Complexity of hypergraph coloring and Seidel's switching

Jan Kratochvíl

Department of Applied Mathematics and
Institute for Theoretical Computer Science*
Charles University
Malostranské nám. 25
118 00 Praha 1
Czech Republic
honza@kam.mff.cuni.cz

Abstract. Seidel's switching of a vertex in a given graph results in making the vertex adjacent to precisely those vertices it was nonadjacent before, while keeping the rest of the graph unchanged. Two graphs are called switching equivalent if one can be transformed into the other one by a sequence of Seidel's switchings. We consider the computational complexity of deciding if an input graph can be switched into a graph having a desired graph property. Among other results we show that switching to a regular graph is NP-complete. The proof is based on an NP-complete variant of hypergraph bicoloring that we find interesting in its own.

1 Seidel's switching

We consider undirected graphs without loops or multiple edges. The vertex set and the edge set of a graph G are denoted by V_G and E_G, respectively. Edges are considered as two-element subsets of the vertex set. A graph is called r-regular if every vertex has degree r. A bipartite graph is called (k,r)-biregular if every vertex in one bipartition class has degree k and every vertex in the other bipartition class has degree r. A hypergraph is a family of subsets (called hyperedges) of its vertex set. A hypergraph is called k-regular if every vertex belongs to exactly k hyperedges, and it is called k-uniform if every hyperedge has size k.

* Project LN00A056 supported by the Ministry of Education of the Czech Republic.

Let G be a graph. *Seidel's switching* of a vertex $v \in V_G$ results in a graph called $S(G,v)$ whose vertex set is the same as of G and the edge set is the symmetric difference of E_G and the full star centered in v, i.e.,

$$V_{S(G,v)} = V_G$$

$$E_{S(G,v)} = (E_G \setminus \{xv : x \in V_G, xv \in E_G\}) \cup \{xv : x \in V_G, x \neq v, xv \notin E_G\}.$$

Graphs G and H are called *switching equivalent* if G can be transformed into a graph isomorphic to H by a sequence of Seidel's switches. It can be easily seen that only the parity of the number of times a particular vertex is switched matters. Denote $A \subseteq V_G$ the set of vertices which are switched odd number of times. The resulting switched graph is then

$$S(G,A) = (V_G, E_G \div \{xy : x \in A, y \in V_G \setminus A\}),$$

and G is switching equivalent to H if and only if H is isomorphic to $S(G,A)$ for some $A \subseteq V_G$ ($\div$ denoting the symmetric difference of sets).

The concept of Seidel's switching was introduced by the Dutch mathematician J.J. Seidel in connection with symmetric structures, often of algebraic flavor, such as systems of equiangular lines, strongly regular graphs, or the so called two-graphs. For more structural properties of two-graphs cf. [13–15].

Colbourn and Corneil [1] (and independently but later Kratochvil et al. [10]) proved that deciding if two graphs are switching equivalent is an isomorphism complete problem. This observation can be extended to the fact that deciding if an input graph is switching equivalent to its complement is also isomorphism complete [11].

Several authors asked the question of how difficult it is to decide if a given graph is switching equivalent to a graph having some prescribed property (this property becomes the parameter of the problem). It was noted in [10], and later also in [3], that there is no correlation between the complexity of this problem and the complexity of the property itself. For instance, it is shown that every graph is switching equivalent to a graph containing a hamiltonian path, and it is polynomial to decide if a given graph is switching equivalent to a graph containing a hamiltonian cycle (while both deciding if the graph itself contains a hamiltonian path or cycle are NP-complete problems). This result was extended in [2] to graph pancyclicity. Hage et al. further show in [5] that switching equivalence to Euler graphs and to bipartite graphs are polynomially solvable problems.

An elegant characterization by forbidden induced subgraphs for the property "not being perfect" is proven in [7, 8]. This characterization yields a polynomial time decision algorithm. (Equivalently, this provides a characterization of graphs such that every switching equivalent graph is perfect.) Also graphs such that all switching equivalent ones contain perfect dominating sets (perfect codes) can be characterized by a finite set of forbidden induced subgraphs and hence are recognizable in polynomial time [9].

Another direction is avoiding forbidden induced subgraphs. It is proved in [10] that deciding if a given input graph can be switched to a P_3-free graph (i.e., a graph not containing and induced copy of the path on 3 vertices) is polynomially solvable. This means deciding if the input graph is switching equivalent to the disjoint union of complete graphs. R. Hayward [6] (and later Hage et al. [5]) showed that deciding switching equivalence to triangle-free graphs is also polynomial. Somewhat surprisingly, this result is a core of the polynomial algorithm for recognizing P_3-structures of graphs [6]. (The H-structure of a graph is the hypergraph on the same vertex set whose hyperedges are those subsets of vertices that induce graphs isomorphic to H.) The question of recognizing P_3-structures was motivated by the class of perfect graphs and P_4-structures, since induced P_4's play an important role for perfectness. Recognition of P_4-structures is still an open problem.

It was also announced (but not proved) in [10] that deciding switching equivalence to a regular graph is NP-complete. Given that recognizing regular graphs is linear, this result may seem somewhat unexpected. In this paper, we present a proof significantly simpler than the original (and so far unpublished) one. The proof is based on NP-hardness of a special hypergraph bicoloring problem (Theorem 1) which is proven in the next section. We believe that this result is interesting in its own, also for the consequence of balancing biregular bipartite graphs. In the last section we conclude with observations on switching to graphs of bounded minimum degree.

2 Bicoloring uniform regular hypergraphs

Satisfiability of Boolean formulas is considered one of the basic and most useful (for further consequences) NP-complete problems. Special variants as 3-SAT, 1-in-3-SAT etc. are known to be NP-complete, and are often used in deriving other hardness results. A remarkable theorem of Schaefer [12] gives a complete characterization of predicates that determine polynomially solvable or NP-complete instances. Schaefer, however, does not take into account further regularity requirements on the input formula.

Let us have a closer look at formulas without negative literals. We can view such a formula as a hypergraph – vertices being the variables and hyperegdes being the clauses. A truth valuation of the variables is then nothing else than a coloring of the vertices by two colors. A clause is not-all-equal satisfied if and only if the corresponding hyperedge is colored so that it is not monochromatic (contains at least one vertex colored true and at least one vertex colored false). Hence Not-All-Equal-SAT of formulas without negations is exactly the hypergraph bicoloring problem, well known to be NP-complete.

Another well known NP-complete problem is 1-in-3-SAT where we require that every clause (all of which are of size 3) contains exactly one true variable. In the hypergraph setting this problem asks if the vertices can be colored by two colors so that every hyperedge contains exactly one vertex colored by a favored color. In the dual setting this is equivalent to the question if the hyperedges can be colored by two colors so that every vertex belongs to exactly one hyperedge of a favored color, ergo, if one can choose a set of hyperedges which cover every vertex exactly once, the problem known as EXACT COVER. This problem is NP-complete even if every hyperedge has size 3 and every vertex belongs to exactly 3 hyperdges [4].

The aim of this section is to further explore regular instances of the hypergraph bicoloring problem. Recall that a hypergraph is called *q-regular* if every vertex belongs to exactly q hyperedges, and *m-uniform* if every hyperedge has size m. Thus the bipartite incidence graph of a q-regular m-uniform hypergraph is (q, m)-biregular. We say that an m-uniform hypergraph is *$(k$-in-$m)$-colorable* if the vertices can be colored by two colors (say black or white) so that every hyperedge contains exactly k white vertices. We prove the following result.

Theorem 1. *For every $q \geq 3$, $m \geq 3$ and $1 \leq k \leq m-1$, deciding $(k$-in-$m)$-colorability of q-regular m-uniform hypergraphs is NP-complete.*

Proof. The case $q = 3, k = 1, m = 3$ is the variant of EXACT COVER mentioned above. We use it to show NP-completeness of the case $q = 3$ (k, m arbitrary). Assume that we have a 3-regular 3-uniform hypergraph $H = (V, E)$ and we ask if it is 1-in-3-colorable. Let $n = |E|$ be the number of hyperedges of H. Without loss of generality we may assume that n is divisible by 3 (if not, we simply reject H as not 1-in-3-colorable).

First we construct a hypergraph B called *a block* as follows:

$$V_B = \{a_i : i = 1, 2, \ldots, m-1\} \cup \{x_1, x_2, x_3\}$$

$$E_B = \{R_1, R_2, R_3\}$$

where

$$R_i = (a_1, a_2, \ldots, a_{m-1}, x_i).$$

Straightforwardly, if B is (k-in-m)-colored then $x_i, i = 1, 2, 3$ have the same color, and vice versa, B can be (k-in-m)-colored with all x_i's having the same color.

Now blocks are combined into chains: Take $r = \frac{2}{3}n(k-1) + 2k$ copies of B, say $B^j, j = 1, 2, \ldots, r$, with the x vertices denoted by x_i^j, $i = 1, 2, 3$, $j = 1, 2, \ldots, r$. Then identify $x_1^j \cong x_3^{j+1}$, $j = 1, 2, \ldots, r$ (counting in subscripts modulo r) and $x_2^{2j} \cong x_2^{2j-1}$ for $j = 1, 2, \ldots, \frac{r}{2}$ (note that r is even). These newly created x vertices must receive the same color in every k-in-m-coloring. Similarly, take $s = \frac{2}{3}n(m-k-2) + 2(m-k)$ copies of B, say $C^j, j = 1, 2, \ldots, s$, with the x vertices denoted by y_i^j, $i = 1, 2, 3$, $j = 1, 2, \ldots, s$. Then identify $y_1^j \cong y_3^{j+1}$, $j = 1, 2, \ldots, s$ and $y_2^{2j} \cong y_2^{2j-1}$ for $j = 1, 2, \ldots, \frac{s}{2}$ (note that s is again even). Again these newly created y vertices must receive the same color in every k-in-m-coloring. Since each of the x or y vertices is contained in exactly 2 hyperedges of $\bigcup_{j=1}^{r} B^j \cup \bigcup_{j=1}^{s} C^j$, we have created $n(k-1) + 3k$ x-vertices and $n(m-k-2) + 3(m-k)$ y-vertices. (The a vertices of the blocks B^j, C^j are all mutually distinct.)

For every hyperedge $e \in E$, add $k-1$ distinct x vertices and $m-k-2$ distinct y vertices to this edge, thus creating a hyperedge $\widetilde{e}$. Distinct hyperedges are completed by mutually distinct vertices. From the remaining $3k$ x-vertices and $3(m-k)$ y-vertices, create 3 auxiliary hyperedges A_1, A_2, A_3, each containing exactly k x-vertices and $m-k$ y-vertices. Denote this hypergraph $\widetilde{H} = (\widetilde{V}, \widetilde{E})$.

Obviously, every hyperedge of $\widetilde{E}$ has exactly m vertices and every vertex of $\widetilde{V}$ is contained in exactly 3 hyperedges. We claim that $(\widetilde{V}, \widetilde{E})$ is (k-in-m)-colorable if and only if (V, E) is (1-in-3)-colorable.

Suppose $\phi : \widetilde{V} \to \{\mathsf{black}, \mathsf{white}\}$ is a coloring such that in every $\widetilde{e} \in \widetilde{E}$ there are exactly k white vertices. Since all x-vertices have the same color, and so do all the y-vertices, and since A_1 should have exactly k white vertices, necessarily $\phi(x) =$white and $\phi(y) =$black (or possibly $\phi(x) =$black and $\phi(y) =$white when $m = 2k$). If $\phi(x) =$white then in every clause $\widetilde{e}$ there are exactly $k-1$ white vertices among those added with respect to e, and hence e has exactly 1 white vertex, i.e., H is (1-in-3)-colorable. (If $m = 2k$ and $\phi(x) =$black then $\widetilde{e}$ would have $m-k-2 = k-2$ white vertices among the new ones, and e would have to contain exactly 2 white vertices, i.e. (V,E) would be (2-in-3)-colorable, and hence also (1-in-3)-colorable.)

The opposite implication, i.e., that $(\widetilde{V},\widetilde{E})$ is (k-in-m)-colorable provided (V,E) is (1-in-3)-colorable, is obvious.

For the case of $q>3$, we reduce from $q=3$. Let $H=(V,E)$ be a hypergraph with every hyperedge of size m and every vertex occurring in 3 hyperedges. We construct $H'=(V',E')$ by taking q copies of H with vertex $v\in V$ called v^i in the i-th copy, for $i=1,2,\ldots,q$. Then for each vertex $v\in V$, we add $(q-3)$ copies of the block $B(q)$, with vertices $x_1,x_2,\ldots,x_q$ being identified with $v^1,v^2,\ldots,v^q$ (respectively). The block $B(q)$ has vertices $\{a_1,a_2,\ldots,a_{m-1},x_1,x_2,\ldots,x_q\}$ and hyperedges $e_i=\{a_1,a_2,\ldots,a_{m-1},x_i\}$ for $i=1,2,\ldots,q$. Every vertex a_j has degree q, every vertex x_j has degree one and in any (k-in-m)-coloring, the vertices $x_1,x_2,\ldots,x_q$ receive the same color, and both colors are feasible. This construction is just the garbage collection, since each copy of H already has to be (k-in-m)-colored. Now H' is q-regular and an (k-in-m)-coloring of H' is obtained by taking the same (k-in-m)-coloring of each copy of H.

Corollary 1. *For every $q\geq 3$, $m\geq 3$ and $1\leq k\leq m-1$, k-in-m-SAT is NP-complete for formulas without negations with every variable occurring in exactly q clauses.*

Note that the bounds $q\geq 3$ and $m\geq 3$ are sharp, since the case $m=2$ is a special instance of 2-SAT and hence polynomially solvable, while in the case of $q=2$ the variables can be regarded as edges of an m-regular graph and the question is equivalent to deciding if an m-regular graph contains a spanning k-regular subgraph, a problem well known to be polynomially solvable by matching techniques.

Next we prove an auxiliary result on bipartite graphs which we will use in the proof of Theorem 2. We say that a graph is *balanced* if its vertices can be colored by two colors (say black and white) so that every vertex has the same number of white and black neighbors (obviously, a necessary condition is that every vertex has even degree). Deciding if a graph is balanced is a hard problem, even in a very restricted case.

Proposition 1. *For all $p,q\geq 2$, it is NP-complete to decide if a $(2p,2q)$-biregular bipartite graph is balanced.*

Proof. Assume first that one of p, q is greater than 2, say $q \geq 3$. Consider a q-regular $2p$-uniform hypergraph H as an instance of (p-in-$2p$)-colorability, which has been just shown NP-complete. Duplicate every hyperedge, obtaining a $2q$-regular $2p$-uniform hypergraph H' whose (p-in-$2p$)-colorability status is the same as of H. We can color the hyperedges of H white and their duplicates black to balance the vertices of V_H, and hence the incidence graph of H' is balanced if and only if H' is (p-in-$2p$)-colorable.

Also in the case $p = q = 2$ we reduce from (2-in-4)-colorability of 3-regular hypergraphs, just the construction is slightly more technical. Given a 3-regular 4-uniform hypergraph $H = (V, E)$, we construct a hypergraph $H' = (V', E')$ by replacing every hyperedge of H by four new hyperedges and four new vertices as follows:

$$V' = V \cup \{v_e : v \in e \in E\}$$

$$E' = \{e_v : v \in e \in E\}$$

with the newly introduced hyperedges being

$$e_v = \{v\} \cup \{x_e : x \in e, x \neq v\}.$$

We claim that H' is (2-in-4)-colorable if and only if H is, and that the hyperedges of H' can be bicolored so that every vertex belongs to at least one hyperedge of each color.

Suppose H is (2-in-4)-colored by a coloring $\varphi : V \longrightarrow \{\mathsf{black}, \mathsf{white}\}$. We extend it to a (2-in-4)-coloring φ' of H' by setting

$$\varphi'(x_e) = \varphi'(x) = \varphi(x)$$

for every $x \in e \in E$. Consider a hyperedge $e \in E$. It has two black and two white vertices, say $e = \{u, v, w, z\}$ with $\varphi(u) = \varphi(v) = \mathsf{black}$ and $\varphi(w) = \varphi(z) = \mathsf{white}$. Then e_v has two black vertices v, u_e and two white vertices w_e, z_e. Similarly for the other three edges e_u, e_w, e_z.

Next suppose that the vertices of H' are bicolored by a bicoloring φ so that every hyperedge has two white and two black vertices. We argue that the restriction of φ to V is a (2-in-4)-coloring of H. Consider a hyperedge $e \in E$, say again $e = \{u, v, w, z\}$. For any three of u_e, v_e, w_e, z_e, E' contains a hyperedge containing all these three vertices, and hence no three of u_e, v_e, w_e, z_e may have the same color. It follows that exactly two of u_e, v_e, w_e, z_e are black and exactly two are white. But then $\varphi(x_e) = \varphi(x)$ for every $x \in e$, and e has two black and two white vertices as claimed.

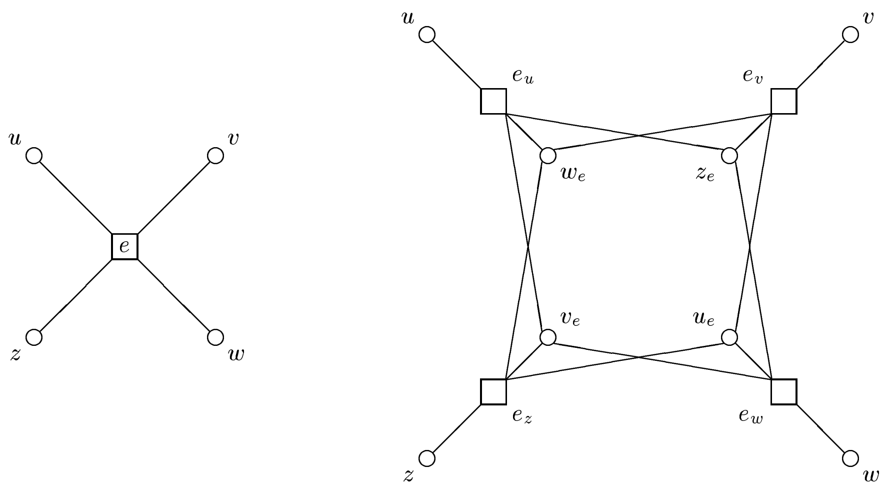

Fig. 1. Construction of H' from H, depicted as the incidence graphs B and B'.

Now we show how to color the hyperedges of H'. Let $B = (V \cup E, \{xe : x \in e \in E\})$ be the bipartite incidence graph of H. We orient the edges of B so that every V-vertex is incident with at least one ingoing and at least one outgoing edge, and each E-vertex has indegree and outdegree two. (To find such an orientation, add to B a dummy vertex adjacent to all V-vertices. This enriched graph has all vertices of even degrees and hence allows an Eulerian walk. Orienting the edges along this walk yields an orientation in which every vertex of $V \cup E$ has indegree and outdegree two.) Color $e_v \in E'$ black if the edge ve of B is oriented from v to e, and color it white otherwise. If e, f, g are the hyperedges of H that contain a vertex $v \in V$, at least one of the edges ve, vf, vg is oriented out of v and at least one is oriented toward v, and hence at least one of the hyperedges e_v, f_v, g_v of H' is colored white and at least one is colored black. For the new vertices of $V' \setminus V$, consider $e = \{u, v, w, z\} \in E$. Since e has indegree and outdegree two (in the orientation of B), exactly two of the hyperedges e_u, e_v, e_w, e_z are black and exactly two are white. Thus each of the vertices v_e, u_e, w_e, z_e is incident with at least one black and at least one white hyperedge.

Finally we take four disjoint copies of H' with corresponding vertices and hyperedges indexed $v^i, e^i, i = 1, 2, 3, 4$, and we add transversal hyperedges

$$t_v = \{v^1, v^2, v^3, v^4\}$$

for all $v \in V'$. Obviously, the hypergraph $\widetilde{H}$ constructed in this way is 4-regular and 4-uniform, and hence its bipartite incidence graph $\widetilde{B}$ is 4-regular. Since $\widetilde{H}$ contains H' as an induced subhypergraph, $\widetilde{H}$ may be (2-in-4)-colorable only if H' is. On the other hand, using a (2-in-4)-coloring of H' on two copies of H' and its reverse on the other two copies gives a (2-in-4)-coloring of $\widetilde{H}$. The hyperedges of $\widetilde{H}$ can always be colored so that every vertex is incident with two black and white hyperedges. Use the same coloring of edges (guaranteed by the previous claim) on each copy of H', and color the transversal hyperedges as forced (e.g., if v belongs to two black and one white hyperedge in the coloring of the edges of H', color the transversal hyperedge t_v white).

The bipartite incidence graph $\widetilde{B}$ of $\widetilde{H}$ is then balanced if and only if H is (2-in-4)-colorable, which concludes the proof.

Also in Proposition 1 the bounds on p and q are tight. For one can decide in polynomial time if a $(2, 2q)$-biregular bipartite graph is balanced for every q. The neighbors of the vertices of degree 2 can be properly colored if and only if the graph does not contain a cycle of length 2 modulo 4. And if this is satisfied, we contract these vertices into edges and decide via matching algorithm if the resulting $2q$-regular graph contains a spanning q-regular subgraph.

3 Complexity of switching to regular graphs

Theorem 2. *Deciding if an input graph is switching equivalent to a regular graph is NP-complete.*

The theorem follows directly from Proposition 1 and the following result.

Proposition 2. *Let G be a $(2p, 2q)$-biregular bipartite graph with $n > 2(p+q)$ vertices. If $p \neq q$, then G is switching equivalent to a regular graph if and only if G is balanced. The resulting switched graph is then $\frac{n}{2}$-regular.*

Proof. Let $G = (V_1 \cup V_2, E)$ be a bipartite $(2p, 2q)$-biregular graph with bipartition classes V_1 and V_2, such that $\deg_G u = 2p$ for every $u \in V_1$ and $\deg_G u = 2q$ for every $u \in V_2$. Suppose that switching a subset A of vertices of G results in a D-regular graph. Denote

$$X = A \cap V_1, \quad Y = A \cap V_2,$$

$$Z = V_1 \setminus A, \quad W = V_2 \setminus A$$

and

$$x = |X|,\ y = |Y|,\ z = |Z|,\ w = |W|.$$

We will refer to X, Y, Z, W as *blocks*.

Denote by $\deg_G(u, S) = |N_G(u) \cap S|$ the number of neighbors of a vertex u in a set S. Consider a vertex $u \in X$. Its degree in the switched graph $S(G, A)$ is $\deg_{S(G,A)} u = \deg_G(u, Y) + z + w - \deg_G(u, W) = w + z + 2\deg_G(x, Y) - 2p$, since $\deg_G(u, Y) + \deg_G(u, W) = 2p$. If $S(G, A)$ is D-regular, $\deg_G(u, Y) = \frac{1}{2}(D + 2p - w - z)$ and hence all vertices in X have the same number of neighbors in Y. Similarly for the number of neighbors in W and for vertices in other blocks. It follows that the partition $V_G = X \cup Y \cup Z \cup W$ is regular, i.e., any two vertices from the same block have the same number of neighbors in each other particular block (adjacency considered in the original graph G). It follows that there exist numbers a, b, c, d, $-p \le a, b \le p$, $-q \le c, d \le q$, such that

$$\deg_G(u, Y) = p - a, \quad \deg_G(u, W) = p + a, \quad \deg_G(u, Z) = 0, \quad \text{for } u \in X,$$

$$\deg_G(u, Y) = p - b, \quad \deg_G(u, W) = p + b, \quad \deg_G(u, X) = 0, \quad \text{for } u \in Z,$$

$$\deg_G(u, X) = q - c, \quad \deg_G(u, Z) = q + c, \quad \deg_G(u, W) = 0, \quad \text{for } u \in Y,$$

$$\deg_G(u, X) = q - d, \quad \deg_G(u, Z) = q + d, \quad \deg_G(u, Y) = 0, \quad \text{for } u \in W.$$

This is illustrated in Figure 2, where the symbol close to each set denotes the size of the set and the symbol at the beginning of an edge ST close to S denotes the number of neighbors of a vertex from S in T, with $S, T \in \{X, Y, Z, W\}$.

The subgraph $G[X \cup Y]$ is $(p-a, q-c)$-biregular, and hence (by counting the number of its edges)

$$x(p - a) = y(q - c), \tag{1}$$

and similarly

$$w(q - d) = x(p + a), \tag{2}$$

$$y(q + c) = z(p - b), \tag{3}$$

$$z(p + b) = w(q + d). \tag{4}$$

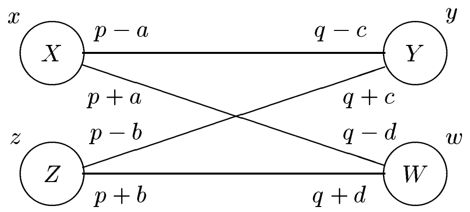

Fig. 2. Schematic diagram of the regular partition $V_G = X \cup Y \cup Z \cup W$.

Multiplying these four equations we get

$$xyzw(p-a)(p+b)(q+c)(q-d) = xyzw(p+a)(p-b)(q-c)(q+d). \quad (5)$$

We will show that all four blocks X, Y, Z, W are nonempty, i.e., that we can divide this equation by $xyzw \neq 0$. But first we list the equations for the degrees of vertices from each block in the switched graph, which should by assumption equal D. For instance, for $u \in X$, we have

$$\deg_{S(G,A)} u = p - a + z + w - (p + a) = z + w - 2a, \quad (6)$$

and similarly

$$u \in Y : \quad \deg_{S(G,A)} u = z + w - 2c, \quad (7)$$

$$u \in Z : \quad \deg_{S(G,A)} u = x + y + 2b, \quad (8)$$

$$u \in W : \quad \deg_{S(G,A)} u = x + y + 2d. \quad (9)$$

Once we prove that $x \neq 0 \neq y$, i.e., that X and Y are nonempty, it will follow from (6) and (7) that $a = c$. Similarly, from $zw \neq 0$ will follow that $b = d$.

To prove that $xyzw \neq 0$, assume that one of the blocks is empty. Say $X = \emptyset$ (by symmetry, this will cover all cases). Since $X \cup Z = V_1$, this means that $Z \neq \emptyset$. Also Y must be nonempty, since otherwise the switching set $A = X \cup Y$ would be empty and the switched graph $S(G, A) = G$ would not be regular. Hence $c = d = q$ and we deduce from (7,8) that

$$z + w - 2q = y + 2b. \quad (10)$$

From (3,4) we get

$$y = \frac{z(p-b)}{2q}, \quad w = \frac{z(p+b)}{2q}.$$

Substituting into (10) yields

$$z + \frac{z(p+b)}{2q} - 2q = \frac{z(p-b)}{2q} + 2b$$

which implies

$$(z - 2q)(b + q) = 0.$$

Now $z = 2q$ implies $y + w = 2p$ and $n = |V_G| = 2(p + q)$, contradicting the assumption of the theorem. Also $b = -q$ leads to a contradiction, since in this case $p + b = p - q \neq 0$ implies that $W \neq \emptyset$ and hence (9,8) give $b = q$.

From now on we will assume without loss of generality that $p < q$. We further note that $-p < a, b < p$. For $a = p$ would imply $c = q$, which is impossible since $c = a = p < q$, while $a = -p$ would imply $b = q > p$. Similarly, $b \neq -p, p$.

Dividing (5) by $xyzw$ and using the fact that $c = a$ and $d = b$ we get

$$(p-a)(p+b)(q+a)(q-b) = (p+a)(p-b)(q-a)(q+b)$$

which yields

$$\frac{(p-a)(q+a)}{(p+a)(q-a)} = \frac{(p-b)(q+b)}{(p+b)(q-b)}.$$

Consider the function $f(b) = \frac{(p-b)(q+b)}{(p+b)(q-b)}$ as a function of b. Its derivative is

$$f'(b) = \frac{2(p-q)(pq+b^2)}{(p+b)^2(q-b)^2} < 0$$

and so $f(b)$ is decreasing in the interval $(-p, q)$. Hence $b = a$ is the only solution of $f(a) = f(b)$ in the interval $b \in (-p, p)$.

Therefore, (1–4) yield

$$y = x\frac{p-a}{q-a}, \quad z = x\frac{q+a}{q-a}, \quad w = x\frac{p+a}{q-a}$$

and substituting these into

$$w + z - x - y = 4a$$

(which is obtained by comparing (7) and (8)) gives

$$4a(x - q + a) = 0.$$

Now $x = q - a$ would imply $y = p - a$, $z = q + a$ and $w = p + a$, hence $n = x + y + z + w = 2(p + q)$ contradicting the assumption. Thus we derive $a = 0$ as the only possibility, yielding

$$\deg_G(u, A) = \deg_G(u, V_G \setminus A) = p \text{ for every } u \in V_1 = X \cup Z$$

and

$$\deg_G(u, A) = \deg_G(u, V_G \setminus A) = q \text{ for every } u \in V_2 = Y \cup W.$$

This shows that G is balanced by coloring the vertices of A white and the vertices of $V_G \setminus A$ black.

On the other hand, if G is balanced by a coloring $\phi : V_G \to \{\mathsf{black}, \mathsf{white}\}$, then switching the set of white vertices yields an $\frac{n}{2}$-regular graph. To see this, denote $A = \{u \in V_G : \phi(u) = \mathsf{white}\}$ and set $X = A \cap V_1$, $Y = A \cap V_2$, $Z = V_1 \setminus A$, $W = V_2 \setminus A$. Since $G[X \cup Y]$ and $G[Z \cup Y]$ are both (p, q)-biregular, we have $p|X| = q|Y| = p|Z|$ and hence $|X| = |Z|$. Similarly, $|Y| = |W|$. For any vertex $u \in X$, we have $\deg_{S(G,A)} u = p + |Z| + |W| - p = |Z| + |W| = \frac{|V_G|}{2}$, and by symmetry, $\deg_{S(G,A)} u = \frac{|V_G|}{2}$ for every $u \in V_G$.

4 Switching to graphs of bounded minimum degree

One can easily see that for every fixed k, one can decide switchability to a k-regular graph in polynomial time. We frame this fact in a more general observation.

Proposition 3. *Let $\mathcal{A}$ be an isomorphism-closed class of graphs such that every graph of $\mathcal{A}$ contains a vertex of degree at most k, for some fixed number k. If $\mathcal{A}$ can be recognized in polynomial time, then it can also be decided in polynomial time if an input graph is switching equivalent to a graph belonging to $\mathcal{A}$. More precisely, this can be decided in time $O(n^{k+3}p(n))$, where $p(n)$ is the worst case time complexity of recognizing graphs of $\mathcal{A}$.*

Proof. If the input graph is switched to a graph of $\mathcal{A}$, one of its vertices is switched to a vertex of degree at most k. There are n possible choices of which vertex will this be. For each such vertex, there are $1 + (n-1) + \binom{n-1}{2} + \ldots + \binom{n-1}{k} = O(n^k)$ possible sets of at most k neighbors for that vertex (in the switched graph). So for each vertex u and for every choice of a set S of at most k other vertices, we switch the set $N_G(u) \div S$ and check if the resulting graph is in $\mathcal{A}$. Each switching can only affect $O(n^2)$ edges, hence the upper bound.

Corollary 2. *Switching equivalence to a tree, acyclic graph, planar graph, outerplanar graph, graph of bounded genus, graph of bounded tree-width, k-regular graph (for fixed k) are all polynomially decidable problems.*

One can argue about the necessity of the assumption that graphs of $\mathcal{A}$ are polynomial time recognizable as follows. It is NP-complete to decide if an input graph of maximum degree 4 is 3-colorable. For any graph G, the disjoint union of G and a triangle is switchable to a 3-colorable graph if and only if G itself is 3-colorable. Hence deciding if an input graph is switchable to a 3-colorable graph of maximum degree at most 4 (as well as deciding switchability to 3-colorable graphs) is NP-complete.

References

1. C.J. Colbourn, D.G. Corneil: On deciding switching equivalence of graphs, Discrete Appl. Math. 2 (1980) 181-184
2. A. Ehrenfeucht, J. Hage, T. Harju, G. Rozenberg: Pancyclicity in switching classes, Inform. Process. Letters 73 (2000) 153-156
3. A. Ehrenfeucht, J. Hage, T. Harju, G. Rozenberg: Complexity Issues in switching of graphs, in: Theory and Application to Graph Transformations, LNCS 1764, Springer-Verlag, 2000, pp. 59-70
4. T.F. Gonzales: Clustering to minimize the maximum intercluster distance, Theoret. Comput. Sci. 38 (1985), 293-306
5. J. Hage, T. Harju, E. Welzl: Euler graphs, triangle-free graphs and bipartite graphs in switching classes, in: Proceedings ICGT 2002, LNCS 2505, Springer-Verlag, 2002, pp. 148-160
6. R.B. Hayward: Recognizing P-3-structure: A switching approach, J. Combin. Th. Ser. B 66 (1996) 247-262
7. A.Hertz: On perfect switching classes, Discr. Appl. Math. 89 (1998) 263-267
8. A.Hertz: On perfect switching classes, Discr. Appl. Math. 94 (1999) 3-7
9. J. Kratochvil: Perfect codes and two-graphs, Comment. Math. Univ. Carolin. 30 (1989), 755-760
10. J. Kratochvíl, J. Nešetřil, O. Zýka: On the computational complexity of Seidel's switching, Combinatorics, graphs and complexity, Proc. 4th Czech. Symp., Prachatice 1990, Ann. Discrete Math. 51 (1992) 161-166

11. J. Kratochvil, M. Rosenfeld: unpublished
12. T. J. Schaefer: The complexity of satisfiability problems, Proc. of the Tenth Annual ACM Symposium on Theory of Computing (STOC), 1978, 216–226
13. J.J. Seidel: A survey of two-graphs, Teorie combinatorie, Atti Conv. Lincei, Vol 17, Academia Nazionale dei Lincei, Rome, 1973, pp. 481-511
14. J.J. Seidel, D.E. Taylor: Two-graphs, a second survey, Algebraic methods in graph theory, Vol. II, Conf. Szeged 1978, Colloq. Math. Janos Bolyai 25 (1981) 689-711
15. J.J. Seidel: More about two-graphs, Combinatorics, graphs and complexity, Proc. 4th Czech. Symp., Prachatice 1990, Ann. Discrete Math. 51 (1992) 297-308