On the hardness of determining the irregularity strength of graphs

Julien Bensmail

▶ **To cite this version:**

Julien Bensmail. On the hardness of determining the irregularity strength of graphs. Theoretical Computer Science, 2022, 937, pp.96-107. 10.1016/j.tcs.2022.09.033 . hal-03614796v2

HAL Id: hal-03614796

https://hal.science/hal-03614796v2

Submitted on 17 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the hardness of determining the irregularity strength of graphs

Julien Bensmail[a]

[a] *Université Côte d'Azur, CNRS, Inria, I3S, France*

Abstract

Let G be a graph, and $\ell : E(G) \to \{1, \dots, k\}$ be a k-labelling of G, *i.e.*, an assignment of labels from $\{1, \dots, k\}$ to the edges of G. We say that ℓ is irregular if no two distinct vertices of G are incident to the same sum of labels. The irregularity strength of G, denoted by $s(G)$, is the smallest k such that irregular k-labellings of G exist. These notions were introduced in the late 1980s as an alternative way to deal with an optimisation problem where one aims at making a graph irregular by multiplying its edges in an optimal way. Since then, the irregularity strength has received a lot of attention, focusing mainly on proving bounds and investigating side aspects and variants.

In this work, we consider the algorithmic complexity of determining the irregularity strength of a given graph. We prove that two close variants of this problem are NP-hard, which we suspect might indicate that the original problem is hard too. Namely, we prove that determining the distant irregularity strength, where only vertices within a certain distance are required to be incident to different sums of labels, and the multiset irregularity strength, where any two distinct vertices are required to be incident to different multisets of labels, are NP-hard problems.

Keywords: irregularity strength; distant irregularity strength; detectable colouring; graph labelling; algorithmic complexity.

---

1. Introduction

In this paper, we deal with the so-called **irregularity strength of graphs**, which can be defined through the following notions. Let G be a graph. A *k-labelling* $\ell : E(G) \to \{1, \dots, k\}$ of G is an assignment of *labels* from $\{1, \dots, k\}$ to the edges of G. Now, from ℓ, one can compute several metrics of interest for the vertices of G. In particular, for every vertex v, one can compute $\sigma(v)$, the *sum* of labels assigned to the edges incident to v. In case these sums of labels for the vertices turn out to distinguish all of them, *i.e.*, we have $\sigma(u) \neq \sigma(v)$ for every two distinct vertices u and v of G, then ℓ is said to be *irregular*. The *irregularity strength* of G, denoted by $s(G)$, is now the smallest k such that G admits irregular k-labellings, if any.

The irregularity strength of graphs was first introduced in late 1980s by Chartrand *et al.* in [5], as another way to deal with a particular optimisation problem where one aims at making a graph irregular by multiplying some of its edges. Among the important properties of interest, let us mention that $s(G)$ is well defined whenever G does not have K_2 as a connected component (a property we thus implicitly assume for every graph considered throughout this work). Remark also that considering the irregularity strength of non-connected graphs makes sense. In terms of magnitude, note that there is no absolute constant bounding $s(G)$ for all graphs G, as illustrated by the fact that $s(G)$ is at least the number of degree-1 vertices in G. It is known, however, that $|V(G)|$ is always an upper bound on $s(G)$, as proved by Nierhoff [12]. Most of the investigations on the topic are actually about establishing better bounds in general and for graphs fulfilling particular conditions. Apart from that, the irregularity strength of graphs ramified into many more or less distant variants, which, since then, have been studied for their own interest. A few of these variants will be mentioned later in the current work, but, for more details on this wide topic, we refer the interested reader to the dynamic survey [7] by Gallian, which gives a good insight into the vast and tremendous topic that distinguishing labellings have become over the years.

In this work, we focus on an aspect of the irregularity strength of graphs that, surprisingly, seems to be missing in the literature. This aspect is the algorithmic complexity of determining the irregularity strength of graphs. One reason why this is surprising is that, for many variants of this parameter, this question has actually been quite investigated. In particular, there exist papers dedicated solely to complexity aspects of very local variants, see e.g. [1, 6, 8], and other less local ones [3]. This being said, irregular labellings are types of objects that are hard to deal with, as, whenever designing one, one has to keep in mind that choices made in certain local places of the graph have high chances to interfere in distant places to be considered later on.

The following question is, thus, our guiding thread throughout this work:

Question 1.1. *Given a graph G, how hard is it to determine $s(G)$?*

Although we do not manage to give an ultimate answer to Question 1.1 in this work, we prove results[1] on several variants of the irregularity strength parameter, which might stand as hints that, perhaps, determining the irregularity strength of a given graph is NP-hard. Namely:

- In Section 2, we first consider a variant of the irregularity strength introduced by Przybyło in [13] where a distance $d \geq 1$ is fixed, and, by a labelling, it is required to have $\sigma(u) \neq \sigma(v)$ only for pairs of vertices u and v that are at distance at most d from each other. Note that this variant is very flexible, as it encapsulates the two most extreme cases, being all notions related to the so-called 1-2-3 Conjecture (when $d = 1$, see [9]) and the irregularity strength of graphs itself (when, in some sense, $d = \infty$). Some of the references we mentioned earlier, namely [1, 6], actually answer Question 1.1 for the case $d = 1$, showing that the problem is NP-hard. Through Theorem 2.3, we prove a similar result for every fixed $d \geq 2$.

- In Section 3, we then consider a variant of the irregularity strength introduced by Chartrand *et al.* in [4] in which all pairs of distinct vertices must be distinguished by a labelling, but through their multisets of incident labels. Note that distinguishing vertices through their multisets is easier than through their sums, as two multisets of labels are different whenever their sums are, but the converse is not always true. We answer Question 1.1 in this setting, showing, in Theorem 3.3, that determining this variant of the irregularity strength is NP-hard.

We believe our two main results complement each other, as, in the first variant of the irregularity strength we consider, the distinguishing parameter is the same while the distance requirement is weaker, while, in the second variant we consider, the distance requirement is the same while the distinguishing parameter is weaker. In particular, we have good hope that the reduction mechanisms we provide in our proofs could help in progressing towards answering Question 1.1. We discuss further on this point in concluding Section 4.

2. Distant irregularity strength

The terminology we use throughout this section is the following. Let G be a graph, and ℓ be a labelling of G. For any $d \geq 1$, we say that ℓ is *d-irregular* if $\sigma(u) \neq \sigma(v)$ for every two vertices u and v that are at distance at most d from each other in G. We denote by $s^d(G)$ the smallest k such that G admits d-irregular k-labellings.

As mentioned earlier, these notions were first considered by Przybyło in [13] as a distinguishing labelling notion generalising both the notions revolving around the 1-2-3 Conjecture (case where $d = 1$) and the irregularity strength itself (case where, in some sense, $d = \infty$). Note that, this time, if d is fixed, then it makes sense to consider the parameter s^d only in connected graphs.

Regarding complexity aspects, it was proved that determining $s^1(G)$ for a given graph G is NP-hard [1, 6]. A stronger result was actually proved in [1], namely that this complexity result remains true when G is regular (actually cubic); this is a very useful result of the field, as many close labelling problems tend to behave similarly in regular graphs, and this NP-hardness result consequently implies the NP-hardness of several other related problems (see e.g. [2]).

[1] More precisely, we prove NP-hardness results of the form "*Given a graph G, it is NP-hard to decide if G can be labelled in a certain way with labels 1 and 2 only*". This indeed implies that determining the corresponding chromatic parameter is NP-hard, as, for each of the labelling notions we consider in this work, deciding if label 1 suffices for a graph G can clearly be determined in polynomial time, by checking whether G has certain properties.

In this section, through upcoming Theorem 2.3, we actually prove that determining $s^d(G)$ for a given graph G is NP-hard for every $d \geq 2$. Before that, we need to introduce two results. In these two results and later on, we deal with labellings that verify some distinguishing properties (such as being d-irregular) except for some fixed vertices $u_1, \dots, u_k$, meaning that the u_i's might be involved in conflicts. Another way to phrase this is that every conflict (with respect to the desired distinguishing properties) must involve a u_i. This is what we mean when saying that, *omitting some vertices*, a labelling has certain properties.

Lemma 2.1. *Let G be a complete graph on $d+1 \geq 3$ vertices $w, v_1, \dots, v_d$. By every 2-labelling of G that is irregular when omitting w, the set $\{\sigma(v_1), \dots, \sigma(v_d)\}$ is either $\{d, \dots, 2d-1\}$ or $\{d+1, \dots, 2d\}$. Furthermore, for every $s \in \{d, 2d\}$, there exist irregular 2-labellings ℓ where:*

- *$s \notin \{\sigma(w), \sigma(v_1), \dots, \sigma(v_d)\}$; and*
- *$\sigma(w) = \frac{3d}{2}$ if d is even (regardless of s), or*
  - *$\sigma(w) = \frac{3d-1}{2}$ if d is odd and $s = 2d$, or*
  - *$\sigma(w) = \frac{3d+1}{2}$ if d is odd and $s = d$.*

Proof. Since the v_i's have degree d, then, by every irregular 2-labelling ℓ of G, we have $\sigma(v_i) \in \{d, \dots, 2d\}$ for every $i \in \{1, \dots, d\}$. Note also that the $\sigma(v_i)$'s must be pairwise distinct since ℓ is irregular, and that having $\sigma(v_i) = d$ for some i implies, due to the edge v_iv_j being assigned label 1, that we cannot have $\sigma(v_j) = 2d$ for any $j \neq i$. Similarly, having $\sigma(v_i) = 2d$ for some i implies we cannot have $\sigma(v_j) = d$ for any $j \neq i$. This implies the first part of the statement.

Let us now focus on proving the rest of the statement. We start by proving the claim for $s = 2d$. We prove it by induction on d. For the base case where $d = 2$, we have three vertices w, v_1, v_2. Consider ℓ, the 2-labelling of G where $\ell(v_1v_2) = 1$, $\ell(wv_1) = 1$, and $\ell(wv_2) = 2$. Then we have $\sigma(v_1) = 2$, $\sigma(v_2) = 3$, and $\sigma(w) = 3$. Thus, omitting w, the vertices are distinguished. Also, $2d = 4 \notin \{2, 3\} = \{\sigma(w), \sigma(v_1), \sigma(v_2)\}$, and we have $\sigma(w) = \frac{3d}{2} = 3$.

Consider now a general value as d, and assume the claim holds for smaller values than d. Consider ℓ, a 2-labelling of $G' = G[v_1, \dots, v_d]$ verifying all conditions of the claim, which exists by induction. In particular, there is an integer α such that, in G', all $\sigma(v_i)$'s but $\sigma(v_\alpha)$ are pairwise distinct, the value $2(d-1)$ does not appear in $\{\sigma(v_1), \dots, \sigma(v_d)\}$, and $\sigma(v_\alpha) = \frac{3(d-1)}{2}$ (if $d-1$ is even) or $\sigma(v_\alpha) = \frac{3(d-1)-1}{2}$ (otherwise). Assume $\sigma(v_1) < \dots < \sigma(v_\alpha) = \sigma(v_{\alpha+1}) < \dots < \sigma(v_d)$. We extend ℓ to the edges incident to w, thus from G' to G, by setting $\ell(wv_1) = \dots = \ell(wv_\alpha) = 1$ and $\ell(wv_{\alpha+1}) = \dots = \ell(wv_d) = 2$. As a result, note that, in G, we now have $\sigma(v_1) < \dots < \sigma(v_d)$. Also, no v_i has all its incident edges assigned label 2 by ℓ, and neither does w, so no vertex has sum $2d$. Actually, regarding w, due to the value of α, it is either incident to $\frac{d}{2}$ edges assigned label 2 (if d is even) or $\frac{d-1}{2}$ edges assigned label 2 (otherwise). Thus, $\sigma(w) = \frac{3d}{2}$ or $\sigma(w) = \frac{3d-1}{2}$, depending on the parity of d. So ℓ fulfills all desired properties, and the induction is proved.

To see that the claim also holds for $s = d$, just consider the same arguments as above, but, for the base case, consider ℓ, the 2-labelling of G where $\ell(v_1v_2) = 1$, $\ell(wv_1) = 2$, and $\ell(wv_2) = 2$. $\square$

We need some special graph modifications for the next preliminary result. Given a graph G with a vertex v, by *attaching a k-clique at v* we mean adding a clique on k vertices to G, and joining v and each of the k vertices of that clique. Remark that, as a result, all vertices of the clique get degree exactly k. Now, for any $k \geq 7$, assuming v has degree 2, by *attaching a k-fan at v* we mean modifying G in the following way (see Figure 1 for an illustration):

- we add $k-2$ new vertices $u_1, \dots, u_{k-2}$, which we join to v through the edges $vu_1, \dots, vu_{k-2}$;
- we attach both a k-clique Q_1 and a $(2k+1)$-clique Q_2 at u_1, and set $n_1 = 3k+2 = d(u_1)$;
- we attach a $(2n_1+1)$-clique Q_3 and a $(2(2n_1+1)+1)$-clique Q_4 at u_2 and set $n_2 = d(u_2)$; then we attach a $(2n_2+1)$-clique Q_5 and a $(2(2n_2+1)+1)$-clique Q_6 at u_3 and set $n_3 = d(u_3)$; then we attach a $(2n_3+1)$-clique Q_7 and a $(2(2n_3+1)+1)$-clique Q_8 at u_4 and set $n_4 = d(u_4)$; and so on. That is, we treat the u_i's one by one in order. For any $i \geq 5$ treated that way, if we define n_{i-1} as the degree of u_{i-1}, then we attach, at u_i, a $(2n_{i-1}+1)$-clique Q_{2i-1} and a $(2(2n_{i-1}+1)+1)$-clique Q_{2i} at u_i. Note that the order of any Q_i is a function of k only.

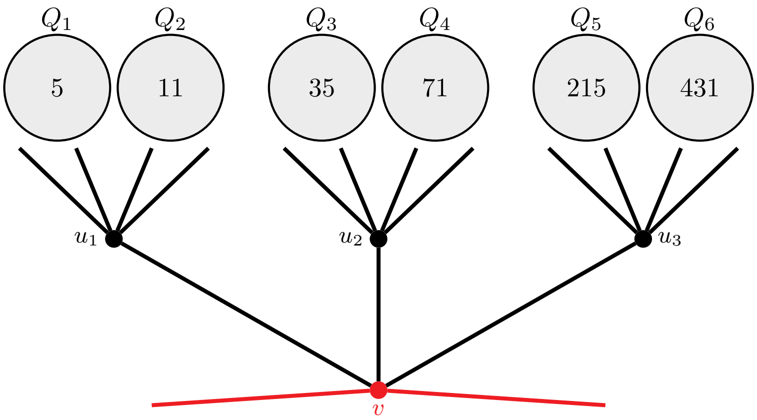

Figure 1: Attaching a k-fan (in black) at a degree-2 vertex v (in red). Circles represent cliques on the indicated numbers of vertices. Here, $k = 5$ for simplicity. Q_1 is a k-clique, thus a 5-clique, while Q_2 is a $(2k+1)$-clique, thus an 11-clique. Setting $n_1 = d(u_1) = 17$, Q_3 is a $(2n_1+1)$-clique, thus a 35-clique, while Q_4 is a $(2(2n_1+1)+1)$-clique, thus a 71-clique. Setting $n_2 = d(u_2) = 107$, Q_5 is a $(2n_2+1)$-clique, thus a 215-clique, while Q_6 is a $(2(2n_2+1)+1)$-clique, thus a 431-clique. Note that, after the attachment, v is of degree 5.

Note also that, through attaching a k-fan at v, the degree of v becomes exactly k.

With respect to d-irregular 2-labellings, attaching fans has the following properties:

Lemma 2.2. *Let G be a graph obtained from the path of uvw length 2 by attaching a k-fan at v for some $k \geq 7$. Omitting u and w, if ℓ is a d-irregular 2-labelling of G for some $d \geq 2$, then $\sigma(v) \in \{k, 2k\}$, i.e., all edges incident to v must be assigned the same label. Furthermore, omitting u and w, for every $s \in \{k, 2k\}$, there exist d-irregular 2-labellings of G for which $\sigma(v) = s$.*

Proof. Note that for a vertex x of degree k, by any 2-labelling of any graph we have $\sigma(x) \in \{k, \dots, 2k\}$. Thus, the first part of the claim follows from the fact that, in G, the k vertices of Q_1 (following the terminology above) are adjacent with degree k, and they are at distance at most d from v, since $d \geq 2$, while v also has degree k. This implies that if we denote by $v_1, \dots, v_k$ the vertices of Q_1, then we must have $\{\sigma(v), \sigma(v_1), \dots, \sigma(v_k)\} = \{k, \dots, 2k\}$. Particularly, one vertex v_i of Q_1 must verify $\sigma(v_i) \in \{k, 2k\}$, and, for this to happen, all edges incident to v_i must be assigned the same label by ℓ. If this label is 1, then this implies that we have $\sigma(v_j) \neq 2k$ for every $j \neq i$ since $\ell(v_iv_j) = 1$. Similarly, if this label is 2, then we must have $\sigma(v_j) \neq k$ for every $j \neq i$ since $\ell(v_iv_j) = 2$. Thus, in the former case we must have $\sigma(v) = 2k$, while in the latter case we must have $\sigma(v) = k$. All edges incident to v must thus be assigned the same label.

We now prove the last part of the statement. Consider the 2-labelling ℓ of G obtained as follows.

- First, setting $Q_1 = \{v_1, \dots, v_k\}$, label the edges of $G[Q_1 \cup \{u_1\}]$ as described in Lemma 2.1, that is, so that $\sigma(v_1) = k, \sigma(v_2) = k+1, \dots, \sigma(v_k) = 2k-1$, and either $\sigma(u_1) = \frac{3k}{2}$ (if k is even) or $\sigma(u_1) = \frac{3k-1}{2}$ (otherwise).
- Similarly, setting $Q_2 = \{w_1, \dots, w_{2k+1}\}$, label the edges of $G[Q_2 \cup \{u_1\}]$ next, as described in Lemma 2.1, so that $\sigma(w_1) = 2k+1, \sigma(w_2) = 2k+2, \dots, \sigma(w_{2k+1}) = 4k+1$, and $\sigma(u_1) = \frac{3(2k+1)-1}{2}$ (which is mandatory, since $2k+1$ is odd).
- For every $i \in \{3, \dots, 2k-4\}$, do the same for $G[Q_i \cup \{u_j\}]$, where u_j is the vertex to which Q_i was attached. That is, setting $|Q_i| = r$ and $Q_i = \{w_1, \dots, w_r\}$, then, as described in Lemma 2.1, label the edges of $G[Q_i \cup \{u_j\}]$ so that $\sigma(w_1) = r, \dots, \sigma(w_r) = 2r-1$, and, since r is odd by construction, $\sigma(u_j) = \frac{3r-1}{2}$.
- Lastly, assign label 2 to all edges incident to v, *i.e.*, to uv, vw, and the vu_i's.

We claim that, omitting u and w, all vertices of G get distinguished by ℓ (not only those at distance at most d from each other). By how ℓ was obtained, recall first that two vertices from a single Q_i cannot be in conflict. The important property, now, is that any two different Q_i's cannot have their vertices being in conflict, due to their degrees being too different. Indeed, if Q_i and Q_j are two cliques, then, assuming $i < j$, note that if the vertices of Q_i have degree r, and thus sum at most $2r$, then, by construction, the vertices of Q_j have degree at least $2r+1$, and thus sum at least $2r+1$. Similarly, note that, by how ℓ was obtained, we have $\sigma(v) = 2k$, while the vertices from Q_1 have sum at most $2k-1$ and all other vertices have sum at least $2k+1$. Also, note that, for any i such that u_{i+1} exists, we have $d(u_{i+1}) > 2d(u_i)$ by construction, which implies that the u_i's cannot be in conflict. Lastly, some u_j cannot be in conflict with the vertices of a Q_i attached at u_j by construction of ℓ, and due to the degrees of the vertices in the other Q_i's, vertex u_j cannot be in conflict with any of these vertices neither. In particular, note that, regardless of the parity of k, if some u_j is incident to two attached cliques having r vertices of degree r and $2r+1$ vertices of degree $2r+1$, respectively, then $d(u_j) = 3r+2$ (recall that vu_j is an edge), and we have $\sigma(u_j) \geq \frac{3r-1}{2} + \frac{6r+2}{2} + 2 = \frac{9r+5}{2}$, which is strictly greater than $4r+1$, the maximum sum of a vertex adjacent to u_j, since $r \geq k \geq 7$. Thus, omitting u and w, we have that ℓ is d-irregular.

It can be checked, that, by a similar construction as above (but labelling the edges of $G[Q_1 \cup \{u_1\}]$ and $G[Q_2 \cup \{u_1\}]$ in the second way described in Lemma 2.1, so that no vertex has sum k), we can also obtain a d-irregular 2-labelling of G where all edges incident to v are assigned label 1. □

We are now ready to prove our main result in this section.

Theorem 2.3. *For every $d \geq 2$, deciding if $s^d(G) \leq 2$ for a given graph G is* NP*-complete.*

Proof. For any fixed value of $d \geq 2$ the problem is clearly in NP, so let us focus on proving it is NP-hard. This is done by reduction from the problem of deciding whether $s^1(H) \leq 2$ for a given cubic graph H, which was shown to be NP-hard in [1]. That is, given a cubic graph H, we construct, in polynomial time, a graph G such that $s^1(H) \leq 2$ if and only if $s^d(G) \leq 2$.

So that we can describe the construction of G from H, we need some colouring of the edges of H. Since H is cubic, then, by Vizing's Theorem [14], it admits a proper 4-edge-colouring ϕ, *i.e.*, an assignment of colours from $\{1,2,3,4\}$ to the edges such that no two adjacent edges are assigned the same colour. Furthermore, ϕ can be obtained in polynomial time according to [10].

We now construct G from H as follows (see Figure 2). We start from H, and subdivide every edge uw of H exactly $d-1$ times, resulting in $d-1$ new degree-2 vertices forming a set which we denote by $S(uw)$. Note that once all edges of H have been subdivided this way, then, for every two adjacent vertices u and w of H, in the current graph u and w are now at distance exactly d. We now modify the current graph further as follows.

- We first consider all edges of H being assigned colour 1 by ϕ. For each such edge uw with $\phi(uw) = 1$, let us set $S(uw) = \{v_1, \dots, v_{d-1}\}$. We start by attaching a 7-fan at v_1. Denoting by α_1 the maximum degree of a vertex in that 7-fan, we then attach a $(2\alpha_1+1)$-fan at v_2. Denoting by α_2 the maximum degree of a vertex in that fan attached at v_2, we then attach a $(2\alpha_2+1)$-fan at v_3. We go on like this for every $i \in \{4, \dots, d-1\}$ in turn, that is, denoting by α_{i-1} the maximum degree of a vertex in the fan attached at v_{i-1}, we then attach a $(2\alpha_{i-1}+1)$-fan at v_i. We denote by α_{d-1} the maximum degree of a vertex in the last fan, attached at v_{d-1}.

- We then consider all edges of H assigned colour 2 by ϕ. For every such edge uw with $\phi(uw) = 2$, set $S(uw) = \{v_1, \dots, v_{d-1}\}$. We start by attaching a $(2\alpha_{d-1}+1)$-fan at v_1. Then, for every $i \in \{2, \dots, d-1\}$ in turn, denoting by β_{i-1} the maximum degree of a vertex in the fan attached at v_{i-1}, we attach a $(2\beta_{i-1}+1)$-fan at v_i. Eventually, we denote by β_{d-1} the maximum degree of a vertex in the fan attached at v_{d-1}.

- We next consider all edges of H assigned colour 3 by ϕ. For every such edge uw with $\phi(uw) = 3$, we set $S(uw) = \{v_1, \dots, v_{d-1}\}$. We first attach a $(2\beta_{d-1}+1)$-fan at v_1. Then, for every $i \in \{2, \dots, d-1\}$ in turn, defining γ_{i-1} as the maximum degree of a vertex in the fan attached at v_{i-1}, we attach a $(2\gamma_{i-1}+1)$-fan at v_i. Once v_{d-1} has been treated, we denote by γ_{d-1} the maximum degree of a vertex in the fan attached at v_{d-1}.

- Lastly, we consider all edges of H assigned colour 4 by ϕ. For every edge uw with $\phi(uw) = 4$, we set $S(uw) = \{v_1, \dots, v_{d-1}\}$. We begin by attaching a $(2\gamma_{d-1}+1)$-fan at v_1. Then, for

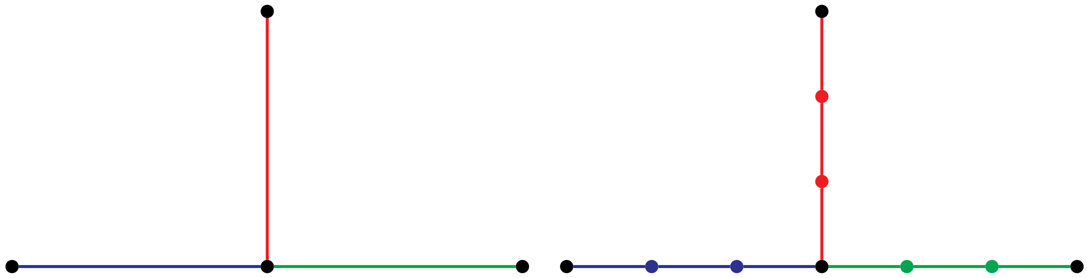

(a) Initial graph (b) Subdividing edges

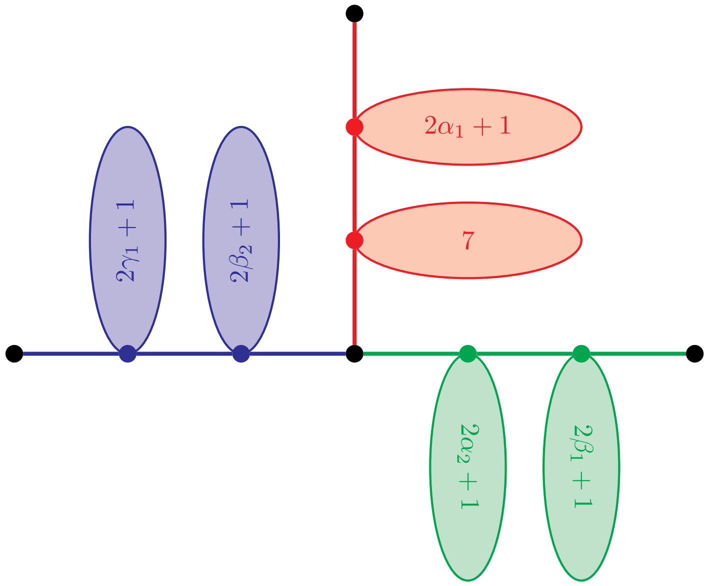

(c) Attaching fans

Figure 2: Illustration of the reduction in the proof of Theorem 2.3, for $d = 3$. (a) represents the neighbourhood of a vertex in H, its three incident edges being assigned colour 1 (red), 2 (green), and 3 (blue) by ϕ. In the first step of the reduction (b), all edges are subdivided $d - 1 = 2$ times. Then, in the second step (c), the subdivided vertices resulting from the subdivision of a red edge are being attached a 7-fan and a $(2\alpha_1 + 1)$-fan, respectively, the subdivided vertices resulting from the subdivision of a green edge are being attached a $(2\alpha_2 + 1)$-fan and a $(2\beta_1 + 1)$-fan, respectively, and the subdivided vertices resulting from the subdivision of a blue edge are being attached a $(2\beta_2 + 1)$-fan and a $(2\gamma_1 + 1)$-fan, respectively. Ellipses represent k-fans on the indicated values of k.

every $i \in \{2, \dots, d-1\}$ in turn, assuming δ_{i-1} denotes the maximum degree of a vertex in the fan attached at v_{i-1}, we attach a $(2\delta_{i-1} + 1)$-fan at v_i.

The resulting graph is G. Note that G is essentially obtained from H by subdividing all edges exactly $d - 1$ times, and attaching a particular fan (following ϕ) at each vertex resulting from a subdivision, which we call a *subdivision vertex*. So, the number of subdivision vertices in G is exactly $(d-1)|E(H)|$. Now, notice that the fans we have attached at the subdivision vertices grew exponentially. Particularly, there are exactly $4(d-1)$ distinct types of fans, that is, k-fans for all k in

$$(7, 2\alpha_1 + 1, \dots, 2\alpha_{d-1} + 1, 2\beta_1 + 1, \dots, 2\beta_{d-1} + 1, 2\gamma_1 + 1, \dots, 2\gamma_{d-1} + 1, 2\delta_1 + 1, \dots, 2\delta_{d-2} + 1)\,.$$

For every two consecutive values x and y of this ordered set, note that the number of vertices in a y-fan is a function only of y and of the number of vertices of an x-fan. Since the initial value, 7, and d are constant, and there are $4(d-1)$ types of fans, the maximum number of vertices in a fan is bounded above by a function of 7 and d only, and is thus constant. Since the number of fans in G is exactly $(d-1)|E(H)|$, we deduce that G is obtained in polynomial time from H.

We now prove that we have the desired equivalence between H and G.

- Assume first that G admits a d-irregular 2-labelling ℓ. The key property is the following. Consider an edge uw of H, and the corresponding set $S(uw) = \{v_1, \dots, v_{d-1}\}$ of subdivision vertices of G. Assume $P = (u, v_1, \dots, v_{d-1}, w)$ is the d-path[2] joining u and w in G. By Lemma 2.2, note that the fan attached at v_1 implies that $\ell(uv_1) = \ell(v_1v_2)$. Similarly, the fan attached at v_2 implies that $\ell(v_1v_2) = \ell(v_2v_3)$. More generally, for every $i \in \{1, \dots, d-1\}$, the fan attached at v_i implies that the two edges of P incident to v_i must be assigned the same label by ℓ. Thus, all edges of P must be assigned the same label by ℓ.

  Now, consider the 2-labelling ℓ' of H obtained by considering every edge uw of H, and, as above, denoting by $(u, v_1, \dots, v_{d-1}, w)$ the corresponding d-path in G, setting $\ell'(uw) = \ell(uv_1)$. By the property we have pointed out above, it can be noted that, for every vertex v of H, the value of $\sigma(v)$ by ℓ' in H is the same as the value of $\sigma(v)$ by ℓ in G. Now, since ℓ is d-irregular and, for every edge uw of H, the vertices u and w are at distance exactly d in G, we have $\sigma(u) \neq \sigma(w)$ by ℓ. Then, we also have $\sigma(u) \neq \sigma(w)$ by ℓ'. Since this holds for every edge uw of H, we deduce that ℓ' is 1-irregular.

- Conversely, assume H admits a 1-irregular 2-labelling ℓ. We consider ℓ', the 2-labelling of G obtained as follows. For every edge uw of H, consider $P = (u, v_1, \dots, v_{d-1}, w)$, the corresponding d-path in G. To every edge of P, we assign label $\ell(uw)$ by ℓ'. Now, for every $i \in \{1, \dots, d-1\}$, we extend ℓ' to the edges of the fan attached at v_i in a d-irregular way and so that all edges incident to v_i are assigned label $\ell(uw)$, which is possible by Lemma 2.2. Note that, once all edges uw of H have been considered, ℓ' labels all edges of G.

  We claim that ℓ' is d-irregular. First off, as previously, note that for every vertex v of H, the value of $\sigma(v)$ by ℓ is the same as the value of $\sigma(v)$ by ℓ'. Note now that G has two main types of vertices: those originating from H, and the other ones we have added when constructing G, which are each part of an attached fan (assuming a subdivision vertex to which a fan was attached is part of that fan). Thus, we split the analysis into the following two cases:

  - If v is a vertex of $V(H) \cap V(G)$, then, in G, note that v has degree 3 by construction (since H is cubic). Note that every vertex that is part of a fan has degree at least 7, and thus it cannot be in conflict with v by ℓ'. So v can only be in conflict with another (degree-3) vertex u of $V(H) \cap V(G)$. By construction, v and u are at distance at most d in G if and only if v and u are adjacent in H. Due to how ℓ' was obtained, and because ℓ is 1-irregular and v and u have the same sum by ℓ and ℓ', we have $\sigma(v) \neq \sigma(u)$ by ℓ'.
  - Assume now that v is part of a k-fan. Due to how ℓ' was constructed, v cannot be in conflict with a vertex from the same fan. Also, as mentioned earlier, v cannot be in conflict with a vertex u of $V(H) \cap V(G)$, since u has degree 3 while v has degree at least 7. So v can only be in conflict with a vertex u from another k'-fan. By construction, note that, when $k \neq k'$, the degrees of v and u are so different that they cannot be in conflict. So we must have $k = k'$, in which case, still by construction, the two fans were attached on subdivision vertices resulting from the subdivision of two distinct edges assigned the same colour by ϕ. Since no two adjacent edges of H are assigned the same colour by ϕ, by construction it can be noted that, in G, actually v and u are at distance strictly more than d. Thus, having $\sigma(v) \neq \sigma(u)$ is not required for ℓ' to be d-irregular.

  All these arguments imply that ℓ' is indeed d-irregular, as claimed.

Thus, the equivalence between H and G holds. □

3. Multiset irregularity strength

Throughout this section, we deal with the following notions. Let G be a graph, and ℓ be a k-labelling of G. For every vertex v of G, we denote by $\mu(v)$ the *colour code* of v, being the *multiset* of labels assigned to the edges incident to v. For convenience, we will represent colour codes in a

[2] Throughout this work, a *d-path* refers to a path of length d.

compact way, namely through the notation $\mu(v) = (1^{n_1}, 2^{n_2}, \dots, k^{n_k})$, where every ith element i^{n_i} of the code tells that v is incident to n_i edges assigned label i by ℓ. Note that $d(v) = n_1 + \dots + n_k$. Now, we say that ℓ is *m-irregular* if we have $\mu(u) \neq \mu(v)$ for every two distinct vertices u and v of G, and we define $s_m(G)$ as the smallest k such that m-irregular k-labellings of G exist.

Note that an irregular labelling is always m-irregular (while the converse does not have to be true), so, as mentioned earlier, we always have $s_m(G) \leq s(G)$ for a graph G. Chartrand *et al.* first introduced m-irregular labellings in [4], under the name "detectable colourings".

In this section, we prove that determining $s_m(G)$ for a given graph G is NP-hard. To that end, we will make use of some of the preliminary results we introduced in previous Section 2, and of a few more easy results which we prove now. In these ones, for any $d \geq 1$ and $k \geq 1$, we denote by $\mathcal{U}(k, d)$ the set containing all colour codes $(1^{n_1}, \dots, k^{n_k})$ that can be obtained for a degree-d vertex through a k-labelling of some graph, *i.e.*, $n_1 + \dots + n_k = d$. Note that $|\mathcal{U}(2, d)| = d + 1$, as $\mathcal{U}(2, d) = \{(1^0, 2^d), (1^1, 2^{d-1}), \dots, (1^d, 2^0)\}$.

Lemma 3.1. *Let G be a complete graph on $d + 1 \geq 3$ vertices $w, v_1, \dots, v_d$, and ℓ be a 2-labelling of G. Omitting w, if ℓ is m-irregular, then $\{\mu(v_1), \dots, \mu(v_d)\}$ is either $\mathcal{U}(2, d) \setminus (1^0, 2^d)$ or $\mathcal{U}(2, d) \setminus (1^d, 2^0)$. Furthermore, omitting w, for every $S \in \{\mathcal{U}(2, d) \setminus (1^0, 2^d), \mathcal{U}(2, d) \setminus (1^d, 2^0)\}$, there exist m-irregular 2-labellings of G for which $S = \{\mu(v_1), \dots, \mu(v_d)\}$.*

Proof. This follows directly from Lemma 2.1, since an irregular labelling is m-irregular. □

Lemma 3.2. *Let G be a graph obtained from a complete graph on $p \geq 3$ vertices $v_1, \dots, v_p$ by attaching a degree-1 vertex u at v_p. Omitting u, for every $s \in \{1, 2\}$, there exist m-irregular 2-labellings of G where uv_p is assigned label s.*

Proof. By Lemma 3.1, there are 2-labellings of $G[\{v_1, \dots, v_p\}]$ where $\{\mu(v_1), \dots, \mu(v_{p-1})\} = \mathcal{U}(2, p-1) \setminus (1^0, 2^{p-1})$. Consider ℓ, such a 2-labelling of $G[\{v_1, \dots, v_p\}]$, and extend ℓ to G by assigning any label $s \in \{1, 2\}$ to $v_p u$. Note that, omitting u, this results in an m-irregular 2-labelling of G, since $d(v_p) > d(v_{p-1}), \dots, d(v_1)$ which implies that $\mu(v_p) \neq \mu(v_{p-1}), \dots, \mu(v_1)$, while the other $p - 1$ v_i's cannot be involved in multiset conflicts due to the main property of ℓ. □

We are now ready to prove our main result in this section.

Theorem 3.3. *Deciding if $s_m(G) = 2$ for a given graph G is NP-complete.*

Proof. The problem is clearly in NP, so we focus on proving its NP-hardness. This is done by reduction from Monotone Cubic 1-in-3 SAT, which is known to be NP-hard [11]. In that problem, a 3CNF formula F, in which all clauses have three distinct non-negated variables and every variable appears in exactly three distinct clauses, is given, and the question is whether there is a truth assignment to the variables such that F is satisfied in a 1-in-3 way, *i.e.*, so that every clause contains exactly one variable set to true. From an instance F of Monotone Cubic 1-in-3 SAT, we construct, in polynomial time, a graph G such that F can be satisfied in a 1-in-3 way if and only if G admits m-irregular 2-labellings. Let us denote by n the number of clauses and variables of F (note that these two numbers are indeed equal).

The construction of G starts as follows:

- We start from G_F, the cubic bipartite graph modelling the structure of F. That is, G_F has a *clause vertex* v_C for every clause C of F, a *variable vertex* v_x for every variable x of F, and a *formula edge* $v_C v_x$ whenever a variable x is contained in a clause C of F.
- We add a special vertex w to the graph, isolated for now.

In the rest of the construction below, we will need to increase the degrees of some vertices of G_F by a particular amount, through incident edges which, eventually, will have to be assigned a particular label by any m-irregular 2-labelling of G. As will be made clear below, such edges can be generated at will, mainly through attaching certain cliques at w. The only issue we need to be careful with is that, when eventually designing m-irregular 2-labellings, the vertices of these cliques might interfere with other vertices (having the same degree) of the whole graph (since no two vertices are allowed to have the same colour code). To make sure to avoid such problems, we will mainly make sure to consider vertices with sufficiently different palettes of degrees. To better track the vertex degrees throughout the proof, we will maintain some sets $\mathcal{C}$ (clause vertices), $\mathcal{V}$ (variable vertices), $\mathcal{T}$ (trail vertices), and $\mathcal{U}$[3] (clean-up vertices) containing the degree values of certain types of vertices. Initially, these sets are empty.

We start by defining the sets $\mathcal{C}$ and $\mathcal{V}$, which we do as follows:

- Assuming the clauses of F are $C_1, \dots, C_n$, we will, later on, add edges incident to every v_{C_i} so that its degree, which is currently precisely 3, is eventually $e(v_{C_i}) = i + 2$. Thus, the degrees of $v_{C_1}, \dots, v_{C_n}$ will eventually be $3, \dots, n + 2$, respectively. We add these n values $e(v_{C_1}), \dots, e(v_{C_n})$ to the set $\mathcal{C}$. We define C as the maximum value in $\mathcal{C}$, which is $n + 2$.
- Assuming the variables of F are $x_1, \dots, x_n$, through adding edges incident to the v_{x_i}'s, we will make sure that each v_{x_i} eventually has degree $e(v_{x_i}) = C+i$. In other words, $v_{x_1}, \dots, v_{x_n}$ will eventually have degree $C+1, \dots, C+n$, respectively. We add these n values $e(v_{x_1}), \dots, e(v_{x_n})$ to $\mathcal{V}$, and we define V as the maximum value in $\mathcal{V}$, which is $C + n$.

So that each v_C and each v_x eventually has degree $e(v_C)$ and $e(v_x)$, respectively, we will need to generate $e(v_C) - 3$ and $e(v_x) - 3$ edges, respectively, and attach them at v_C and v_x, respectively. Recall indeed that each of these vertices is currently incident to three formula edges of G_F. This will be done by generating such edges for which the label by any m-irregular 2-labelling of G is known. This way, the colour code of v_C and v_x will, eventually, be mostly known by any m-irregular 2-labelling, as only the formula edges incident to these vertices will fluctuate. More precisely, since v_C and v_x are incident to exactly three incident edges, their eventual colour codes will each be one of four possible values only (depending on whether all, two, one, or no incident formula edges are assigned label 1, while the others are assigned label 2). To have the desired equivalence between G and F, we will need the colour code of every v_C to be precisely one of the four possible values, while we will need the colour code of every v_x to be one of two of the four possible values. To force that, we will also generate three and two vertices of degree $e(v_C)$ and $e(v_x)$, respectively, and force their incident edges to be assigned some labels by any m-irregular 2-labelling, so that the colour codes of these vertices are precisely those we want to forbid for v_C and v_x.

In what follows, we will thus require some particular edges that we will "attach" at other vertices. Formally, assuming xy is a pending edge of some graph, *i.e.*, $d(x) > 1$ and $d(y) = 1$, then, by *attaching xy at some vertex $z \notin \{x, y\}$*, we mean identifying y and z. Note that the attachment operation involves an edge with exactly one degree-1 vertex; this means a pending edge can be attached only once. We need to attach edges which, eventually, we know which label they will be assigned by an m-irregular 2-labelling; to get such edges, we add some more structure to the graph (see Figure 3 for an illustration of that part of the reduction).

- We grow a sufficiently long (see later) 1-*forcing trail* in the following way. We start by adding a *trail vertex* v to the graph which we join to $V + 1$ new degree-1 vertices $u_1, \dots, u_{V+1}$ (so that v has degree $V + 1$), then attach a $(V + 1)$-clique at w, and finally add the value $V + 1$ to $\mathcal{T}$. Now, to make the 1-forcing trail one step longer, we proceed as follows. Assume u is the last trail vertex we have added to the 1-forcing trail, and that this vertex u is of degree d. Then, to make the 1-forcing trail longer, we add a new trail vertex v joined to each of $d + 1$ new degree-1 vertices (so that v has degree $d + 1$), attach a $(d + 1)$-clique at w, add the value $d + 1$ to $\mathcal{T}$, and, lastly, identify one pending edge ux and one pending edge yv (that is, we identify u and y, and, similarly, x and v, and keep only one edge joining u and v). In other words, the 1-forcing trail is obtained from a main path by attaching degree-1 vertices to every of its inner vertices (being the trail vertices) so that certain consecutive degrees are attained. We denote by T_1 the maximum value added to $\mathcal{T}$.

  Additionally, once the 1-forcing trail is long enough (see below), we eventually identify any two of its degree-1 vertices with distinct neighbours, resulting in a degree-2 vertex t_1.
- We also grow a sufficiently long (disjoint) 2-*forcing trail* similarly as the 1-forcing trail, the difference being that the first trail vertex is of degree $T_1 + 1$. We denote by T_2 the maximum

[3] Not to be mistaken with the function $\mathcal{U}(k, d)$ from the beginning of Section 3.

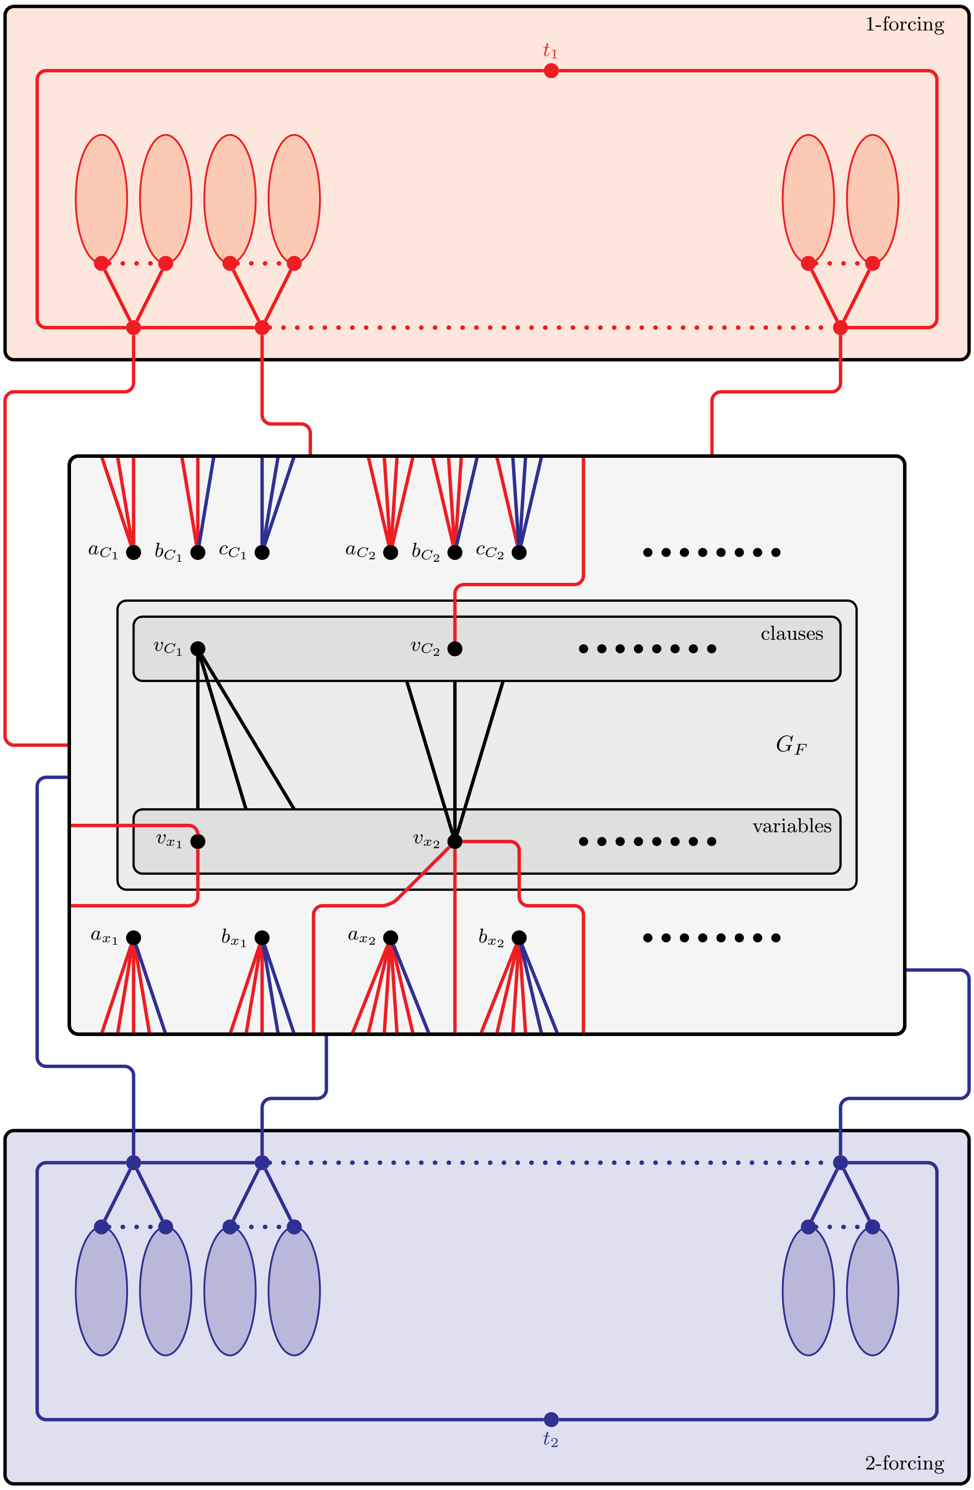

Figure 3: Part of the reduction in the proof of Theorem 3.3. The connected component containing w and all cliques attached to it are not represented. Red edges are 1-edges by any m-irregular 2-labelling, while blue edges are 2-edges. Ellipses represent cliques. Note that in each of the sets $\{v_{C_1}, a_{C_1}, b_{C_1}, c_{C_1}\}$, $\{v_{C_2}, a_{C_2}, b_{C_2}, c_{C_2}\}$, $\{v_{x_1}, a_{x_1}, b_{x_1}\}$, and $\{v_{x_2}, a_{x_2}, b_{x_2}\}$, all vertices have the same degree due to incident 1-edges and 2-edges, and that any two of these sets contain vertices with distinct degrees. For every $i \in \{1, 2\}$, due to a_{C_i}, b_{C_i}, and c_{C_i}, note that, in an m-irregular 2-labelling, it must be that one formula edge incident to v_{C_i} is assigned label 1 while the other two are assigned label 2. Likewise, for every $i \in \{1, 2\}$, due to a_{x_i} and b_{x_i}, all formula edges incident to v_{x_i} must be assigned the same label.

value added to $\mathcal{T}$ in the process, once the trail is long enough. We also identify two degree-1 vertices of the 2-forcing trail, and denote by t_2 the resulting degree-2 vertex.

The *length* of a forcing trail is its number of trail vertices. For every $i \in \{1,2\}$, a pending edge of the i-forcing trail is called an *i-edge*. In what follows, we assume the two forcing trails are long enough so that they each contain $R = 8nV$ trail vertices. Particularly, this implies that the number of i-edges is at least R for every $i \in \{1,2\}$, and even that there are at least R i-edges that are pairwise disjoint (*i.e.*, which originate from different trail vertices) for every $i \in \{1,2\}$.

We can now continue the construction of G:

- We start by considering every clause vertex v_C first, and attach at v_C exactly $e(v_C) - 3$ 1-edges so that the degree of v_C becomes exactly $e(v_C)$. We then add a new vertex a_C to the graph, and attach exactly $e(v_C)$ 1-edges at a_C. Similarly, we then add a new vertex b_C to the graph, and attach exactly $e(v_C) - 1$ 1-edges and one 2-edge at b_C. Lastly, we also add a new vertex c_C to the graph, and attach exactly $e(v_C) - 3$ 1-edges and three 2-edges at c_C. Note that all of a_C, b_C, and c_C have degree precisely $e(v_C)$.
- We then consider every clause vertex v_x, and attach at v_x $e(v_x) - 3$ 1-edges so that the degree of v_x becomes $e(v_x)$. We then add a new vertex a_x to the graph, at which we attach $e(v_x) - 1$ 1-edges and one 2-edge. Lastly, we add a new vertex b_x, at which we attach $e(v_x) - 2$ 1-edges and two 2-edges. Note, here, that a_x and b_x have degree $e(v_x)$.

Since there are a total of $2n$ clause vertices and variable vertices, we have added at most three new vertices for each clause vertex and variable vertex, and the maximum value of some $e(v_C)$ or $e(v_x)$ is V, note that, due to the value of R, we indeed have sufficiently many 1-edges and 2-edges in hands to go through the whole process above. Also, since the two forcing trails contain R trail vertices each, note that, whenever attaching several 1-edges or 2-edges to a single vertex, we can always make sure to not create parallel edges; in other words, the whole process above can be achieved in such a way that G remains simple.

We finish off the construction of G with some final clean-up:

- Some 1-edges and 2-edges were, perhaps, not attached to any vertex of the graph, and are thus pending. We consider every such edge xy (where $d(x) > 1$ and $d(y) = 1$) in turn, choose the largest value $d \in \{T_2\} \cup \mathcal{U}$, attach a $(d+1)$-clique at y, and add $d+1$ and $d+2$ to $\mathcal{U}$.
- After this, *i.e.*, once no pending 1-edges and 2-edges remain, assuming U is the maximum value in $\mathcal{U}$, we attach a $(U+1)$-clique at w, and add $U+1$ to $\mathcal{U}$.

We denote by G the resulting graph. We claim that G is obtained in polynomial time from F, due to the following arguments. Note first that the number of vertices of G_F is $2n$, and, since we have added three vertices a_C, b_C, and c_C for every clause vertex v_C, and two vertices a_x and b_x for every variable vertex v_x, the total number of all these vertices is $\mathcal{O}(n)$. Next, every i-forcing trail contains R trail vertices; thus, in total, there are $\mathcal{O}(n^2)$ trail vertices. Since the first trail vertex we have added has degree $V+1$ and all trail vertices we have added after that have consecutive degrees, the maximum degree of a trail vertex is $\mathcal{O}(n^2)$. For every trail vertex of degree d, we have also attached a d-clique at w. These cliques attached at w thus have $\mathcal{O}(n^4)$ vertices, since there are $\mathcal{O}(n^2)$ trail vertices, all of which have degree $\mathcal{O}(n^2)$. Finally, for every trail vertex of degree d, at most d of its incident pending edges were not attached to a_C's, b_C's, c_C's, a_x's, and b_x's, and, at the end of each such pending edge, we have attached a k-clique for some $k \in \mathcal{U}$, resulting in a vertex of degree $k+1$. Recall that d is $\mathcal{O}(n^2)$. Regarding k, the smallest possible value is T_2+1, which is $\mathcal{O}(n^2)$, and we have added consecutive pairs of consecutive values or single values to $\mathcal{U}$ only. Since there are $\mathcal{O}(n^2)$ trail vertices all of which have degree $\mathcal{O}(n^2)$, in total we have attached $\mathcal{O}(n^4)$ cliques to pending edges, and thus added $\mathcal{O}(n^4)$ values to $\mathcal{U}$. So the maximum value in $\mathcal{U}$ is $\mathcal{O}(n^4)$. Thus, in total, all these cliques attached to pending edges have $\mathcal{O}(n^8)$ vertices. The very last clique we have attached at w is a $(U+1)$-clique, which is thus of order $\mathcal{O}(n^4)$. So the order of G is $\mathcal{O}(n^8)$, and the construction of G from F is thus performed in polynomial time.

We now prove that we have the desired equivalence between F and G. Assume first that G admits an m-irregular 2-labelling ℓ. We analyse how ℓ must behave in G. Let us start by considering the two forcing trails first. Let v be a trail vertex of degree d. Recall that we have attached, at w, a d-clique, and no other vertex has degree d. There are thus, in G, exactly $d+1$ vertices of degree d. Furthermore, by Lemma 3.1, we deduce that, by ℓ, the colour code of v must be $(1^0, 2^d)$ or $(1^d, 2^0)$. In other words, all edges incident to v must be assigned the same label. Now, since the same argument applies to every trail vertex, by repeatedly considering pairs of adjacent trail vertices (*i.e.*, sharing an incident edge), we deduce that, by ℓ, in each forcing trail all trail vertices must have all their incident edges being assigned the same label. Now, recall that, in each i-forcing trail, we have identified two degree-1 vertices to form the degree-2 vertex t_i. By the previous remarks, the colour code of t_i is either $(1^0, 2^2)$ or $(1^2, 2^0)$. Thus, so that $\mu(t_1) \neq \mu(t_2)$, we deduce that the edges of the 1-forcing trail are assigned the same label that must be different from that assigned to the edges of the 2-forcing trail. Free to swap labels 1 and 2 by ℓ, we may assume that all 1-edges are assigned label 1, while all 2-edges are assigned label 2.

- Now consider every clause vertex v_C. Recall that $d(v_C) = e(v_C)$, and, assuming this value is d, that three other vertices, a_C, b_C, and c_C, also have degree d. Furthermore, because of incident 1-edges and 2-edges, $d-3$ edges incident to v_C are assigned label 1, while, by construction, the colour codes of a_C, b_C, and c_C are exactly $(1^d, 2^0)$, $(1^{d-1}, 2^1)$, and $(1^{d-3}, 2^3)$, respectively. This means that the colour code of v_C is precisely $(1^{d-2}, 2^2)$, and, thus, that there must be exactly one formula edge incident to v_C assigned label 1.

- Similar deductions can be made for every variable vertex v_x. Setting $d = d(v_x) = e(v_x)$, recall that exactly two other vertices, a_x and b_x, have degree d. Also, due to how 1-edges and 2-edges were attached, v_x is incident to $d-3$ 1-edges, while the colour codes of a_x and b_x are $(1^{d-1}, 2^1)$ and $(1^{d-2}, 2^2)$. Thus the colour code of v_x is either $(1^d, 2^0)$ or $(1^{d-3}, 2^3)$, meaning that either all formula edges incident to v_x are assigned label 1 by ℓ, or they are all assigned label 2.

In G, imagine that having $\ell(v_x v_C) = 1$ models, in F, the fact that variable x brings truth value true to clause C, while having $\ell(v_x v_C) = 2$ models that x brings truth value false to C. The fact, in G, that all formula edges incident to v_x must be assigned the same label by ℓ thus models, in F, the fact that, by a truth assignment, x brings the same truth value to all clauses that contain it. Similarly, the fact, in G, that, for every clause vertex v_C, exactly one incident formula edge is assigned label 1 while the other two are assigned label 2 by ℓ, models, in F, the fact that C is considered satisfied by a truth assignment if and only if it has one true variable and two false ones. From this, we directly deduce, from ℓ, a 1-in-3 truth assignment to the variables of F.

Let us now focus on the converse direction, *i.e.*, assume F admits a 1-in-3 truth assignment ϕ to its variables. We construct an m-irregular 2-labelling ℓ of G in the following way. The key point is that vertices with the same degree in G form very particular sets, which is crucial since vertices with distinct degrees cannot have the same colour code. In particular, no two cliques we have attached to some vertices (either to w or to degree-1 vertices during the eventual clean-up process) are both k-cliques, meaning that these cliques can be labelled independently.

We start by considering any one trail vertex v of the 1-forcing trail, and assign label 1 by ℓ to all its incident edges. As described above, assuming v has degree d, we have attached a d-clique at w, and it can be checked that the vertices from that d-clique are the only vertices, besides v, that have degree d. According to Lemma 3.1, we can 2-label the edges of that clique in an m-irregular way and so that none of its vertices has colour code $(1^d, 2^0)$, thereby avoiding any conflict with v. Since the same arguments apply to all trail vertices, by arguments we have used earlier, we deduce that all edges of the 1-forcing trail must be assigned label 1, and for every trail vertex v of degree d, by Lemma 3.1 we can 2-label the edges of the d-clique attached at w in an m-irregular way and so that none of its vertices gets colour code $(1^d, 2^0)$. Now, since the two edges incident to t_1 are assigned label 1, the colour code of t_1 is $(1^2, 2^0)$.

We now apply all the exact same arguments to the 2-forcing trail, but with labelling all its edges with label 2. Note that this raises no conflict, since all trail vertices of the 2-forcing trail and their associated cliques have degrees more than those of the 1-forcing trail. Also, the colour code of t_2 is $(1^0, 2^2)$, and, thus, t_1 and t_2, the only two vertices of degree 2 of G, are not in conflict.

So, as previously, the 1-edges are all assigned label 1 by ℓ, while the 2-edges are all assigned label 2. By construction, every vertex a_C with degree d has colour code $(1^d, 2^0)$, every vertex b_C or

a_x with degree d has colour code $(1^{d-1}, 2^1)$, every vertex c_C with degree d has colour code $(1^{d-3}, 2^3)$, and every vertex b_x with degree d has colour code $(1^{d-2}, 2^2)$. This implies, since every clause vertex or variable vertex with degree d is incident to $d-3$ 1-edges, that a clause vertex must have colour code $(1^{d-2}, 2^2)$ while a variable vertex must have colour code $(1^d, 2^0)$ or $(1^{d-3}, 2^3)$. If we can achieve this, note that this will not raise any further conflict, since no other vertex of G has degree d. To achieve this, we just consider every variable x of F, and assign, in G, label 1 to all formula edges incident to v_x if ϕ sets x to true, while we assign label 2 to these edges otherwise. This results in v_x (of degree d) getting colour code $(1^d, 2^0)$ (if x is set to true by ϕ) or $(1^{d-3}, 2^3)$ (otherwise), while, due to ϕ 1-in-3-satisfying F, this results in every clause vertex v_C to get colour code $(1^{d-2}, 2^2)$.

To finish off the construction of ℓ, we lastly consider every clique attached at a pending 1-edge or 2-edge during the clean-up process, and the $(U+1)$-clique attached at w, and 2-label their edges arbitrarily in an m-irregular way, following Lemma 3.1. This results in ℓ being m-irregular. Indeed, w is the only vertex with maximum degree, so it cannot be involved in conflicts. For every degree value $d \in \mathcal{U}$, note that the only degree-d vertices are either 1) the vertices of a d-clique attached at w or at a (previously) pending vertex of a forcing trail, or 2) a single (previously pending) vertex to which a $(d-1)$-clique was attached; thus, in both cases, they cannot be involved in conflicts due to how G was labelled. Similarly, for a degree value $d \in \mathcal{T}$, only the vertices of a d-clique attached at w and one trail vertex have degree d, and they cannot be in conflict due to how ℓ was constructed. For a $d \in \mathcal{V}$, only one variable vertex v_x and two vertices a_x and b_x have degree d, and they are not in conflict due to how ℓ was obtained. Similarly, for a $d \in \mathcal{C}$, only one clause vertex v_C and three vertices a_C, b_C, and c_C have degree d, and they are not in conflict. Lastly, only t_1 and t_2 are of minimum degree, 2, and, as pointed out earlier, they are not in conflict. Thus, ℓ is m-irregular.

We thus have the desired equivalence between F and G. □

4. Discussion

Our goal in this work was to provide evidence, towards Question 1.1, that the problem of determining the irregularity strength of a given graph might be NP-hard. To that aim, we have considered two close variants of the problem, related to the distant irregularity strength and the multiset irregularity strength of graphs, and proved that these two problems are NP-hard. The interesting fact is that these two variants are close to the original one for different reasons: in one of the two variants, the distance requirement is weaker while the distinguishing requirement is similar, while, for the second variant, it is the other way round. Thus, in a sense, these two variants complement each other, with respect to the original problem.

Looking at the forcing mechanisms we employed in the reductions for proving Theorems 2.3 and 3.3, we can definitely come up with similar ones for the original irregularity strength. In particular, in both reductions, note that an important tool are cliques, which are very convenient for forbidding sums or multisets and making sure a vertex must be incident to certain labels by a labelling. A way to use cliques for the irregularity strength is for instance as follows. Let G be a graph with a vertex w, and, for some k, attach both a $(k+1)$-clique Q_1 and a $(2k+1)$-clique Q_2 at w. Assume further that, somewhere in G, there is another vertex v of degree k. Then note that, by all irregular 2-labellings of G, we must have $\sigma(v) = k$, *i.e.*, all edges incident to v must be assigned label 1. This is because the set of sums of the $k+1$ vertices of Q_1 must be either $\{k+1, \dots, 2k+1\}$ or $\{k+2, \dots, 2k+2\}$ while the set of sums of the $2k+1$ vertices of Q_2 must be either $\{2k+1, \dots, 4k+1\}$ of $\{2k+2, \dots, 4k+2\}$ (by Lemma 2.1), which forces the set of sums of the $3k+2$ vertices of Q_1 and Q_2 to be $\{k+1, \dots, 4k+2\}$ so that there is no conflict. Thus, since the sum of v must lie in $\{k, \dots, 2k\}$, so that there is no conflict, we indeed must have $\sigma(v) = k$.

The downside of this method, now, is that to force a single vertex v to have a particular sum, k, by any irregular 2-labelling, we had to make all values in $\{k+1, \dots, 4k+2\}$ appear as sums of some vertices. Then, if we want to force another vertex to have a particular sum $k' > k$, then the smallest k' we can consider is $4k+3$, which, by the method above, requires to attach both a $(4k+4)$-clique and a $(8k+7)$-clique at w. After that, if we want to force another vertex to have some sum $k'' > k'$, then the smallest k'' we can consider is $16k+15$, which requires to attach both a $(16k+16)$-clique and a $(32k+31)$-clique at w. And so on. This means that, with this method, the number of vertices of the resulting reduced graph would be exponential in the number of things we need to force, which is most probably a function of the input of the problem we reduce from. In other words, a reduction using such forcing mechanisms would not run in polynomial time.

Note that we did not run into this problem in the proofs of Theorems 2.3 and 3.3, as, for the distant irregularity strength, it is possible to have several cliques with the same number of vertices provided they are pairwise far from each other (which limits the growth of the reduced graph), while, for the multiset irregularity strength, the convenient property is that a k-clique and a k'-clique cannot have conflicting vertices provided $k \neq k'$ (which limits the size of the cliques, thus that of the reduced graph). In other words, while definitely close to the original irregularity strength, these two variants have peculiarities that make them way easier to work with.

To go further, probably a good direction could be to investigate the existence of other forcing graphs which, with respect to irregular 2-labellings, would be more permissive and less demanding than cliques. It would be crucial, for instance, to come up with graphs of reasonable size admitting certain sets S of values such that, by every irregular 2-labelling, no vertex has sum in S.

On a different note, several generalisations and restrictions of Theorems 2.3 and 3.3 could also be worth considering. For instance, Theorems 2.3 and 3.3 are about designing distinguishing 2-labellings, and one could wonder about a generalisation of these results to distinguishing k-labellings for any fixed $k \geq 3$. Other examples include restrictions of Theorems 2.3 and 3.3 to classes of graphs that are important in the field. In particular, with a bit of extra efforts, we believe our proof of Theorem 3.3 could be modified so that the result holds for connected graphs. Additionally, note that our reductions in the proofs of Theorems 2.3 and 3.3 provide graphs with unbounded maximum degree and that are very far from being planar (due to the presence of large cliques). Graphs with bounded maximum degree and planar graphs having received a lot of attention in this field, wondering about restrictions of our results to these classes of graphs would definitely be an appealing direction to consider.

References

[1] A. Ahadi, A. Dehghan, M.-R. Sadeghi. Algorithmic complexity of proper labeling problems. *Theoretical Computer Science*, 495:25–36, 2013.

[2] O. Baudon, J. Bensmail, T. Davot, H. Hocquard, J. Przybyło, M. Senhaji, É. Sopena, M. Woźniak. A general decomposition theory for the 1-2-3 Conjecture and locally irregular decompositions. *Discrete Mathematics & Theoretical Computer Science*, 21(1), 2019, #2

[3] J. Bensmail, H. Hocquard, P.-M. Marcille. Going Wide with the 1-2-3 Conjecture. *Discrete Applied Mathematics*, 320:15–32, 2022.

[4] G. Chartrand, H. Escuardro, F. Okamoto, P. Zhang. Detectable colorings of graphs. *Utilitas Mathematica*, 69(20):13–32, 2006.

[5] G. Chartrand, M.S. Jacobson, J. Lehel, O.R. Oellermann, S. Ruiz, F. Saba. Irregular networks. *Congressus Numerantium*, 64:197–210, 1988.

[6] A. Dudek, D. Wajc. On the complexity of vertex-coloring edge-weightings. *Discrete Mathematics and Theoretical Computer Science*, 13(3):45–50, 2011.

[7] J.A. Gallian. A dynamic survey of graph labeling. *Electronic Journal of Combinatorics*, 6, 1998.

[8] F. Havet, N. Paramaguru, R. Sampathkumar. Detection number of bipartite graphs and cubic graphs. *Discrete Mathematics & Theoretical Computer Science*, 16(3):642, 2014.

[9] M. Karoński, T. Łuczak, A. Thomason. Edge weights and vertex colours. *Journal of Combinatorial Theory, Series B*, 91:151–157, 2004.

[10] J. Misra, D. Gries. A constructive proof of Vizing's Theorem. *Information Processing Letters*, 41(3):131–133, 1992.

[11] C. Moore, J.M. Robson. Hard Tiling Problems with Simple Tiles. *Discrete and Computational Geometry*, 26(4):573–590, 2001.

[12] T. Nierhoff. A tight bound on the irregularity strength of graphs. *SIAM Journal on Discrete Mathematics*, 13(3):313–323, 2000.

[13] J. Przybyło. Distant irregularity strength of graphs. *Discrete Mathematics*, 313(24):2875–2880, 2013.

[14] V.G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskret. Analiz.*, 3:25–30, 1964.