Algorithms and Almost Tight Results for 3-Colorability of Small Diameter Graphs

George B. Mertzios[1,⋆] and Paul G. Spirakis[2,⋆⋆]

[1] School of Engineering and Computing Sciences, Durham University, UK
george.mertzios@durham.ac.uk
[2] Computer Technology Institute and University of Patras, Greece
spirakis@cti.gr

Abstract. The 3-coloring problem is well known to be NP-complete. It is also well known that it remains NP-complete when the input is restricted to graphs with diameter 4. Moreover, assuming the Exponential Time Hypothesis (ETH), 3-coloring can not be solved in time $2^{o(n)}$ on graphs with n vertices and diameter at most 4. In spite of the extensive studies of the 3-coloring problem with respect to several basic parameters, the complexity status of this problem on graphs with small diameter, i.e. with diameter at most 2, or at most 3, has been a longstanding and challenging open question. In this paper we investigate graphs with small diameter. For graphs with diameter at most 2, we provide the first subexponential algorithm for 3-coloring, with complexity $2^{O(\sqrt{n \log n})}$. Furthermore we present a subclass of graphs with diameter 2 that admits a polynomial algorithm for 3-coloring. For graphs with diameter at most 3, we establish the complexity of 3-coloring, even for the case of triangle-free graphs. Namely we prove that for every $\varepsilon \in [0, 1)$, 3-coloring is NP-complete on triangle-free graphs of diameter 3 and radius 2 with n vertices and minimum degree $\delta = \Theta(n^{\varepsilon})$. Moreover, assuming ETH, we use three different amplification techniques of our hardness results, in order to obtain for every $\varepsilon \in [0, 1)$ subexponential asymptotic lower bounds for the complexity of 3-coloring on triangle-free graphs with diameter 3 and minimum degree $\delta = \Theta(n^{\varepsilon})$. Finally, we provide a 3-coloring algorithm with running time $2^{O(\min\{\delta\Delta,\ \frac{n}{\delta}\log\delta\})}$ for arbitrary graphs with diameter 3, where n is the number of vertices and δ (resp. Δ) is the minimum (resp. maximum) degree of the input graph. To the best of our knowledge, this algorithm is the first subexponential algorithm for graphs with $\delta = \omega(1)$ and for graphs with $\delta = O(1)$ and $\Delta = o(n)$. Due to the above lower bounds of the complexity of 3-coloring, the running time of this algorithm is asymptotically almost tight when the minimum degree of the input graph is $\delta = \Theta(n^{\varepsilon})$, where $\varepsilon \in [\frac{1}{2}, 1)$.

Keywords: 3-coloring, graph diameter, graph radius, subexponential algorithm, NP-complete, exponential time hypothesis.

⋆ Partially supported by EPSRC Grant EP/G043434/1.
⋆⋆ Partially supported by the ERC EU Project RIMACO and by the EU IP FET Project MULTIPLEX.

P. van Emde Boas et al. (Eds.): SOFSEM 2013, LNCS 7741, pp. 332–343, 2013.

1 Introduction

A *proper k-coloring* (or *k-coloring*) of a graph G is an assignment of k different colors to the vertices of G, such that no two adjacent vertices receive the same color. That is, a k-coloring is a partition of the vertices of G into k independent sets. The corresponding *k-coloring problem* is the problem of deciding whether a given graph G admits a k-coloring of its vertices, and to compute one if it exists. Furthermore, the minimum number k of colors for which there exists a k-coloring is denoted by $\chi(G)$ and is termed the *chromatic number* of G. The *minimum coloring problem* is to compute the chromatic number of a given graph G, and to compute a $\chi(G)$-coloring of G if one exists.

One of the most well known complexity results is that the k-coloring problem is NP-complete for every $k \geq 3$, while it can be solved in polynomial time for $k = 2$ [10]. Therefore, since graph coloring has numerous applications besides its theoretical interest, there has been considerable interest in studying how several graph parameters affect the tractability of the k-coloring problem, where $k \geq 3$. In view of this, the complexity status of the coloring problem has been established for many graph classes. It has been proved that 3-coloring remains NP-complete even when the input graph is a line graph [13], a triangle-free graph with maximum degree 4 [18], or a planar graph with maximum degree 4 [10].

On the positive side, one of the most famous result in this context has been that the minimum coloring problem can be solved in polynomial time for perfect graphs using the ellipsoid method [11]. Furthermore, polynomial algorithms for 3-coloring have been also presented for classes of non-perfect graphs, such as AT-free graphs [23] and P_6-free graphs [22] (i.e. graphs that do not contain any path on 6 vertices as an induced subgraph). Furthermore, although the minimum coloring problem is NP-complete on P_5-free graphs, the k-coloring problem is polynomial on these graphs for every fixed k [12]. Courcelle's celebrated theorem states that every problem definable in Monadic Second-Order logic (MSO) can be solved in linear time on graphs with bounded treewidth [8], and thus also the coloring problem can be solved in linear time on such graphs.

For the cases where 3-coloring is NP-complete, considerable attention has been given to devise exact algorithms that are faster than the brute-force algorithm (see e.g. the recent book [9]). In this context, asymptotic lower bounds of the time complexity have been provided for the main NP-complete problems, based on the *Exponential Time Hypothesis (ETH)* [14,15]. ETH states that there exists no deterministic algorithm that solves the 3SAT problem in time $2^{o(n)}$, given a boolean formula with n variables. In particular, assuming ETH, 3-coloring can not be solved in time $2^{o(n)}$ on graphs with n vertices, even when the input is restricted to graphs with diameter 4 and radius 2 (see [17,21]). Therefore, since it is assumed that no subexponential $2^{o(n)}$ time algorithms exist for 3-coloring, most attention has been given to decrease the multiplicative factor of n in the exponent of the running time of exact exponential algorithms, see e.g. [4,9,20].

One of the most central notions in a graph is the distance between two vertices, which is the basis of the definition of other important parameters, such as the diameter, the eccentricity, and the radius of a graph. For these graph parameters, it is known that 3-coloring is NP-complete on graphs with diameter at most 4 (see e.g. the standard proof of [21]). Furthermore, it is straightforward to check that k-coloring is NP-complete for graphs with diameter at most 2, for every $k \geq 4$: we can reduce 3-coloring on arbitrary graphs to 4-coloring on graphs with diameter 2, just by introducing to an arbitrary graph a new vertex that is adjacent to all others.

In contrast, in spite of the extensive studies of the 3-coloring problem with respect to several basic parameters, the complexity status of this problem on graphs with small diameter, i.e. with diameter at most 2 or at most 3, has been a longstanding and challenging open question, see e.g. [5, 7, 16]. The complexity status of 3-coloring is open also for triangle-free graphs of diameter 2 and of diameter 3. It is worth mentioning here that a graph is triangle-free and of diameter 2 if and only if it is a maximal triangle free graph. Moreover, it is known that 3-coloring is NP-complete for triangle-free graphs [18], however it is not known whether this reduction can be extended to maximal triangle free graphs. Another interesting result is that almost all graphs have diameter 2 [6]; however, this result can not be used in order to establish the complexity of 3-coloring for graphs with diameter 2.

Our Contribution. In this paper we provide subexponential algorithms and hardness results for the 3-coloring problem on graphs with low diameter, i.e. with diameter 2 and 3. As a preprocessing step, we first present two reduction rules that we apply to an arbitrary graph G, such that the resulting graph G' is 3-colorable if and only G is 3-colorable. We call the resulting graph *irreducible* with respect to these two reduction rules. We use these reduction rules to reduce the size of the given graph and to simplify the algorithms that we present.

For graphs with diameter at most 2, we first provide a subexponential algorithm for 3-coloring with running time $2^{O(\min\{\delta, \frac{n}{\delta}\log\delta\})}$, where n is the number of vertices and δ is the minimum degree of the input graph. This algorithm is simple and has worst-case running time $2^{O(\sqrt{n \log n})}$, which is asymptotically the same as the currently best known time complexity of the graph isomorphism problem [3]. To the best of our knowledge, this algorithm is the first subexponential algorithm for graphs with diameter 2. We demonstrate that this is indeed the worst-case of our algorithm by providing, for every $n \geq 1$, a 3-colorable graph $G_n = (V_n, E_n)$ with $\Theta(n)$ vertices, such that G_n has diameter 2 and both its minimum degree and the size of a minimum dominating set is $\Theta(\sqrt{n})$. In addition, this graph is triangle-free and irreducible with respect to the above two reduction rules. Finally, we present a subclass of graphs with diameter 2, called *locally decomposable* graphs, which admits a polynomial algorithm for 3-coloring. In particular, we prove that whenever an irreducible graph G with diameter 2 has at least one vertex v such that $G - N(v) - \{v\}$ is disconnected, then 3-coloring on G can be decided in polynomial time.

For graphs with diameter at most 3, we establish the complexity of deciding 3-coloring, even for the case of triangle-free graphs. Namely we prove that 3-coloring is NP-complete on irreducible and triangle-free graphs with diameter 3 and radius 2, by providing a reduction from 3SAT. In addition, we provide a 3-coloring algorithm with running time $2^{O(\min\{\delta\Delta,\ \frac{n}{\delta}\log\delta\})}$ for arbitrary graphs with diameter 3, where n is the number of vertices and δ (resp. Δ) is the minimum (resp. maximum) degree of the input graph. To the best of our knowledge, this algorithm is the first subexponential algorithm for graphs with $\delta = \omega(1)$ and for graphs with $\delta = O(1)$ and $\Delta = o(n)$. Table 1 summarizes the current state of the art of the complexity of k-coloring, as well as our algorithmic and NP-completeness results.

Table 1. Current state of the art and our algorithmic and NP-completeness results for k-coloring on graphs with diameter $diam(G)$. Our results are indicated by an asterisk.

$k \setminus diam(G)$	2	3	≥ 4
3	(*) $2^{O(\min\{\delta, \frac{n}{\delta}\log\delta\})}$ time algorithm (*) polynomial alg. for locally decomposable graphs	(*) NP-complete for min. degree $\delta = \Theta(n^{\varepsilon})$, for every $\varepsilon \in [0, 1)$, even if $rad(G) = 2$ and G is triangle-free (*) $2^{O(\min\{\delta\Delta,\ \frac{n}{\delta}\log\delta\})}$ time algorithm	NP-complete [21], no $2^{o(n)}$ algorithm
≥ 4	NP-complete	NP-complete	NP-complete

Furthermore, we provide three different amplification techniques that extend our hardness results for graphs with diameter 3. In particular, we first show that 3-coloring is NP-complete on irreducible and triangle-free graphs G of diameter 3 and radius 2 with n vertices and minimum degree $\delta(G) = \Theta(n^{\varepsilon})$, for every $\varepsilon \in [\frac{1}{2}, 1)$ and that, for such graphs, there exists no algorithm for 3-coloring with running time $2^{o(\frac{n}{\delta})} = 2^{o(n^{1-\varepsilon})}$, assuming ETH. This lower bound is asymptotically almost tight, due to our above algorithm with running time $2^{O(\frac{n}{\delta}\log\delta)}$, which is subexponential when $\delta(G) = \Theta(n^{\varepsilon})$ for some $\varepsilon \in [\frac{1}{2}, 1)$. With our second amplification technique, we show that 3-coloring remains NP-complete also on irreducible and triangle-free graphs G of diameter 3 and radius 2 with n vertices and minimum degree $\delta(G) = \Theta(n^{\varepsilon})$, for every $\varepsilon \in [0, \frac{1}{2})$. Moreover, we prove that for such graphs, when $\varepsilon \in [0, \frac{1}{3})$, there exists no algorithm for 3-coloring with running time $2^{o(\sqrt{\frac{n}{\delta}})} = 2^{o(n^{(\frac{1-\varepsilon}{2})})}$, assuming ETH. Finally, with our third amplification technique, we prove that for such graphs, when $\varepsilon \in [\frac{1}{3}, \frac{1}{2})$, there exists no algorithm for 3-coloring with running time $2^{o(\delta)} = 2^{o(n^{\varepsilon})}$, assuming ETH. Table 2 summarizes our lower time complexity bounds for 3-coloring on irreducible and triangle-free graphs with diameter 3 and radius 2, parameterized by their minimum degree δ.

Organization of the Paper. We provide in Section 2 the necessary notation and terminology, as well as our two reduction rules and the notion of an irreducible graph. In Sections 3 and 4 we present our results for graphs with diameter 2 and 3, respectively. Detailed proofs have been omitted due to space limitations; a full version can be found in [19].

Table 2. Our lower time complexity bounds for deciding 3-coloring on irreducible and triangle-free graphs G with n vertices, diameter 3, radius 2, and minimum degree $\delta(G) = \Theta(n^{\varepsilon})$, where $\varepsilon \in [0, 1)$, assuming ETH. The lower bound for $\varepsilon \in [\frac{1}{2}, 1)$ is asymptotically almost tight, as there exists an algorithm for arbitrary graphs with diameter 3 with running time $2^{O(\frac{n}{\delta}\log\delta)} = 2^{O(n^{1-\varepsilon}\log n)}$ by Theorem 4.

$\delta(G) = \Theta(n^{\varepsilon})$:	$0 \leq \varepsilon < \frac{1}{3}$	$\frac{1}{3} \leq \varepsilon < \frac{1}{2}$	$\frac{1}{2} \leq \varepsilon < 1$
Lower time complexity bound:	no $2^{o(n^{(\frac{1-\varepsilon}{2})})}$ time algorithm	no $2^{o(n^{\varepsilon})}$ time algorithm	no $2^{o(n^{1-\varepsilon})}$ time algorithm

2 Preliminaries and Notation

In this section we provide some notation and terminology, as well as two reduction (or "cleaning") rules that can be applied to an arbitrary graph G. Throughout the article, we assume that any given graph G of low diameter is *irreducible* with respect to these two reduction rules, i.e. that these reduction rules have been iteratively applied to G until they can not be applied any more. Note that the iterative application of these reduction rules on a graph with n vertices can be done in time polynomial in n.

Notation. We consider in this article simple undirected graphs with no loops or multiple edges. In a graph G, the edge between vertices u and v is denoted by uv. Given a graph $G = (V, E)$ and a vertex $u \in V$, denote by $N(u) = \{v \in V : uv \in E\}$ the set of neighbors (or the *open neighborhood*) of u and by $N[u] = N(u) \cup \{u\}$ the *closed neighborhood* of u. Whenever the graph G is not clear from the context, we will write $N_G(u)$ and $N_G[u]$, respectively. Denote by $\deg(u) = |N(u)|$ the *degree* of u in G and by $\delta(G) = \min\{\deg(u) : u \in V\}$ the *minimum degree* of G. Let u and v be two non-adjacent vertices of G. Then, u and v are called (false) *twins* if they have the same set of neighbors, i.e. if $N(u) = N(v)$. Furthermore, we call the vertices u and v *siblings* if $N(u) \subseteq N(v)$ or $N(v) \subseteq N(u)$; note that two twins are always siblings.

Given a graph $G = (V, E)$ and two vertices $u, v \in V$, we denote by $d(u, v)$ the *distance* of u and v, i.e. the length of a shortest path between u and v in G. Furthermore, we denote by $diam(G) = \max\{d(u, v) : u, v \in V\}$ the *diameter* of G and by $rad(G) = \min_{u\in V}\{\max\{d(u, v) : v \in V\}\}$ the *radius* of G. Given a subset $S \subseteq V$, $G[S]$ denotes the *induced* subgraph of G on the vertices in S. We denote for simplicity by $G - S$ the induced subgraph $G[V \setminus S]$ of G. A complete graph (i.e. clique) with t vertices is denoted by K_t. A graph G that contains no K_t as an induced subgraph is called *K_t-free*. Furthermore, a subset $D \subseteq V$ is a *dominating set* of G if every vertex of $V \setminus D$ has at least one neighbor in D. For simplicity, we refer in the remainder of the article to a proper k-coloring of a graph G just as a *k-coloring* of G. Throughout the article we perform several times the *merging* operation of two (or more) independent vertices, which is defined as follows: we *merge* the independent vertices $u_1, u_2, \ldots, u_t$ when we replace them by a new vertex u_0 with $N(u_0) = \cup_{i=1}^{k} N(u_i)$.

Observe that, whenever a graph G contains a clique K_4 with four vertices as an induced subgraph, then G is not 3-colorable. Furthermore, we can check easily in polynomial time (e.g. with brute-force) whether a given graph G contains a K_4. Therefore we assume in the following that all given graphs are K_4-free. Furthermore, since a graph is 3-colorable if and only if all its connected components are 3-colorable, we assume in the following that all given graphs are connected. In order to present our two reduction rules of an arbitrary K_4-free graph G, recall first that the *diamond* graph is a graph with 4 vertices and 5 edges, i.e. it consists of a K_4 without one edge. Suppose that four vertices u_1, u_2, u_3, u_4 of a given graph $G = (V, E)$ induce a diamond graph, and assume without loss of generality that $u_1u_2 \notin E$. Then, it is easy to see that in any 3-coloring of G (if such exists), u_1 and u_2 obtain necessarily the same color. Therefore we can merge u_1 and u_2 into one vertex, as the next reduction rule states, and the resulting graph is 3-colorable if and only if G is 3-colorable.

Reduction Rule 1 (diamond elimination). *Let $G = (V, E)$ be a K_4-free graph. If the quadruple $\{u_1, u_2, u_3, u_4\}$ of vertices in G induces a diamond graph, where $u_1u_2 \notin E$, then merge vertices u_1 and u_2.*

Note that, after performing a diamond elimination in a K_4-free graph G, we may introduce a new K_4 in the resulting graph. Suppose now that a graph G has a pair of siblings u and v and assume without loss of generality that $N(u) \subseteq N(v)$. Then, we can extend any proper 3-coloring of $G - \{u\}$ (if such exists) to a proper 3-coloring of G by assigning to u the same color as v. Therefore, we can remove vertex u from G, as the next reduction rule states, and the resulting graph $G - \{u\}$ is 3-colorable if and only if G is 3-colorable.

Reduction Rule 2 (siblings elimination). *Let $G = (V, E)$ be a K_4-free graph and $u, v \in V$, such that $N(u) \subseteq N(v)$. Then remove u from G.*

Definition 1. *Let $G = (V, E)$ be a K_4-free graph. If neither Reduction Rule 1 nor Reduction Rule 2 can be applied to G, then G is* irreducible.

Due to Definition 1, a K_4-free graph is irreducible if and only if it is diamond-free and siblings-free. Given a K_4-free graph G with n vertices, clearly we can iteratively execute Reduction Rules 1 and 2 in time polynomial on n, until we either find a K_4 or none of the Reduction Rules 1 and 2 can be further applied. If we find a K_4, then clearly the initial graph G is not 3-colorable. Otherwise, we transform G in polynomial time into an irreducible (K_4-free) graph G' of smaller or equal size, such that G' is 3-colorable if and only if G is 3-colorable.

Observation 1. *Let $G = (V, E)$ be a connected K_4-free graph and $G' = (V', E')$ be the irreducible graph obtained from G. If G' has more than two vertices, then $\delta(G') \geq 2$, $diam(G') \leq diam(G)$, $rad(G') \leq rad(G)$, and G' is 3-colorable if and only if G is 3-colorable. Moreover, for every $u \in V'$, $N_{G'}(u)$ induces in G' a graph with maximum degree 1.*

3 Algorithms for 3-Coloring on Graphs with Diameter 2

In this section we present our results on graphs with diameter 2. In particular, we provide in Section 3.1 our subexponential algorithm for 3-coloring on such graphs. We then provide, for every n, an example of an irreducible and triangle-free graph G_n with $\Theta(n)$ vertices and diameter 2, which is 3-colorable, has minimum dominating set of size $\Theta(\sqrt{n})$, and its minimum degree is $\delta(G_n) = \Theta(\sqrt{n})$. Furthermore, we provide in Section 3.2 our polynomial algorithm for irreducible graphs G with diameter 2, which have at least one vertex v, such that $G - N[v]$ is disconnected.

3.1 An $2^{O(\sqrt{n \log n})}$-Time Algorithm for Any Graph with Diameter 2

We first provide in the next lemma a well known algorithm that decides the 3-coloring problem on an arbitrary graph G, using a dominating set (DS) of G.

Lemma 1 (the DS-approach). *Let $G = (V, E)$ be a graph and $D \subseteq V$ be a dominating set of G. Then, the 3-coloring problem can be decided in $O^*(3^{|D|})$ time on G.*

In an arbitrary graph G with n vertices and minimum degree δ, it is well known how to construct in polynomial time a dominating set D with cardinality $|D| \leq n\frac{1+\ln(\delta+1)}{\delta+1}$ [2] (see also [1]). On the other hand, in a graph with diameter 2, the neighborhood of every vertex is a dominating set. Thus we can use Lemma 1 to provide in the next theorem an improved 3-coloring algorithm for the case of graphs with diameter 2.

Theorem 1. *Let $G = (V, E)$ be an irreducible graph with n vertices. Let $diam(G) = 2$ and δ be the minimum degree of G. Then, the 3-coloring problem can be decided in $2^{O(\min\{\delta, \frac{n}{\delta} \log \delta\})}$ time on G.*

Corollary 1. *Let $G = (V, E)$ be an irreducible graph with n vertices and let $diam(G) = 2$. Then, the 3-coloring problem can be decided in $2^{O(\sqrt{n \log n})}$ time on G.*

Given the statements of Lemma 1 and Theorem 1, a question that arises naturally is whether the worst case complexity of the algorithm of Theorem 1 is indeed $2^{O(\sqrt{n \log n})}$ (as given in Corollary 1). That is, do there exist 3-colorable irreducible graphs G with n vertices and $diam(G) = 2$, such that both $\delta(G)$ and the size of the minimum dominating set of G are $\Theta(\sqrt{n \log n})$, or close to this value? We answer this question to the affirmative, thus proving that, in the case of 3-coloring of graphs with diameter 2, our analysis of the DS-approach (cf. Lemma 1 and Theorem 1) is asymptotically almost tight. In particular, we provide in the next theorem for every n an example of an irreducible 3-colorable graph G_n with $\Theta(n)$ vertices and $diam(G_n) = 2$, such that both $\delta(G_n)$ and the size of the minimum dominating set of G are $\Theta(\sqrt{n})$. In addition, each of these graphs G_n is triangle-free, as the next theorem states. The construction of the graphs G_n is based on a suitable and interesting matrix arrangement of the vertices of G_n.

Theorem 2. *Let $n \geq 1$. Then there exists an* irreducible *and* triangle-free 3-colorable *graph $G_n = (V_n, E_n)$ with $\Theta(n)$ vertices, where $diam(G_n) = 2$ and $\delta(G_n) = \Theta(\sqrt{n})$. Furthermore, the size of the minimum dominating set of G_n is $\Theta(\sqrt{n})$.*

3.2 A Tractable Subclass of Graphs with Diameter 2

In this section we present a subclass of graphs with diameter 2, which admits an efficient algorithm for 3-coloring. We first introduce the definition of *locally decomposable* graphs.

Definition 2. *Let $G = (V, E)$ be a graph. If there exists a vertex $v_0 \in V$ such that $G - N[v_0]$ is disconnected, then G is a* locally decomposable *graph.*

We prove in Theorem 3 that, given an irreducible and locally decomposable graph G with $diam(G) = 2$, we can decide 3-coloring on G in polynomial time. Note here that there exist instances of K_4-free graphs G with diameter 2, for which $G - N[v]$ is connected for every vertex v of G, but in the irreducible graph G' obtained by G (by iteratively applying the Reduction Rules 1 and 2), $G' - N_{G'}[v_0]$ becomes disconnected for some vertex v_0 of G'. That is, G' may be locally decomposable, although G is not. Therefore, if we provide as input to the algorithm of Theorem 3 the irreducible graph G' instead of G, this algorithm decides in polynomial time the 3-coloring problem on G' (and thus also on G). The crucial idea of this algorithm is that, since G is irreducible and locally decomposable, we can prove in Theorem 3 that every connected component of $G - N[v_0]$ is bipartite, and that in every proper 3-coloring of G, all connected components of $G - N[v_0]$ are colored using only two colors.

Theorem 3. *Let $G = (V, E)$ be an irreducible graph with n vertices and $diam(G) = 2$. If G is a locally decomposable graph, then we can decide 3-coloring on G in time polynomial on n.*

A question that arises now naturally by Theorem 3 is whether there exist any irreducible 3-colorable graph $G = (V, E)$ with $diam(G) = 2$, for which $G - N[v]$ is connected for *every* $v \in V$. A negative answer to this question would imply that we can decide the 3-coloring problem on *any* graph with diameter 2 in polynomial time using the algorithm of Theorem 3. However, the answer to that question is positive: for every $n \geq 1$, the graph $G_n = (V_n, E_n)$ that has been presented in Theorem 2 is irreducible, 3-colorable, has diameter 2, and $G_n - N[v]$ is connected for every $v \in V_n$. Therefore, the algorithm of Theorem 3 can not be used in a trivial way to decide in polynomial time the 3-coloring problem for an arbitrary graph of diameter 2. We leave the tractability of the 3-coloring problem of arbitrary diameter-2 graphs as an open problem.

4 Almost Tight Results for Graphs with Diameter 3

4.1 An $2^{O(\min\{\delta\Delta,\ \frac{n}{\delta}\log\delta\})}$-Time Algorithm for Any Graph with Diameter 3

In the next theorem we use the DS-approach of Lemma 1 to provide an improved 3-coloring algorithm for the case of graphs with diameter 3.

Theorem 4. *Let $G = (V, E)$ be an irreducible graph with n vertices and $diam(G) = 3$. Let δ and Δ be the minimum and the maximum degree of G, respectively. Then, the 3-coloring problem can be decided in $2^{O(\min\{\delta\Delta,\ \frac{n}{\delta}\log\delta\})}$ time on G.*

To the best of our knowledge, the algorithm of Theorem 4 is the first subexponential algorithm for graphs with diameter 3, whenever $\delta = \omega(1)$, as well as whenever $\delta = O(1)$ and $\Delta = o(n)$. As we will later prove in Section 4.3, the running time provided in Theorem 4 is asymptotically almost tight whenever $\delta = \Theta(n^{\varepsilon})$, for any $\varepsilon \in [\frac{1}{2}, 1)$.

4.2 The 3-Coloring Problem Is NP-Complete on Graphs with Diameter 3 and Radius 2

In this section we provide a reduction from the 3SAT problem to the 3-coloring problem of triangle-free graphs with diameter 3 and radius 2. Let ϕ be a 3-CNF formula with n variables $x_1, x_2, \ldots, x_n$ and m clauses $\alpha_1, \alpha_2, \ldots, \alpha_m$. We can assume in the following without loss of generality that each clause has three distinct literals. We now construct an irreducible and triangle-free graph $H_\phi = (V_\phi, E_\phi)$ with diameter 3 and radius 2, such that ϕ is satisfiable if and only if H_ϕ is 3-colorable. Before we construct H_ϕ, we first construct an auxiliary graph $G_{n,m}$ that depends only on the number n of the variables and the number m of the clauses in ϕ, rather than on ϕ itself.

We construct the graph $G_{n,m} = (V_{n,m}, E_{n,m})$ as follows. Let v_0 be a vertex with $8m$ neighbors $v_1, v_2, \ldots, v_{8m}$, which induce an independent set. Consider also the sets $U = \{u_{i,j} : 1 \leq i \leq n + 5m, 1 \leq j \leq 8m\}$ and $W = \{w_{i,j} : 1 \leq i \leq n+5m, 1 \leq j \leq 8m\}$ of vertices. Each of these sets has $(n + 5m)8m$ vertices. The set $V_{n,m}$ of vertices of $G_{n,m}$ is defined as $V_{n,m} = U \cup W \cup \{v_0, v_1, v_2, \ldots, v_{8m}\}$. That is, $|V_{n,m}| = 2 \cdot (n + 5m)8m + 8m + 1$, and thus $|V_{n,m}| = \Theta(m^2)$, since $m = \Omega(n)$.

The set $E_{n,m}$ of the edges of $G_{n,m}$ is defined as follows. Define first for every $j \in \{1, 2, \ldots, 8m\}$ the subsets $U_j = \{u_{1,j}, u_{2,j}, \ldots, u_{n+5m,j}\}$ and $W_j = \{w_{1,j}, w_{2,j}, \ldots, w_{n+5m,j}\}$ of U and W, respectively. Then define $N(v_j) = \{v_0\} \cup U_j \cup W_j$ for every $j \in \{1, 2, \ldots, 8m\}$, where $N(v_j)$ denotes the set of neighbors of vertex v_j in $G_{n,m}$. For simplicity of the presentation, we arrange the vertices of $U \cup W$ on a rectangle matrix of size $2(n + 5m) \times 8m$, cf. Figure 1(a). In this matrix arrangement, the (i, j)th element is vertex $u_{i,j}$ if $i \leq n + 5m$, and vertex $w_{i-n-5m,j}$ if $i \geq n + 5m + 1$. In particular, for every $j \in \{1, 2, \ldots, 8m\}$, the jth column of this matrix contains exactly the vertices of $U_j \cup W_j$, cf. Figure 1(a). Note that, for every $j \in \{1, 2, \ldots, 8m\}$, vertex v_j is adjacent to all vertices of the jth column of this matrix. Denote now by $\ell_i = \{u_{i,1}, u_{i,2}, \ldots, u_{i,8m}\}$ (resp. $\ell'_i = \{w_{i,1}, w_{i,2}, \ldots, w_{i,8m}\}$) the ith (resp. the $(n+5m+i)$th) row of this matrix, cf. Figure 1(a). For every $i \in \{1, 2, \ldots, n+5m\}$, the vertices of ℓ_i and of ℓ'_i induce two independent sets in $G_{n,m}$. We then add between the vertices of ℓ_i and the vertices of ℓ'_i all possible edges, except those of $\{u_{i,j}w_{i,j} : 1 \leq j \leq 8m\}$. That is, we add all possible $(8m)^2 - 8m$ edges between the vertices of ℓ_i and of ℓ'_i, such that they induce a complete bipartite graph without a perfect matching between the vertices of ℓ_i and of ℓ'_i, cf. Figure 1(b). Note by the construction of $G_{n,m}$ that both U and W are independent sets in $G_{n,m}$. Furthermore note that the minimum degree in $G_{n,m}$ is $\delta(G_{n,m}) = \Theta(m)$ and the maximum degree is $\Delta(G_{n,m}) = \Theta(n+m)$. Thus, since $m = \Omega(n)$, we have that $\delta(G_{n,m}) = \Delta(G_{n,m}) = \Theta(m)$. The construction of the graph $G_{n,m}$ is illustrated in Figure 1. Moreover, we can prove that $G_{n,m}$ has diameter 3 and radius 2, and furthermore that it is irreducible and triangle-free.

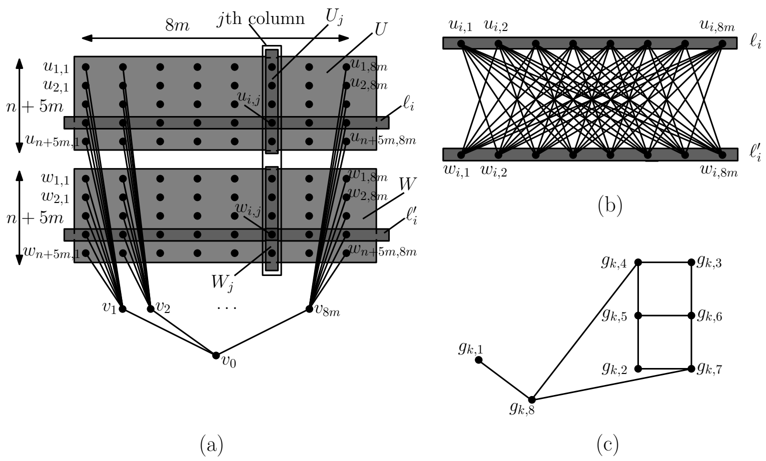

Fig. 1. (a) The $2(n+5m) \times 8m$-matrix arrangement of the vertices $U \cup W$ of $G_{n,m}$ and their connections with the vertices $\{v_0, v_1, v_2, \ldots, v_{8m}\}$, (b) the edges between the vertices of the ith row ℓ_i and the $(n+5m+i)$th row ℓ'_i in this matrix, and (c) the gadget with 8 vertices and 10 edges that we associate in H_ϕ to the clause α_k of ϕ, where $1 \leq k \leq m$.

We now construct the graph $H_\phi = (V_\phi, E_\phi)$ from ϕ by adding $10m$ edges to $G_{n,m}$ as follows. Let $k \in \{1, 2, \ldots, m\}$ and consider the clause $\alpha_k = (l_{k,1} \vee l_{k,2} \vee l_{k,3})$, where $l_{k,p} \in \{x_{i_{k,p}}, \overline{x_{i_{k,p}}}\}$ for $p \in \{1, 2, 3\}$ and $i_{k,1}, i_{k,2}, i_{k,3} \in \{1, 2, \ldots, n\}$. For this clause α_k, we add on the vertices of $G_{n,m}$ an isomorphic copy of the gadget in Figure 1(c), which has 8 vertices and 10 edges, as follows. Let $p \in \{1, 2, 3\}$. The literal $l_{k,p}$ corresponds to vertex $g_{k,p}$ of this gadget. If $l_{k,p} = x_{i_{k,p}}$, we set $g_{k,p} = u_{i_{k,p},8k+1-p}$. Otherwise, if $l_{k,p} = \overline{x_{i_{k,p}}}$, we set $g_p = w_{i_{k,p},8k+1-p}$. Furthermore, for $p \in \{4, \ldots, 8\}$, we set $g_{k,p} = u_{n+5k+4-p,8k+1-p}$.

Note that, by construction, the graphs H_ϕ and $G_{n,m}$ have the same vertex set, i.e. $V_\phi = V_{n,m}$, and that $E_{n,m} \subset E_\phi$. Therefore $diam(H_\phi) = 3$ and $rad(H_\phi) = 2$, since $diam(G_{n,m}) = 3$ and $rad(G_{n,m}) = 2$. Observe now that every positive literal of ϕ is associated to a vertex of U, while every negative literal of ϕ is associated to a vertex of W. In particular, each of the $3m$ literals of ϕ corresponds by this construction to a different column in the matrix arrangement of the vertices of $U \cup W$. If a literal of ϕ is the variable x_i (resp. the negated variable $\overline{x_i}$), where $1 \leq i \leq n$, then the vertex of U (resp. W) that is associated to this literal lies in the ith row ℓ_i (resp. in the $(n+5m+i)$th row ℓ'_i) of the matrix. Moreover, note by the above construction that each of the $8m$ vertices $\{q_{k,1}, q_{k,2}, \ldots, q_{k,8}\}_{k=1}^{m}$ corresponds to a different column in the matrix of the vertices of $U \cup W$. Finally, each of the $5m$ vertices $\{q_{k,4}, q_{k,5}, q_{k,6}, q_{k,7}, q_{k,8}\}_{k=1}^{m}$ corresponds to a different row in the matrix of the vertices of U.

Observation 2. *The gadget of Figure 1(c) has no proper* 2*-coloring, as it contains an induced cycle of length* 5*.*

Observation 3. *Consider the gadget of Figure 1(c). If we assign to vertices $g_{k,1}, g_{k,2}, g_{k,3}$ the same color, we can not extend this coloring to a proper* 3*-coloring of the gadget. Furthermore, if we assign to vertices $g_{k,1}, g_{k,2}, g_{k,3}$ in total two or three colors, then we can extend this coloring to a proper* 3*-coloring of the gadget.*

Observation 4. *For every $i \in \{1, 2, \ldots, n+5m\}$, there exists no pair of adjacent vertices in the same row ℓ_i or ℓ'_i in H_ϕ.*

Theorem 5. *The formula ϕ is satisfiable if and only if H_ϕ is* 3*-colorable.*

Moreover we can prove that H_ϕ is irreducible and triangle-free, and thus we conclude the main theorem of this section.

Theorem 6. *The* 3*-coloring problem is NP-complete on irreducible and triangle-free graphs with diameter* 3 *and radius* 2*.*

4.3 Lower Time Complexity Bounds and General NP-Completeness Results

In this section we present our three different amplification techniques of the reduction of Theorem 5. In particular, using these three amplifications we extend for every $\varepsilon \in [0, 1)$ the result of Theorem 6 (by providing both NP-completeness and lower time complexity bounds) to irreducible triangle-free graphs with diameter 3 and radius 2 and minimum degree $\delta = \Theta(n^\varepsilon)$. Our extended NP-completeness results, as well as our lower time complexity bounds are given in Tables 1 and 2. A detailed presentation of the results in this section can be found in [19].

References

1. Alon, N.: Transversal numbers of uniform hypergraphs. Graphs and Combinatorics 6, 1–4 (1990)
2. Alon, N., Spencer, J.H.: The Probabilistic Method, 3rd edn. John Wiley & Sons (2008)
3. Babai, L., Luks, E.M.: Canonical labeling of graphs. In: Proceedings of the 15th Annual ACM Symposium on Theory of Computing (STOC), pp. 171–183 (1983)
4. Beigel, R., Eppstein, D.: 3-coloring in time $O(1.3289^n)$. Journal of Algorithms 54(2), 168–204 (2005)
5. Bodirsky, M., Kára, J., Martin, B.: The complexity of surjective homomorphism problems – A survey. CoRR (2011), `http://arxiv.org/abs/1104.5257`
6. Bollobás, B.: The diameter of random graphs. Transactions of the American Mathematical Society 267(1), 41–52 (1981)
7. Broersma, H., Fomin, F.V., Golovach, P.A., Paulusma, D.: Three complexity results on coloring P_k-free graphs. In: Proceedings of the 20th International Workshop on Combinatorial Algorithms, pp. 95–104 (2009)
8. Courcelle, B.: The monadic second-order logic of graphs I: Recognizable sets of finite graphs. Information and Computation 85(1), 12–75 (1990)
9. Fomin, F.V., Kratsch, D.: Exact exponential algorithms. Texts in Theoretical Computer Science. An EATCS Series. Springer (2010)
10. Garey, M.R., Johnson, D.S.: Computers and intractability: A guide to the theory of NP-completeness. W.H. Freeman (1979)
11. Grötschel, M., Lovász, L., Schrijver, A.: Polynomial algorithms for perfect graphs. Topics on Perfect Graphs 88, 325–356 (1984)
12. Hoàng, C.T., Kamiński, M., Lozin, V.V., Sawada, J., Shu, X.: Deciding k-colorability of P_5-free graphs in polynomial time. Algorithmica 57(1), 74–81 (2010)
13. Holyer, I.: The NP-completeness of edge-coloring. SIAM Journal on Computing 10, 718–720 (1981)
14. Impagliazzo, R., Paturi, R.: On the complexity of k-SAT. Journal of Computer and System Sciences 62(2), 367–375 (2001)
15. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity? Journal of Computer and System Sciences 63(4), 512–530 (2001)
16. Kamiński, M.: Open problems from algorithmic graph theory. In: 7th Slovenian International Conference on Graph Theory (2011), `http://rutcor.rutgers.edu/~mkaminski/AGT/openproblemsAGT.pdf`
17. Lokshtanov, D., Marx, D., Saurabh, S.: Lower bounds based on the Exponential Time Hypothesis. Bulletin of the EATCS 84, 41–71 (2011)
18. Maffray, F., Preissmann, M.: On the NP-completeness of the k-colorability problem for triangle-free graphs. Discrete Mathematics 162, 313–317 (1996)
19. Mertzios, G.B., Spirakis, P.G.: Algorithms and almost tight results for 3-colorability of small diameter graphs. CoRR (2012), `http://arxiv.org/abs/1202.4665`
20. Narayanaswamy, N., Subramanian, C.: Dominating set based exact algorithms for 3-coloring. Information Processing Letters 111, 251–255 (2011)
21. Papadimitriou, C.H.: Computational complexity. Addison-Wesley (1994)
22. Randerath, B., Schiermeyer, I.: 3-colorability $\in P$ for P_6-free graphs. Discrete Applied Mathematics 136(2-3), 299–313 (2004)
23. Stacho, J.: 3-colouring AT-free graphs in polynomial time. Algorithmica (to appear)