clause $C = (l_{x,i} \vee l_{y,j} \vee l_{z,k})$ is represented by a triangle whose vertices are the vertices of variable components that correspond to the literals $l_{x,i}, l_{y,j}$ and $l_{z,k}$. Denote the graph obtained in this way by F.

Figure 1: The weak negator gadget with endpoints x, y together with a colouring where the white vertices are in the P_3-free part and the black vertices are in the K_3-free part (left), and its symbolic representation (right).

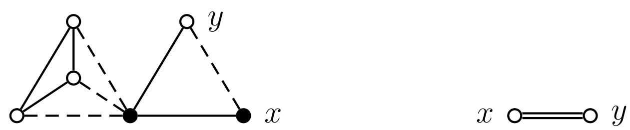

Figure 2: The strong negator gadget with endpoints x, y together with a colouring where the white vertices are in the P_3-free part and the black vertices are in the K_3-free part (left) and its symbolic representation (right)

Lemma 3.1. *F is colourable if and only if θ is satisfiable.*

Proof. By the property of the strong negator gadget, in any colouring of a variable component the positive literal vertices receive one colour and the negative literal vertices receive the other colour.

Suppose θ is satisfiable. If $\theta(x)$ is true, let all the positive literal vertices corresponding to x be in the P_3-free part, and let the negative literal vertices of x are in the K_3-free part. If $\theta(x)$ is false, let all the negative literal vertices corresponding to x be in the P_3-free part, and let the literal vertices of x are in the K_3-free part. Clearly, every variable component is colourable.

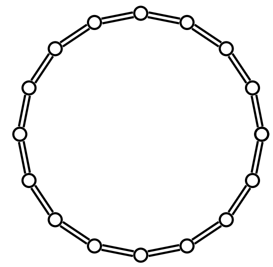

Figure 3: A variable component

Furthermore, every triangle T corresponding to a clause C is colourable, for otherwise all 3 vertices in T belong in the K_3-free part, in which case all three literals in C are false, a contradiction. To ensure that no vertices in the P_3-free part induce a P_3, we can colour each strong negator gadget S occurring in a variable component in such a way that its endpoint in the P_3-free part does not have a neighbour in S in the P_3-free part.

Conversely, suppose F is colourable. If a positive literal vertex corresponding to variable x is in the P_3-free part, set $\theta(x)$ to true. Otherwise, set $\theta(x)$ to false. Observing that every triangle corresponding to a clause must have at least one vertex in the P_3-free part concludes the proof. □

The proof of planarity can be easily derived from [5]. We include it here for completeness:

Lemma 3.2. *F is planar.*

Proof. F can be obtained from the associated graph $G(\theta)$ as follows. For every variable x and vertex v_x occurring in $G(\theta)$, replace v_x by a variable component. For every clause C and vertex v_C occurring in $G(\theta)$, replace v_C by a triangle. There is an edge between a triangle and a variable component whenever the variable represented by the variable component occurs in the clause represented by the triangle. Each node of the triangle is used exactly once. By contracting every edge that goes from a triangle to a variable component we get the graph F as required. □

Conjoining Lemma 3.1 and 3.2, Theorem 1.1 follows.

4 Proof of Theorem 1.2

The problem is clearly in NP. To show NP-hardness, we provide a reduction from Positive 1-in-3-SAT, which is known to be NP-hard [2]. An instance of Positive 1-in-3-SAT is a set of variables $X = \{x_1, x_2, \ldots, x_n\}$ and a set of clauses $C = \{C_i \mid i = 1, 2, \ldots, m\}$, such that each $C_i = (l_{i,1} \vee l_{i,2} \vee l_{i,3})$ consists of 3 positive literals and each literal $l_{i,k}$ is x_p for some $x_p \in X$. The problem is to determine whether there exists a truth assignment to the variables in X such that $\theta = C_1 \wedge C_2 \wedge \ldots \wedge C_m$ is satisfiable with exactly one true literal per clause.

Gadgets. The weak negator gadget (see Figure 1) and the strong negator gadget (see Figure 2) have been described in Section 3.

The *literal gadget* with endpoints x, y, z is presented in Figure 4. The following observation is left as a simple exercise to the reader. It has a colouring, and in every colouring it has at least two endpoints in the P_3-free part.

The *propagator gadget* with endpoints u, v, w is presented in Figure 4. The following observation is left as a simple exercise to the reader. It has a colouring, and in every colouring it has exactly one or three endpoints in the P_3-free part.

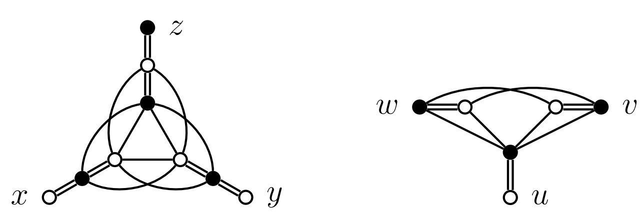

Figure 4: The literal gadget (left) with endpoints x, y, z and the propagator gadget (right) with endpoints u, v, w along with a colouring where the white vertices are in the P_3-free part and the black vertices are in K_3-free part. Note that the propagator gadget is not symmetric.

Given an instance of Positive 1-in-3-SAT, we construct the following reduction graph.

Construction. For each variable x that appears in θ create a variable component V_x (see Figure 3), which is a cycle of length $2m_x$ (where m_x is the number of occurences of x) whose edges are replaced by a strong negator gadget. We number its vertices from 1 to $2m_x$ in clockwise traversal. Its even numbered vertices, denoted by *literal vertices*, are labelled $l_{x,1}, \ldots, l_{x,m_x}$, and its odd numbered vertices, denoted by *propagator vertices*, are labelled $p_{x,1}, \ldots, p_{x,m_x}$. For a clause $C = (x \vee y \vee z)$ where x, y and z are the i'th, j'th and k'th occurrence, respectively, create a copy H_C of the literal gadget whose endpoints are identified with $l_{x,i}, l_{y,j}$ and $l_{z,k}$, and a copy R_C of the propagator gadget whose endpoints are identified with $p_{x,i}, p_{y,j}$ and $p_{z,k}$. H_C and R_C are said to be the *associated literal gadget* and *associated propagator gadget*, respectively, of C. Denote the graph obtained in this way by G.

Lemma 4.1. *G is colourable if and only if θ is satisfiable with exactly one true literal per clause.*

Proof. By the property of the strong negator gadget, in any colouring of a variable component, the set of literal vertices receive one colour and the set of propagator vertices receive the other colour.

Suppose θ is satisfiable with exactly one true literal per clause. If $\theta(x)$ is true, let the literal vertices and the propagator vertices in V_x be in the K_3-free part and P_3-free part, respectively. If $\theta(x)$ is false, let the literal vertices and propagator vertices in V_x be in the P_3-free part and K_3-free part, respectively. Clearly, the variable components are colourable. Consider the associated literal gadget H_C and associated propagator gadget R_C of a clause C. It follows by our colouring that H_C has two endpoints in the P_3-free part, and R_C has one endpoint in the P_3-free part. Consequently, H_C and R_C are colourable. To ensure that no vertices in the P_3-free part induce a P_3, colour each strong negator gadget N in such a way that its endpoint in the P_3-free part has no neighbour in N in the P_3-free part.

Conversely, suppose G is colourable. If the literal vertices in V_x are in the K_3-free part we set the variable $\theta(x)$ to true. Otherwise, we set $\theta(x)$ to false. Consider the associated literal gadget H_C and associated propagator gadget R_C of a clause C. By contradiction, suppose all endpoints of H_C are in the P_3-free part. By the property of the construction, the endpoints of R_C are in the K_3-free part, which contradicts the property of the propagator gadget. It follows that exactly one endpoint of H_C is in the K_3-free part, in which case C has exactly one true literal as required. □

Lemma 4.2. *G has no odd hole.*

Proof sketch. It is routine to verify that the weak negator gadget, the strong negator gadget, the literal gadget, and the propagator gadget contain no odd hole. Also, each induced path P (or Q) between the endpoints of a literal gadget (or propagator gadget, respectively) has even length. P and Q are said to be a literal path and propagator path, respectively.

Let C be an induced cycle of length at least 4 in G. If C occurs in a variable component or a gadget, then C has even length. Otherwise, let S be the set of literal and propagator paths occurring in C, and let $C' = C \setminus S$. By the property of the construction, C' forms a set of disjoint paths where each path is composed of alternating propagator and literal vertices.

We claim that there exists an even number of odd length paths in C'. Suppose otherwise, and let $J = \{J_1, J_2, \ldots, J_{2h+1}\} \subseteq C'$ be the the set of odd length paths occurring in C'. Without loss of generality, let this be the order in which they appear along a clockwise traversal of C. Given that any $J_i \in J$ has odd length, it follows that J_i has one literal vertex endpoint and one propagator vertex endpoint. From this, given that the endpoints of the even length paths in C are of the same type, observe that for any $J_s, J_q \in J$, s and q have the same parity if and only if J_s and J_q both have the same order of appearance of the type of their endpoints along the traversal. Consequently, to get from J_{2l+1} to J_1, there must exist another odd length path J_{2l+2} in C', a contradiction. Hence, the sum of the length of the paths in C' is even. Since the sum of the length of the paths in S is clearly even, it follows that C has even length. □

Observation 4.1. *the weak negator gadget, the strong negator gadget, the literal gadget, and the propagator gadget are short-chorded.*

Lemma 4.3. *Let C be an odd cycle of length at least 5 in G. If C is not a subgraph of a weak negator gadget, a strong negator gadget, a literal gadget, or a propagator gadget, then at least one of the following holds:*

(i) C contains an even length path connecting the endpoints of a weak negator gadget.

(ii) C contains an even length path connecting the endpoints of a strong negator gadget.

(iii) C contains an odd length path connecting two endpoints of a literal gadget.

(iv) C contains an odd length path connecting two endpoints of a propagator gadget.

Proof. If the part of C within each gadget is induced then C has even length by Lemma 4.2. So there exists a part of C within a gadget that has a length whose parity differs from the length of the induced path connecting the endpoints of the gadget under consideration. In any gadget, all induced paths between endpoints have lengths of the same parity. Namely odd for the weak and the strong negator gadgets, and even for the literal and the propagator gadgets. This completes the proof. □

Lemma 4.4. *G is short-chorded.*

Proof. Each of the four paths from Lemma 4.3 has a short chord. Together with Observation 4.1 we get the desired result. □

By conjoining Lemmas 4.1 and 4.4, Theorem 1.2 follows.

5 Further Work

One direction would be to investigate the complexity of the (P_3-free, K_3-free)-colouring problem in other graph classes.

More generally, we can ask for the complexity of ($\mathcal{F}$, $\mathcal{Q}$)-colouring a graph for various properties $\mathcal{F}$ and $\mathcal{Q}$. When $\mathcal{F}$ and $\mathcal{Q}$ are additive induced hereditary properties, the problem is NP-complete on general graphs [3]. Can complexity dichotomies for this problem be established in special classes of graphs?

Acknowledgments

The authors would like to express their gratitude to Matthew Johnson for carefully reviewing the paper and improving its presentation and readability. The first author was supported by the Durham Doctoral Scholarship (DDS).

References

[1] A. Lubiw, Short-chorded and perfect graphs. *Journal Combinatorial Theory B* **51** (1991) 24–33.

[2] T.J. Schaefer, The complexity of satisfiability problems. *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing* (STOC 1978, San Diego, Calif.), 216–226, ACM, New York, 1978.

[3] A. Farrugia, Vertex-partitioning into fixed additive induced-hereditary properties is NP-hard. *The Electronic Journal of Combinatorics* **11** (2004) R46.

[4] R. Diestel, Graph theory, Third edition. *Springer* (2006).

[5] D. Lichtenstein, Planar formulae and their uses. *SIAM Journal on Computing* **11** (1982) 329–343.

[6] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas, The strong perfect graph theorem. *Annals of Mathematics* **164** (2006) 51–229.