Discrete Applied Mathematics 157 (2009) 780–793

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Paintshop, odd cycles and necklace splitting

Frédéric Meunier [a,*], András Sebő [b]

[a] Université Paris Est, LVMT, Ecole Nationale des Ponts et Chaussées, 6-8 avenue Blaise Pascal, Cité Descartes Champs-sur-Marne, 77455 Marne-la-Vallée cedex 2, France

[b] CNRS, Laboratoire G-SCOP (INPG, CNRS, UJF), 46, Avenue Félix Viallet, 38000 Grenoble, France

ARTICLE INFO

Article history:
Received 11 January 2008
Received in revised form 12 June 2008
Accepted 12 June 2008
Available online 29 July 2008

Keywords:
Binary clutter
Max-cut
Odd cycles
Paintshop problem

ABSTRACT

The following problem has been presented in [T. Epping, W. Hochstättler, P. Oertel, Complexity results on a paint shop problem, Discrete Applied Mathematics 136 (2004) 217–226] by Epping, Hochstättler and Oertel: cars have to be painted in two colors in a sequence where each car occurs twice; assign the two colors to the two occurrences of each car so as to minimize the number of color changes. More generally, the "paint shop scheduling problem" is defined with an arbitrary multiset of colors given for each car, where this multiset has the same size as the number of occurrences of the car; the mentioned article states two conjectures about the general problem and proves its NP-hardness. In a subsequent paper in [P. Bonsma, Th. Epping, W. Hochstättler, Complexity results for restricted instances of a paint shop problem for words, Discrete Applied Mathematics 154 (2006) 1335–1343], Bonsma, Epping and Hochstättler proved its APX-hardness and noticed the applicability of some classical results in special cases.

We first identify the problem concerning two colors as a minimum odd circuit cover problem in particular graphs, exactly situating the problem. A resulting two-way reduction to a special minimum uncut problem leads to polynomial algorithms for subproblems, to observing APX-hardness through MAX CUT in 3-regular graphs, and to a solution with at most 3/4th of all possible remaining color changes (when all obliged color changes have been made).

For the general problem concerning an arbitrary number of colors, we realize that the two aforementioned conjectures are corollaries of the celebrated "necklace splitting" theorem of Alon, Goldberg and West.

1. Introduction

1.1. Problem formulation

In [7], Epping, Hochstättler and Oertel introduced the following problem. The origins of the model lie in car manufacturing with individual demands, which is reported to occur often in Europe.

Given a sequence of cars where repetition can occur, and for each car a multiset of colors where the sum of the multiplicities is equal to the number of repetitions of the car in the sequence, decide the color to be applied for each occurrence of each car so that each color occurs with the multiplicity that has been assigned. The goal is to minimize the number of color changes in the sequence. If cars are considered to be letters in an alphabet, the following is a formalization.

PPW [Paint Shop Problem for Words]. Given a finite alphabet Σ whose elements are called *letters*, a word $w = (w_1, \dots, w_n) \in \Sigma^*$, a finite color set F, and a *coloring* $f = (f_1, \dots, f_n)$ of w with $f_i \in F$ for $i = 1, \dots, n$, find a permutation

* Corresponding author. Fax: +33 0 1 64 15 21 40.
E-mail addresses: frederic.meunier@enpc.fr (F. Meunier), andras.sebo@g-scop.inpg.fr (A. Sebő).

0166-218X/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.dam.2008.06.017

σ of $\{1, \ldots, n\}$ such that $w_{\sigma(i)} = w_i$ for $i = 1, \ldots, n$, and the number of color changes within $\sigma(f) = (f_{\sigma(1)}, \ldots, f_{\sigma(n)})$ is minimized.

We say that we have a *color change* in f whenever $f_i \neq f_{i+1}$. The minimum of the number of color changes is denoted $\gamma = \gamma(w; f)$.

In [7] this problem has been solved with a dynamic program which can be implemented to run with a space and time complexity of $O(|F|n^{(|F|-1)|\Sigma|})$: this bound is exponential unless both $|F|$ and $|\Sigma|$ are fixed; if any of them is not fixed the problem is proved to be NP-hard in the same paper.

The problem PPW restricted to instances where the number of colors is c, each color occurs k times with each letter, and accordingly each letter occurs ck times in the input sequence, is denoted by PPW(c, k). Such instances are called *k-regular*. Two conjectures are stated, the following one, and its special case to $c = 2$:

Conjecture 1. *For any instance in* PPW(c, k) *we have* $\gamma \leq |\Sigma|(c-1)$, *independently of* k.

In subsequent papers [5,6], Bonsma, Epping and Hochstättler claim the APX-hardness of PPW(2, 1) and transform the problem into matroid theory. PPW(2, 1) plays a particular role because of several "natural interpretations", and the related relevant new optimization problem: paint a set of objects, for instance cards, with two colors, one color for each face, if the objects arrive to the painting machines in a given order, each object twice. Minimize the number of color-changes. In the manufacturing of cars or other objects the same problem arises for more than one color.

1.2. Main notions and results

Let S be a finite set. A hypergraph $\mathcal{H} \subseteq 2^S$ is a *clutter* if none of the hyperedges contains another. The hypergraph $\mathcal{B} \subseteq 2^S$, $\mathcal{B} \neq \emptyset$, $\mathcal{B} \neq \{\emptyset\}$ is a *binary clutter* if and only if it is a clutter, and the symmetric difference of an odd number of sets in $\mathcal{B}$ contains a set in $\mathcal{B}$. If $\mathcal{H}$ is a hypergraph, $\mathcal{B}(\mathcal{H})$ denotes the binary clutter *generated by* $\mathcal{H}$, that is the (inclusionwise) minimal elements of the family of symmetric differences (mod 2 sums, more precisely sums over GF(2) of the incidence vectors) of any odd number of members of $\mathcal{H}$, provided $\emptyset$ is not among these symmetric differences, in which case, indeed, the generated hypergraph is a binary clutter.

Define for each input $w = (w_1, \ldots, w_{2n})$ of the PPW(2, 1) problem the collection of intervals

$$\mathcal{I}(w) := \{\{i, i+1, \ldots, j-1\} : 1 \leq i < j \leq 2n,\ w_i = w_j\}.$$

In terms of the paintshop problem one can think of the elements of $\{1, \ldots, 2n-1\}$ as possible moments for color-change: if moment i ($i = 1, \ldots, 2n-1$) is chosen, that means changing the color in our painting machine right after the occurrence of i (and before the occurrence of $i+1$).

Consider now the binary clutter on the set $\{1, \ldots, 2n-1\}$ defined as

$$\text{paint}(w) := \mathcal{B}(\mathcal{I}(w)).$$

We will say that paint(w) is the *paintshop clutter* of w. A paintshop clutter is just a binary clutter generated by intervals. In Section 2, we will see that a solution of the PPW(2, 1) problem – referred to in the sequel as *paintshop solution* – can be fruitfully interpreted as a transversal of the binary clutter paint(w). Moreover, one of our main results – Theorem 1 – is that this binary clutter has a very special structure, which will be exploited to derive algorithms, other properties, etc. We state here a slightly simplified version:

Theorem. *If $\mathcal{A}$ is a paintshop clutter, then there exists a signed graph (G, F) whose odd circuit clutter is isomorphic to $\mathcal{A}$. Moreover, G can be chosen to be a 4-regular graph from which an edge has been removed.*

Let us explain some of these terms. A *signed graph* is a pair (G, F), where $G = (V, E)$ is a graph and $F \subseteq E$. Such a set F is called a *signature*. Denote $\mathcal{C} = \mathcal{C}(G)$ the set of *circuits* of G, that is, of connected subgraphs with all degrees equal to 2. A *cycle* is an edge-disjoint union of circuits. In a signed graph, circuits (cycles) C with $|C \cap F|$ odd are called *odd circuits* (cycles) of (G, F).

The set

$$\mathcal{O}(G, F) := \{C \in \mathcal{C} : |C \cap F| \text{ is odd}\},$$

is a binary clutter and it is called the *odd circuit clutter* of (G, F).

Let us state a typical application of this:

The PPW(2, 1) *problem is polynomially solvable if the graph whose vertices are the letters and the edges are the pairs of crossing letters – that is, letters a, b, which occur in the order a, b, a, b – is bipartite.*

Another result (Corollary 1) improves the quite simple bound n stated in [7]:

Theorem. *A* PPW(2, 1) *problem has a solution σ with at most $p + 3/4(n-p)$ color changes, where p is a lower bound of the optimal solution, and both p and σ can be computed in polynomial time.*

This result occurred to us surprisingly long to prove even if the ideas and the main lines are natural.

Other results and remarks establish min–max theorems in special cases or explore the range of validity of these theorems or of some simple greedy-type algorithms, or note that Conjecture 1 is only a reformulation of Alon, Goldberg and West's well-known result in combinatorial topology [10,2,3].

1.3. Organization of the paper

Section 2 establishes the basic correspondence between paintshop solutions and transversals of odd circuit clutters of graphs, providing also a short introduction to the latter. Section 3 continues by situating paintshop clutters more specifically: first, characterizing them as *odd circuit clutters of particular graphs* (Section 3.1, see a simplified version above). As an immediate corollary we characterize the weak max-flow-min-cut (ideal) property of paintshop clutters (from Guenin's general characterization), and a polynomial algorithm (using the ellipsoid method) for *finding the solution of the paintshop problem in polynomial time whenever the weak max-flow-min-cut property holds* (asked in [5]), and including integral dual solutions (disjoint circuits) for ideal paintshop clutters when the strong max-flow-min-cut property holds. This is stated in Section 3.2 as well as some applications (among them the one mentioned in 1.2). Section 3.3 provides the already cited connection of the paintshop problem (with an arbitrary number of colors) to "combinatorial topology". Section 4 is devoted to the proof of the APX-hardness of PPW(2, 1) and the way of finding a "good" solution in this case, even if an approximation algorithm with fixed ratio is yet unknown. In Section 4.1, we use the polynomial equivalence between the binary paintshop problem and a particular case of BIP, which is shown to be still APX-hard. In lack of readily implementable specializations of MAX-CUT and BIP we first have to prove that these are still APX-hard when restricted to 4-regular graphs, possibly interesting for its own right. (A proof of the APX-hardness of the general paintshop problem has already appeared [5].) In Section 4.2 we show Theorem 6, which gives an upper bound for the minimum transversal of the odd circuits of a graph of maximal degree 4, and which implies Corollary 1 (already stated in 1.2).

2. Preliminaries

2.1. Paintshop clutters as binary clutters

For an introduction to binary clutters we refer to Schrijver [17], p. 1406–1418. These will only be used in some remarks of the sequel: in this paper we can mostly restrict ourselves to graphs and well-known notions (like contraction and deletion of edges).

If $\mathcal{A} \subseteq 2^S$ is a clutter, then its *blocker* is the clutter

$$\mathcal{B} := b(\mathcal{A}) := \{B \subseteq S : B \cap A \neq \emptyset \text{ for all } A \in \mathcal{A}, \text{ and } B \text{ is minimal under this condition}\}.$$

The elements of the blocker are called *transversals.*

Binary clutter and their transversals are crucial for the PPW(2, 1) problem. The paintshop problem PPW(2, 1) consists in *designing a minimum number of color changes so that each hyperedge of the clutter* paint(w) *contains an odd number of them.* That is, we are looking for a minimum cardinality set in the blocker. We state and prove the following simple property though for its crucial importance for the paintshop problem.

Proposition 1. *If $\mathcal{A} \subseteq 2^S$ is a binary clutter and $\mathcal{B} \subseteq 2^S$ is its blocker, then for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$, $|A \cap B|$ is odd.*

Proof. By the minimality of B, for all $e \in B$ there exists $A_e \in \mathcal{A}$ that does not meet $B \setminus \{e\}$, that is, $A_e \cap B = \{e\}$. So indeed, if for a contradiction we suppose that $|A \cap B|$ is even, then the mod 2 sum of A and the sets in $\{A_e : e \in A \cap B\}$, altogether an odd number of sets, contains a set in $\mathcal{A}$ which is not met by B, a contradiction with $B \in b(\mathcal{A})$. □

Proposition 1 proves that minimal transversals of paint(w) are paintshop solutions. Conversely, paintshop solutions are minimal transversals, as realized by Tannier [21]. Take any hyperedge of paint(w) and any paintshop solution: the hyperedge is the symmetric difference of an odd number of intervals, each of them containing an odd number of color changes. Hence, PPW(2, 1) consists in designing a minimum number of color changes so that each hyperedge of the clutter paint(w) contains an odd number of them.

In other words, we do not have to care about the parity, inclusionwise minimal transversals of paint(w) have automatically an odd number of elements in each of its defining intervals, and the converse is also obvious. Any result concerning binary clutters can now be applied to paintshop clutters. For those with the strong max-flow-min-cut property [5,7] have stated Seymour's characterization for paintshop clutters. However, in order to identify subclasses of this problem when a solution can be found with a polynomial algorithm – this is asked for strong max-flow-min-cut clutters in [5] – we need to realize that paintshop clutters belong to a more restricted class of binary clutters, namely the odd circuits clutters of a signed graph (Theorem 1).

2.2. Signed graphs and odd circuits

A *graph* in this paper is a pair (V, E), where V is the vertex-set, E the edge-set and the edges are undirected. Loops and parallel edges are allowed. If it is not said otherwise, $n := |V|$. For $X \subseteq V$, $\delta(X)$ denotes the set of edges with exactly one endpoint in X, and $d(X) := |\delta(X)|$. If X is given with its elements in " " the "()" can be omitted. Whenever we want to specify that the notation concerns the graph G, we write G in the index, like this: "δ_G, d_G".

Let (G, F) be a signed graph. If F' is another signature such that $\mathcal{O}(G, F) = \mathcal{O}(G, F')$, then F and F' are said to be *equivalent*. The symmetric difference of two equivalent signatures meets every circuit in an even number of edges, so it is a cut by introductory graph theory. Conversely, replacing F by $F' := F \triangle D := (F \setminus D) \cup (D \setminus F)$ for some cut D does not change the clutter, that is, $\mathcal{O}(G, F) = \mathcal{O}(G, F')$. This operation will be called *switching*. The operation of *switching on a vertex* $v \in V$ means switching on the star of v. We will extensively use switching for decreasing the cardinality of a transversal, whenever $|F \setminus D| < |D \setminus F|$.

Accordingly, BIP(G, F), where (G, F) is a signed graph will stay for the problem of finding a minimum transversal of the odd circuits of (G, F). Such a transversal is called an *uncut* of the signed graph (G, F). According to Proposition 1, an uncut is always a signature equivalent to F, hence looking for a minimum uncut is equivalent to looking for a minimum signature equivalent to F. (This property will be used extensively in the proof of Theorem 6 already mentioned at the end of Section 1.3). An uncut contains all the odd loops.

If all edges are in the signature, we use the notation BIP(G) for BIP(G, E) and in this special case, the problem of finding the minimum uncut is called the *bipartization problem*, denoted by BIP = BIP(G). The set $F \subseteq E$ is an uncut of G if and only if $E - F$ is a bipartite subgraph, that is, a cut of G. Therefore BIP is polynomially equivalent to MAX-CUT (by complementation, that is, the bijection can be computed in $O(|E|)$ time), that is, to the problem of computing the maximum cardinality of a cut in a given graph.

The problem BIP(G, F) is in fact not more general than BIP(G): indeed, introducing two edges in series for every edge not in F, odd circuits of (G, F) are becoming odd (size) circuits of G.

Let $\mathcal{O}_5 := \mathcal{O}(K_5, E(K_5))$. A third binary clutter will also play a role: F_7, the set of lines in the projective plane over GF(2) (the Fano plane).

The *deletion*, respectively, *contraction* of $s \in S$ in a (binary) clutter $(S, \mathcal{A})$ is the operation resulting in $(S \setminus \{s\}, \mathcal{A} \setminus s)$, $(S \setminus \{s\}, \mathcal{A}/s)$, respectively, where:

$$\mathcal{A} \setminus s := \{A \in \mathcal{A} : A \not\ni s\}$$
$$\mathcal{A}/s := \{A' : A' = A \setminus \{s\}, A \in \mathcal{A}, A' \text{ is minimal among such sets}\}.$$

A clutter obtained after a succession of deletions and contractions is called *minor* of $\mathcal{H}$. It can be readily seen that minors of binary clutters are also binary clutters. The following can also be immediately checked, yet it will be very useful to remember:

Proposition 2 ([9,17]). *Minors of odd circuit clutters of graphs are also odd circuit clutters of graphs.*

Indeed, the deletion of $e \in E$ corresponds to deletion of e in G (and also in F if $e \in F$); for the contraction of e, if $e \notin F$, just contract e, if $e \in F$ switch first on a cut containing e and then contract e. (Contraction of $e = uv (u, v \in V)$ in G means the deletion of e and identification of u and v. If there are edges parallel to e we keep them as loops.) It is easy to see that the resulting odd circuit clutter is the same as the one defined by the clutter operations above. It follows then that graph properties can be translated to clutter properties:

Proposition 3 ([18]). *Odd circuit clutters of graphs do not contain an $b(\mathcal{O}_5)$ or F_7 minor.*

We would like to be aware of a particular contraction: *getting rid* of degree two vertices. Indeed, if $d_G(v) = 2$, and exactly one of the edges incident to v are in F, then one which is not in F can be contracted without changing the set of odd circuits and the number of edges in F. (If both edges incident to v are in F, then we can switch on $\delta(v)$.)

3. Situating the problem

Let us first fix some graph theory terminology. A sequence of vertices is called a *walk* if $xy \in A$ whenever y is a successor of x in the sequence; it is a *closed walk* if the last vertex of the sequence coincides with the first; if every vertex occurs at most once, then we use the term *path* and *circuit*, respectively. The circuit C as a set also means the edge-set $E(C)$ of C, or the vertex set in the appropriate order, depending on the context. In this section we reduce the problem of finding minimum transversals of odd circuit clutters of graphs to the paintshop problem, and we also do the converse transformation. The former will allow us to deduce NP-completeness and more, the latter, polynomial solvability in some cases, and bounds.

3.1. Representation

We first represent paintshop clutters paint(w) as odd circuit clutters of graphs, useful both for pedagogical reasons, and to exploit particular graph properties. As an immediate corollary we get both excluded minors of paintshop clutters; furthermore, these simple (and shortly proved) statements provide compact lemmas for complexity results in both directions.

A reformulation in terms of graphs will allow us to conclude immediately a set of excluded minors for paintshop clutters and ultimately to clarifying the complexity of the paintshop problem. Let $P = (V_n, E_n)$ be the ("Hamiltonian Path") graph defined by $V_n := \{w_1, \dots, w_{2n}\}$ and $E_n := \{e_1 := w_1w_2, e_2 := w_2w_3, \dots, e_{n-1} := w_{2n-1}w_{2n}\}$. Let M be the partition of V_n into the n pairs of occurrences: $w_iw_j \in M$ if $w_i = w_j$. (The pairs participating in M may and also may not be in E_n.) See Fig. 1 for an illustration of this construction and of the Proposition 4.

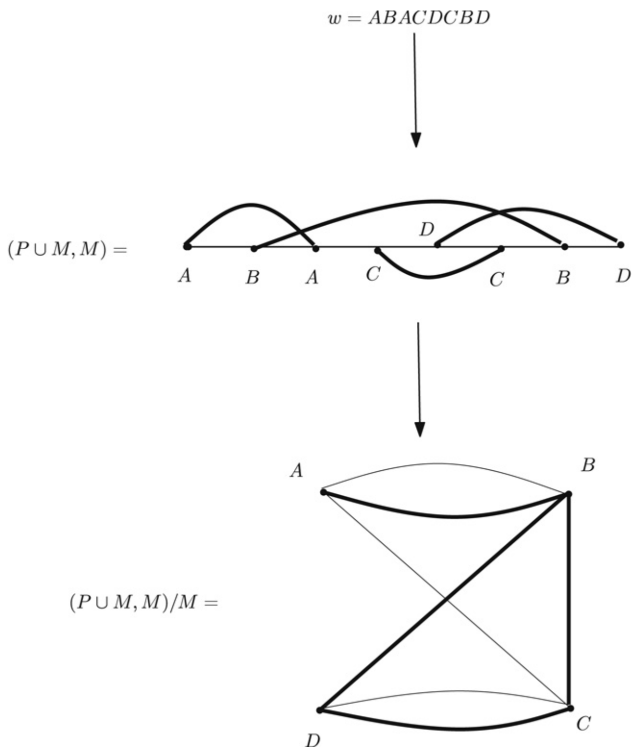

Fig. 1. Illustration of Proposition 4.

Proposition 4. *The mapping $i \mapsto e_i = i(i+1)$ is an isomorphism between* $\text{paint}(w)$ *and* $\mathcal{O}(P \cup M, M)/M$.

Proof. Indeed, the generators of $\text{paint}(w)$ arise by contracting an edge in a circuit containing exactly 1 edge of M. (With an abuse of notation we consider the elements that correspond according to the defined mapping to be the same.) Therefore the generators of $\text{paint}(w)$ are in $\mathcal{O}(P \cup M, M)/M$, so $\text{paint}(w) \subseteq \mathcal{O}(P \cup M, M)/M$. Conversely, every odd circuit of $\mathcal{O}(P \cup M, M)/M$ arises by contracting the M-edges of a circuit C of $\mathcal{O}(P \cup M, M)$. But C is the sum mod 2 of the unique circuits of $P \cup \{e\}$ for all $e \in C \cap M$ (since P is a spanning tree and therefore these unique circuits form a basis of the circuits having a unique combination for each), whence $C \setminus M \in \mathcal{O}(P \cup M, M)/M$ is the mod 2 sum of generators of $\text{paint}(w)$: $\mathcal{O}(P \cup M, M)/M \subseteq \text{paint}(w)$. □

Proposition 5. *Paintshop clutters are odd circuit clutters of signed graphs, and in particular they do not contain $\mathcal{O}_5^*$ and F_7 as minors.*

Proof. The first part is a direct consequence of Proposition 4. Once this is established, the condition of Proposition 3 is verified, and therefore the second part also follows by copying the assertion of Proposition 3. □

A Eulerian Trail is a (closed or open) walk containing every edge of the graph exactly once; there exists such a walk if and only the graph is connected and the number of odd degree vertices is at most 2. It is 0 if and only if the trail is closed, in which case G is called Eulerian. We will call connected graphs with two odd degree vertices *almost Eulerian*.

The *graph contraction of M in $P \cup M$ leads to a graph where all vertices are of degree 4*, except the two extremities which – if they are different – become vertices of degree 3 and the graph becomes almost Eulerian; P becomes an Eulerian trail between the two extremities. If the two extremities are joined by an edge of M, they are contracted to a vertex of degree 2, and the graph becomes Eulerian.

Given an (open or closed) Eulerian Trail $v_1, e_1, v_2, \ldots,$ of a graph $G(V, E)$, $m := |E|$, the *hitches* $C_{v,1}, \ldots, C_{v,\lceil d(v)/2-1\rceil}$ in $v \in V$ (of the given Eulerian Trail) are the subsequences of consecutive edges, starting and ending with the same vertex $v \in V$, and v is not contained any more in the subsequence.

For us the only important case will be the case of connected graphs with two vertices of degree 3 and all other vertices having degree 4, or one vertex of degree 2 and all other of degree 4. We will call such graphs *almost 4-regular.* An Eulerian Trail of such a graph has a hitch C_v for each $v \in V$, a total of n hitches.

In a signed graph (G, F) will say a hitch is *odd* if it contains an odd number of edges of F, otherwise it is *even*.

Theorem 1. *A binary clutter $\mathcal{A}$ is a paintshop clutter if and only if there exists an almost* 4*-regular graph $G = (V, E)$ and $F \subseteq E$ such that (G, F) has an Eulerian trail starting in the minimum degree vertex with only odd hitches, and $\mathcal{A} = \mathcal{O}(G, F)$.*

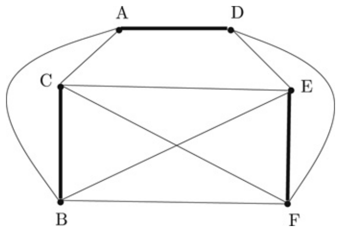

Fig. 2. H_6.

Proof. The "if" direction is easy by letting V to be our alphabet, the list of the letters following the Eulerian trail the given word w. Since $\mathcal{A} = \mathcal{O}(G, F)$, and all hitches are odd, by Proposition 4 we have $\mathcal{A} = \text{paint}(w)$.

To see "only if" direction, suppose first $\mathcal{A} = \text{paint}(w)$, where the length of w is $2n$ and apply Proposition 4 to see that $\mathcal{A}$ is the clutter-minor of the odd circuit clutter of a graph. Then use Proposition 2 to see that $\mathcal{A}$ is an odd circuit clutter in $(P \cup M, M)/M$.

The graph $(P \cup M)/M$ is clearly almost 4-regular, P becomes an Eulerian trail in it, and for all $ij \in M$ the circuits $\{e_i, e_{i+1}, \dots, e_j\} \cup \{ij\}$ have exactly one edge in M, they are therefore odd circuits: $\{e_i, e_{i+1}, \dots, e_j\}$ becomes an odd hitch after (re-signing and) the contraction of ij.

Since there is one hitch for all $e \in M$ and these are all the hitches, we have achieved the proof of the "only if" part. □

The Eulerian trail is provided by the reduction and can of course be determined in polynomial time for each paintshop problem. The following question has no visible practical interest but by the theorem it is polynomially equivalent to the recognition of paintshop clutters, and it might be interesting for its own sake:

Problem 1. Given an almost 4-regular signed graph, does there exist an Eulerian Trail with only odd hitches? Can the answer be "well-characterized"? Is this decision problem polynomially solvable or NP-complete?

Note that in case some cars have to be painted twice with the same color the mathematical model is not really more difficult: then some intervals have to be even, some others odd. Then $\mathcal{O}(P \cup M, M)/M$ has to be replaced by $\mathcal{O}(P \cup M, F)/M$ for some $F \subseteq M$, and Theorem 1 remains the same, except that the clause about the odd hitches can be deleted. Such "generalized" paintshop clutters are exactly the odd circuit clutters of almost 4-regular signed graphs. Some operations on paintshop clutters do lead to such generalized problems.

3.2. Easiness

In this subsection we characterize paintshop clutters for which min–max theorems of different strengths hold, and also those to which particular, naturally arising polynomial algorithms apply.

Let w be an instance of the PPW(2, 1) problem. For convenience, we are considering the binary clutter $\mathcal{O}(P \cup M, M)/M$ isomorphic to paint(w) (see Proposition 4).

A clutter is said to be *packing* if the size of a minimum transversal is equal to the maximum number of pairwise disjoint elements of $\mathcal{A}$. The word w or $\mathcal{O}(P \cup M, M)/M$ will be said to have the *(strong) packing property* if $\mathcal{A}(= \text{paint}(w) = \mathcal{O}(P \cup M, M)/M)$ is packing, and the *weak packing property* if the fractional relaxation of this property holds, that is, if coefficients can be assigned to all $A \in \mathcal{A}$ so that – in the language of $\mathcal{O}(P \cup M, M)/M$ – for all $e \in E(G)$ the sum of the coefficients of $A \in \mathcal{A}$ containing e is at most 1 and the sum of all the coefficients is equal to the minimum transversal. The *strongest packing property* holds for w, if it has the strong packing property, and there exists a maximum set of disjoint elements of $\mathcal{A}$ (=paint$(w) = \mathcal{O}(P \cup M, M)/M$) that are all among the generators (that is, among the initial intervals) of paint(w).

Example. The almost 4-regular signed graph of Fig. 2 is well-known from Seymour's work [19] under the name H_6; it belongs to the word *ABCADECFEBFD*. It is not strongly packing, but it is weakly packing (take the 6 cycles *CBEC*, *EBFE*, *ADECA*, *ADFBA*, *EFDE*, *ACBA*, each of them with a 1/2 coefficient). We introduce the notation $\mathcal{H}_6 := \mathcal{O}(H_6, F)$, where $F = (AD, BC, FE)$.

The size of a minimum transversal of $\mathcal{H}_6$ is 3, but it does not have three disjoint odd circuits. (Otherwise, deleting the 3 circuits there is a fourth one between A and D, which is impossible, see a few lines later.)

The *max-flow-min-cut*, *strong max-flow-min-cut* or *strongest max-flow-min-cut* properties hold for $\mathcal{A} = \text{paint}(w) = \mathcal{O}(P \cup M, M)/M$ if for all $l : E(G) \longrightarrow \mathbb{Z}_+$ the weak packing, strong packing and strongest packing properties hold with the length function l:

The clutter $\mathcal{A}$ has the *strong packing property for l*, if the minimum l-weight of a transversal is equal to the maximum size of a multiset $\mathcal{B}$ of elements of $\mathcal{A}$ such that each $e \in E(G)$ is contained in at most $l(e)$ members of $\mathcal{B}$ (with multiplicities). If the same is required but fractional multiplicities are allowed, we say that $\mathcal{A}$ has the *weak packing property for l*. If the strong packing property holds and in addition the elements of $\mathcal{B}$ are required to be among the intervals generating paint(w), then we say that $\mathcal{A}$ has the *strongest packing property for l*. We also say then that $\mathcal{B}$ is a weak, strong, and strongest l-packing.

A deep result of Guenin [11] characterizes binary clutters with the weak max-flow-min-cut property:

Theorem 2. *The odd circuit clutter of a signed graph has the weak max-flow-min-cut property, if and only if it does not have an $\mathcal{O}_5$ clutter minor.*

The signed graph of Fig. 2 has clearly a $\mathcal{O}_5$ minor (contract edge *AD* after resigning); hence paintshop clutters can indeed contain an $\mathcal{O}_5$ minor. We will see that this clutter does not have the weak max-flow-min-cut property, and it is the (unique, minorwise) minimal such clutter.

For clutters that have the weak or strong max-flow-min-cut property we will need the following algorithms:

Proposition 6. *Given a graph $G = (V, E)$, $F \subseteq E$ and a weight function $l : E \longrightarrow \mathbb{R}_+$, the minimum weight transversal of $\mathcal{O}(G, F)$ or a certificate that $\mathcal{O}(G, F)$ does not have the weak max-flow-min-cut property can be found in polynomial time.*

Proof. We first provide a polynomial algorithm for finding the minimum weight transversal using the ellipsoid method as applied to combinatorial optimization problems in [12]. The separation problem for the polytope

$$P = \{x \in \mathbb{R}^E : x(C) \geq 1 \text{ for every odd circuit } C \text{ of } (G, F),\ x \geq 0\}$$

is exactly the minimum weight odd cardinality circuit problem in graphs: subdivide even edges into two edges having half of the weights each, denote the arising weighted graph by $\hat{G}$.

The solution of this problem with polynomial algorithm is well-known: let V' be a copy of the vertex-set, where the copy of $v \in V$ is $v' \in V'$. For each $uv \in E$ define the two edges uv' and $u'v$, and in weighted graphs define the weight of uv' and $u'v$ be the same as that of uv. Let (G', E') be the defined graph. Clearly, for every odd circuit of G containing $v \in V$, there is a (v, v')-path in G' of the same weight, and conversely, every (v, v')-path corresponds to an odd circuit of the same weight. Thus one can optimize on P in polynomial time by solving a shortest path problem in G'.

Now for finding an obstruction for the max-flow-min-cut property, note that it is sufficient to exhibit a fractional vertex x_0 of P. For certifying that x_0 is a vertex, it is sufficient to exhibit n affinely independent defining inequalities, once *the feasibility of x_0 is established*. However, the latter, that is, deciding $x_0 \in P$ or separating x_0 from P is just the separation problem on P, and we have just noticed that this problem can be solved in polynomial time. □

We saw in Proposition 5 that paintshop clutters are graphic. The following result characterizes those having the weak max-flow-min-cut property:

Proposition 7. *A paintshop clutter has the weak max-flow-min-cut property, if and only if it does not have an $\mathcal{O}_5$ as clutter minor. The optimal solution of such paintshop problems can be found in polynomial time.*

Proof. Suppose $\mathcal{B}$ is a paintshop clutter. By Proposition 5 $\mathcal{B}$ is then the odd circuit clutter *of a graph*, that is, $\mathcal{B} = \mathcal{O}(G, F)$, $(F \subseteq E(G))$. The if part follows from this crucial fact: indeed, if $\mathcal{B}$ does not have the weak max-flow-min-cut property, then by Guenin's Theorem 2 it has an $\mathcal{O}_5$ minor. Conversely, for any binary clutter with an $\mathcal{O}_5$ minor, it is easy to find a weight l [19,17] showing that the weak max-flow-min-cut property does not hold.

A polynomial algorithm (using the ellipsoid method, see [12]) follows from the algorithm stated in the proof of Proposition 6. □

We do not know about a combinatorial algorithm for minimum weight transversals of odd circuits with the weak or strong max-flow-min-cut property. We state the problem for paintshop clutters.

Problem 2. Is there a combinatorial algorithm that solves the paintshop problem for instances that have the weak max-flow-min-cut property, and finds and integer l-packing of sets in a given paintshop clutter if the strong max-flow-min-cut property holds?

The odd circuit clutters of planar graphs are a special case that can arise directly and naturally, we note below that a stronger property holds for them.

Besides Guenin's result, a stronger good characterization theorem (the strong max-flow-min-cut property) holds for a more restricted class of binary clutters, if it does not contain a certain clutter called Q_6 consisting of the stars of the graph K_4 [18], already mentioned in [5]. This means that the minimum transversal of the clutter is equal to the maximum number of disjoint members of the clutter and the appropriate generalization also holds for the weighted generalization of the problem.

Paintshop clutters are not closed under minors, and any graphic binary clutter can be their minor.

There is one more property a paintshop clutter may have, independently of the others, and it is that the greedy algorithm works for it:

Recall that the *greedy algorithm* for a paintshop clutter consists in the coloration of the letters in the given order so as to change the current color only at the second occurrences of letters, and only if necessary.

We finish this section by presenting some examples that separate and enlighten the properties defined in this subsection, before presenting some sufficient conditions for some of them to hold.

Examples. For coloring the word *ABBCCA* the greedy algorithm works, but it does not have the strongest packing property. Indeed, three color changes are sufficient for this word, but the three intervals of paint(*ABBCCA*) are not disjoint. However, it has the strong packing property, since {1, 3, 5} is the mod 2 sum of the three generators of the clutter so it is in paint(*ABBCCA*) and together with {2} and {4} they form a packing.

In the opposite direction, for *ABACDCBD* the greedy algorithm does not find a minimum transversal, but $\{1, 5\}$ is a minimum transversal, and the generators (intervals) that belong to A and C are disjoint, so it has the strongest packing property.

We finally exhibit some (three) simple statements that show how to solve the PPW problem exactly in some nice particular cases. Let us call the set of instances w of PPW$(2, 1)$ *planar*, if the graph $P \cup M$ introduced before Proposition 4 and Theorem 1 is planar. We denote the problem that consists of the planar instances by PPPW. $|w|$ is the length of the word w.

Recall that the vertices of paint(w) are the elements $1, 2, \ldots, |w|-1$ and the (hyper)edges are the symmetric difference of an odd number of intervals $\{i, i+1, \ldots, j-1\}$ with $w_i = w_j$, and this latter set of intervals is denoted by $\mathcal{I}(w)$.

Proposition 8. *The* PPPW *problem is polynomially solvable and the minimum number of color changes for the instance defined by w is equal to the half of the maximum size of a subset $\mathcal{B} \subseteq$ paint(w), where every $i \in \{1, \ldots, |w|-1\}$ is contained in at most two $B \in \mathcal{B}$.*

Proof. Let P be the path and M the matching in the definition of PPW. Since the minors of a planar graph are also planar, the binary clutter we get by contracting M (which means that first we have to switch on one of the endpoints of each arc in M) is planar. Now the statement to prove translates as follows to the obtained planar 4-regular graph : the minimum transversal of the odd circuits of the signed graph $(P \cup M, M)$ is equal to half the maximum size of a set of odd circuits so that every edge is contained in at most two of these. However, this is just a well-known theorem of Lovász ([17], page 489). □

The proof shows that a weaker definition of planarity is also sufficient: we only need that the graph $(P \cup M)/M$ is planar. This definition is strictly more general, as the example *ABCDACBD* shows (contracting the two edges connecting the two occurrences of A and the two occurrences of D leads to a $K_{3,3}$, but contracting also the two occurrences of B and C we get a graph on 4 vertices). However, Proposition 8 has the advantage that it inspires reasonable sufficient conditions that can be simply checked:

Let us recall that two letters a and b are *crossing*, if the order of their two occurrences is a, b, a, b.

Proposition 9. *The* PPW$(2, 1)$ *problem is polynomially solvable and the min–max equality of the preceding theorem holds if the graph of crossing letters (whose vertices are the letters and the edges are the pairs of crossing letters) is bipartite.*

Indeed, in this case one can put the edges in M that belong to one class of the bipartition on one side of P, and the other class on the other side of P, so that none of the edges of M are crossing. Therefore $P \cup M$ is planar.

The following special case is the opposite extreme and has a more practical flavor. We recall that a PPW$(2, 1)$ problem is *fifo* if for any two letters the order of the first occurrences is the same as that of the second occurrences. In other words, in the car manufacturing model, the car that is introduced first is also finished first. For example, the instance *ABACBC* is fifo. In this case the greedy algorithm is optimal:

Proposition 10. *For the subset of fifo instances of* PPW$(2, 1)$ *the greedy algorithm finds the optimum, which is equal to the maximum number of disjoint intervals from among the generators.*

Proof. Let $\mathcal{I}$ be the set of intervals of $\mathcal{I}(w)$ that terminate at color changes. We show that the intervals in $\mathcal{I}$ are pairwise disjoint. If not, then we have two intervals, $I_1 = [a_1, b_1]$, $I_2 = [a_2, b_2]$ which meet, and then, because of the condition, without loss of generality, $a_1 < a_2 < b_1 < b_2$. (Note that in $\mathcal{I}(w)$ the endpoints of the intervals are all different.) Then there is a color change in both b_1 and b_2 and we suppose b_2 is smallest in such a counterexample.

Now by the minimality of b_2 (using again the condition), there is no other color change inside I_2. But then I_2 has exactly one color change besides b_2, a contradiction. □

These two examples show occurrences of the max-flow-min-cut property that have nothing to do with the context of minors.

3.3. Splitting necklaces

Let us imagine an (open) necklace built with n precious stones of t different types and $2a_i$ stones of each type ($n = 2\sum_i a_i$). Now, let us suppose that this necklace has to be divided fairly in two parts (let us say, between two thieves who have stolen it). Fairly means that the each part has the same number of stones of each type. Then we have the following theorem, first proved by Goldberg and West [10].

Theorem 3 (*Necklace Theorem — Weak Version*). *The necklace can be fairly divided in two parts using no more than t cuts.*

Alon and West [2] found a new elegant proof based on the Borsuk-Ulam theorem. Finally, Alon [3] found the following generalization for a necklace with qa_i stones of each type ($n = q\sum_i a_i$):

Theorem 4 (*Necklace Theorem — Strong Version*). *The necklace can be fairly divided in q parts using no more than $t(q-1)$ cuts.*

Theorem 3 is nothing else but Conjecture 1, with $t = |\Sigma|$ and $a_1 = \cdots = a_t = k$, and $|F| = 2$, whereas Theorem 4 is Conjecture 1 for $t = |\Sigma|$, $q = |F|$ and $a_1 = \cdots = a_t = k$. The conjectures were actually also stated in two parts, one containing the other, corresponding exactly to the two theorems.

The question of finding the number of cuts stated in the theorems is an intriguing question from the viewpoint of complexity which is more than open. It is one of several problems of similar nature (SPERNER, KAKUTANI, SECOND HAMILTONIAN CIRCUIT, etc.) and mostly interrelated algorithmically, that serve as illustration for phenomena in the theory of algorithms that do not fit into the usual complexity context (see for instance [4] or [14]): there is no chance to prove their NP-hardness, unless NP = co-NP, since the objects that have to be found always exist.

Even in the easier Theorem 4, and even after some effort, one cannot guarantee a better running time than $O(n^{t-2})$ for $t \geq 3$, and $O(n)$ for $t = 2$ [10].

4. Complexity and bound

4.1. Hardness

In this subsection we prove negative results concerning the complexity and even the approximability of our main results. We are realizing that the NP-hardness and APX-hardness of the paintshop problem have been already proved [5], yet we provide here the simple proofs that come out of our new look on the problem.

We have not found the NP-hardness of MAX-CUT in 4-regular graphs in the literature: in [8] we found it only for 3-regular graphs; the APX-hardness of the problem in general can be found in [16], and for 3-regular graphs in Alimonti, Kann [1]. Our study includes the proof for the 4-regular case, as a necessary step.

We do not provide an introduction to approximability. It is sufficient to know that APX is the set of problems that can be approximated with a constant ratio in polynomial time, and APX-hardness means the nonexistence of polynomial approximation schemes (that is, approximation algorithms with polynomial running time and ratio arbitrary close to 1) unless P = NP. Another definition of an APX-*hard* problem is that a polynomial approximation scheme for such a problem would imply a polynomial approximation scheme for every problem in APX. (The second definition is clearly more restrictive, but there is actually equivalence, even if it is nontrivial, see [15].)

The following so called L-reduction [16,15] of problem A to B is sufficient (but not at all necessary) to conclude polynomial algorithms for constant approximations or approximation schemes for A whenever B has such approximations. Therefore an L-reduction allows to deduce that B is APX-hard, whenever A is APX-hard. It is a refinement of the usual polynomial time reductions that takes into account the approximation ratio - at the same time it is simple and in many cases strong enough to deduce the needed complexity results.

An *L-reduction* [14,15] of an optimization problem A to an optimization problem B, is a pair (R, φ) – where R is a polynomial time reduction from instances of A to instances of B, and $\varphi(x, .)$ is a reduction from feasible solutions to the instance $R(x) \in B$ to feasible solutions to the instance $x \in A$ – with the following two properties satisfied for some $\alpha, \beta \in \mathbb{R}_+$:

$$\text{opt}_B(R(x)) \leq \alpha \text{opt}_A(x),$$
$$|\text{opt}_A(x) - \text{val}_A(x, \varphi(x, s))| \leq \beta |\text{opt}_B(R(x)) - \text{val}_B(R(x), s)|,$$

for all $x \in A$ and feasible solution s for $R(A)$, where $\text{val}_A(x, t)$, $\text{val}_B(y, s)$ denote the values (costs) of the feasible solutions t for instances $x \in A$ or feasible solutions s for instances $y \in B$; $\text{opt}_A(I)$, $\text{opt}_B(R(I))$ is the optimum of instances $I \in A$ or $R(I) \in B$, respectively.

Theorem 5. MAX-CUT *and* BIP *(minimum uncut) restricted to (almost)* 4*-regular graphs, and* PPW(2, 1) *are* NP*-hard; furthermore, all these problems are* APX*-hard.*

Proof. To decide whether a cut of at least a given size, or an uncut of at most a given size exists is clearly in NP. We first present a polynomial reduction of MAX-CUT in 3-regular graphs whose NP-hardness is well known (see [8] noting that vertices of degree 2 can disappear with simple gadgets) to MAX-CUT in 4-regular graphs; then we observe the consequences of the reduction for the complexity and approximability of the problems relevant for us.

Let $G = (V, E)$ be 3-regular, $n := |V|$, let $M(G)$ be the maximum cardinality of a cut in G, and define $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ to be two disjoint copies of G. For $v \in V$, $v_1 \in V_1$ and $v_2 \in V_2$ denote the two corresponding copies of v. For $X \subseteq V, X_i \subseteq V_i, X_i := \{v_i : v \in V\} (i = 1, 2)$. Let finally $T := \{v_1 v_2 : v \in V\}$ be a matching between the corresponding vertices of G_1 and $G_2 (|T| = n)$, and $\widehat{G} := G_1 \cup G_2 \cup T$. Clearly, $\widehat{G}$ is 4-regular.

If $\delta(X), X \subseteq V$ is a cut of G of size d ($d \in \mathbb{N}$), then the cut $\delta(X_1 \cup (V_2 \setminus X_2))$ is a cut of size $2d + n$ of $\widehat{G}$. In particular, if $\delta(X)$ is a maximum cut, then we have a cut of size $2M(G) + n$ in $\widehat{G}$, whence

$$M(\widehat{G}) \geq 2M(G) + n.$$

Conversely, if $\delta(Y)$ with $Y \subseteq V_1 \cup V_2$ is a cut of $\widehat{G}$, then $|\delta_{G_i}(Y \cap V_i)| \leq M(G)$, thus

$$M(\widehat{G}) \leq M(G) + M(G) + |T| = 2M(G) + n,$$

where the equality is satisfied if and only if $\delta_{G_i}(Y \cap V_i)$ is a maximum cut in $G_i (i = 1, 2)$, and T is contained in $\delta(Y)$. Because of the proved opposite inequality we know that the equality is satisfied, and therefore the relation between $M(\widehat{G})$ and $M(G)$ is given by $M(\widehat{G}) = 2M(G) + n$, reducing the MAX-CUT problem of G to the max cut problem of $\widehat{G}$.

We get as a first result that *the* MAX-CUT *problem in a* 4*-regular graph* is NP-hard.

We prove now the APX-hardness of the same problem by refining the above reduction. Let R be the mapping that lets correspond a 3-regular instance to a 4-regular one, like in the first part of the proof, and $\varphi(G, \widehat{C})$ maps for a fixed instance G any cut $\widehat{C} = \delta_{\widehat{G}}(\widehat{X})$ of $\widehat{G}$ to the copy of the biggest of $\delta_{G_1}(\widehat{X} \cap V_1)$ and $\delta_{G_2}(\widehat{X} \cap V_2)$ in G.

Indeed, because of the above proved equality and since a maximum cut contains at least half of the $3n/2$ edges of the 3-regular graph G, that is, $3n/4 \leq M(G)$:

$$M(\widehat{G}) = 2M(G) + n \leq 10M(G)/3 \leq 4M(G),$$

and the difference from the optimum exactly doubles after the reduction, so (R, φ) is an L-reduction with $\alpha = 4$, $\beta = 1/2$.

This reduction immediately extends to the BIP problem in the same graph since complementation is a bijection between the maximum cuts of $\widehat{G}$ and the minimum uncuts in the same graph. Since a maximum cut contains at least half of the edges, the size $m(\widehat{G})$ of the minimum uncut of $\widehat{G}$ is at most half of the edges, and therefore:

$$m(\widehat{G}) \leq M(\widehat{G}) \leq 4M(G)$$

and the distance from the optimum does not change by complementation, so we get again an L-reduction with $\alpha = 4$, $\beta = 1/2$ from MAX CUT in 3-regular graphs.

Last, to extend the proof to PPW(2, 1) problems, we need (see Theorem 1) that the signed graph to which we reduce

– is not 4-regular, but only almost 4-regular, and
– has an Eulerian Trail with only odd hitches.

In order to satisfy the first constraint just delete an arbitrary $e = ab \in T$ in $\widehat{G}$, to get G^{ab}. (The only role of the edges of T in the construction was to increase the degree. This deletion does of course not affect connectivity.) Determine an arbitrary (open) Eulerian Trail (between a and b as endpoints), and define F the signature as the set of all edges of G^{ab}. For each even hitch C do the following: subdivide the last edge f of C into two, f_1 and f_2, with a new vertex, and add a loop; delete f from F and put the loop and f_2 in it. If C is not the last hitch, do the same with the edge g that immediately follows this hitch in the Eulerian Trail (and, as for f, replace g by the loop and g_2 in F). We extend the Eulerian Trail with the two loops (replacing f by f_1, the loop and f_2 in this order, and similarly for g). Every hitch except C contains either both the added loops or neither, and C contains exactly one of them. So the number of even hitches in the fixed Eulerian Trail of G^{ab} decreased by 1 (and the number of odd ones increased by 3). Eventually all even hitches of the Eulerian Trail disappear.

Let k be the number of loops we added to G^{ab}, and denote the resulting graph by $\widetilde{G}$.

Now every hitch is odd, and the minimum transversal of the odd circuits has k more edges than $\widehat{G}$. Since $k \leq 2n \leq 8M(G)/3$, the size $m(\widetilde{G}, F)$ of the minimum transversal of the odd circuits of $(\widetilde{G}, F)$ satisfies (using also the previous series of inequalities) :

$$m(\widetilde{G}, F) = m(\widehat{G}) + k \leq 4M(G) + 8M(G)/3 \leq 7M(G),$$

and the distance from the respective optima does not change by complementation, so we get again an L-reduction with $\alpha = 7$, $\beta = 1/2$ from MAX-CUT in 3-regular graphs. Yet this latter problem is APX-hard according to [1]. The statement for PPW(2, 1) follows from Proposition 4 and the reduction of BIP for signed graphs to BIP. □

4.2. Bound

There is no constant approximation guarantee known for PPW(2, 1) and we can also not prove any such guarantee. However, we can guarantee a solution below a certain bound which may be better if the optimum is large and can be far from the optimum when the optimum is very small (but then we may care less).

In a fixed BIP problem on a signed graph instance (G, F), where G is of maximum degree 4 different from K_5, let us denote by p_0 the total number of odd loops (loops $e \in F$) plus the number of edge-disjoint odd circuits of length 2. These will be called *obliged* loops or parallel classes. No decision is related to these: a transversal has to contain at least p_0 of the edges of these odd circuits, and this lower bound can be clearly reached.

By deleting odd loops and the edge-disjoint odd circuits of length 2, we get a polynomially equivalent minimization problem. The new graph will have degree 1 or degree 2 vertices. By contracting edges, one can get rid of these. Repeating this procedure as many times as possible one eventually gets an optimization problem equivalent to the original one on a simple graph. We denote by p the total number of these obliged edges, which are deleted from the graph. Note that through the process the number of vertices decreases by p.

For the odd circuits of K_5 the ratio transversal/vertex is 4/5 (see Fig. 3). This is the only exception for the upper bound 3/4 (among simple graphs) according to the following theorem:

Theorem 6. *An odd circuit clutter $\mathcal{O}(G, F)$, where $G = (V, E)$ is a connected simple graph on n vertices with maximum degree 4, has a transversal B with $|B| \leq \frac{3}{4}n$ edges unless $(G, F) = \mathcal{O}_5$. Such a B can be found in polynomial time.*

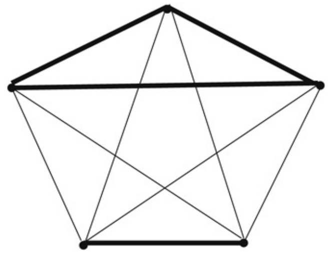

Fig. 3. $|B|/|E| = 4/10$; $|B|/n = 4/5$.

The most relevant fact is the following: in a 4-regular (or almost 4-regular) graph without loops or parallel edges, at least 5/8 of the edges can be eliminated as not participating in a transversal, in particular in an uncut.

In the proof, we use extensively the fact that an uncut is always a signature equivalent to F, as noted in Section 2.

Proof. Suppose for a contradiction that (G, F) is a counterexample with a minimum number of vertices. Let B be a minimum cardinality transversal of the odd circuits of (G, F). The proof will show that even in this case one has $|B| \leq \frac{3}{4}n$.

Claim 1. There are at most 2 edges of B incident to every vertex, and if the equality holds for $v \in V$, then $d_G(v) = 4$.

Otherwise switching on v we could decrease B, contradicting the minimality of B.

Let $D_i := \{v \in V : d_B(v) = i\}$ $(i = 0, 1, 2)$. Note that $2|B| = 2|D_2| + |D_1|$.

Claim 2. D_2 is a stable set in the graph $(V, E \setminus B)$.

Indeed, if $u, v \in D_2$, $uv \in E \setminus B$, then $d_G(u, v) \leq 6$ and 4 of these (at most) 6 edges of a cut are in B, contradicting the minimality of B.

Claim 3. Two vertices in D_2 nonadjacent in B have no common $E \setminus B$ neighbor in D_1.

Indeed, the cut $\delta(a, b, x)$ defined by two vertices $a, b \in D_2$ and a vertex $x \in D_1$ adjacent to both a and b is of size at most 8. On the other hand $\delta(a, b, x)$ contains at least 5 edges of B, contradicting the choice of B.

Claim 4. Three vertices in D_2 that are pairwise nonadjacent in B have no common neighbor in D_0.

Indeed, if $a, b, c \in D_2$ were such vertices, and x their common neighbor in D_0, then $|\delta(a, b, c, x)| \leq 10$ and 6 of these (at most) 10 edges are in B.

Claim 5. If three different vertices $a, b, c \in D_2$ have a common $E \setminus B$ neighbor in D_1, then $ab, bc, ca \in B$.

Indeed, if not, then say $ab \notin B$, and a, b have a common neighbor in D_1, contradicting Claim 3.

Claim 6. There are no three different vertices $a, b, c \in D_2$ with $ab, bc, ca \in B$ that have two common $E \setminus B$ neighbors.

Indeed, for a contradiction suppose d and e are such common neighbors. If $de \in E$, then $\{a, b, c, d, e\}$ induce a K_5, and since the maximum degree in G is 4, and G is connected, $G = K_5$. So $de \in B$ implies $\mathcal{O}(G, B) = \mathcal{O}(G, F) = \mathcal{O}_5$, which is not possible; $de \notin B$ implies that at most 3/10 of the edges are in B, and in this case we are done.

So $de \notin E$. If G is not 2-connected we can proceed by 2-connected blocks. So we can suppose without loss of generality that G is 2-connected. Then it has a cut of two edges dd' and ee', we distinguish two cases depending on whether exactly one or none of these edges are in B, supposing without loss of generality $ee' \notin B$ (Fig. 4). (Both cannot be in B because then it would contradict the minimality of B.) We replace this part of the graph by a path of length 2: d', f, e' – with $d'f$ in B' if $dd' \in B$ and not in B' if $dd' \notin B$ – the other edges of B' are exactly those of B not incident to a, b, c, d, e. It is easy to see that a signature equivalent to B' with fewer edges will lead also to a signature equivalent to B with fewer edges.

By the minimality of G we know the statement for the reduced graph: $|B'| \leq 3/4(n-4)$. The original graph has 4 more vertices, and 3 more edges in B, so the bound also holds for G: $|B| \leq 3/4(n-4) + 3 = 3/4n$. A contradiction with the fact that G is a counterexample, finishing the proof of Claim 6.

We will call a triangle $\{a, b, c\} \subseteq V$, $\{ab, bc, ca\} \subseteq B$ *bad* if a, b, c have exactly one common neighbor x, and $x \in D_1$.

Claim 7. One has $|D_2| \leq 2|D_0| + |D_1|$.

Consider the bipartite graph $H = (U, W, T)$, $U := D_2$, $W := D_0 \cup D_1$, where T is the set of edges in $E \setminus B$ with one endpoint in U, the other in W.

Because of Claim 1 and Claim 2, all vertices of D_2 are of degree exactly 4, whence the number of $E \setminus B$ edges incident to D_2 is exactly $2|D_2| = |T|$. We prove now that $|T| \leq 4|D_0| + 2|D_1|$ by counting the number of edges incident to vertices of W.

If every vertex in $D_1 \subseteq W$ has at most 2 neighbors in H, we have:

$$(*) \qquad 2|D_2| = |T| \leq \sum_{v \in D_0} d_H(v) + \sum_{v \in D_1} d_H(v) \leq \sum_{v \in D_0} 4 + \sum_{v \in D_1} 2 = 4|D_0| + 2|D_1|,$$

and then the claim is proved. Unfortunately, because of the exceptions the proof of Claim 7 will still be long:

Otherwise, there is a vertex in $x \in D_1$ which has at least three neighbors in D_2, and then the estimate in (*) for x is not correct: $d_H(x) = 3 > 2$. (If $x \in D_0$, it is correct.) Then by Claim 5 and Claim 6 the three neighbors of x form a bad triangle $\{a, b, c\} \subseteq V(G)$, $ab, bc, ca \in B$, with x as only common neighbor.

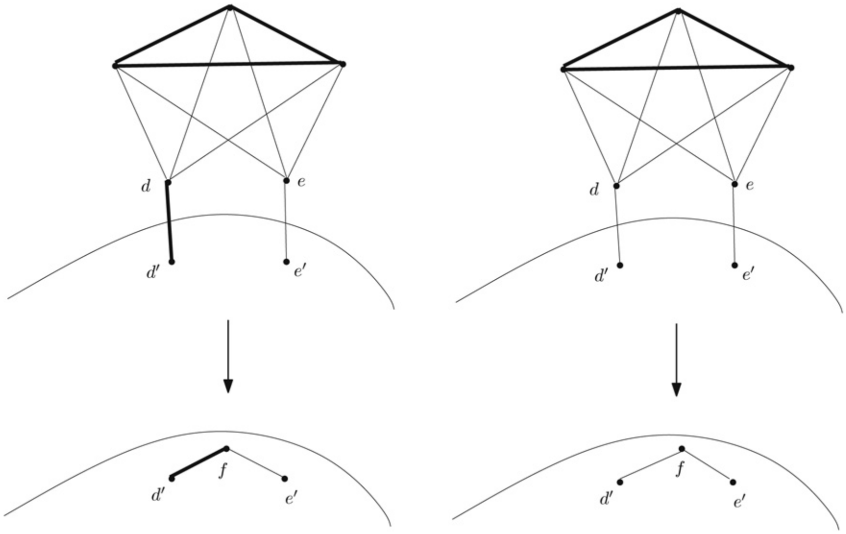

Fig. 4. The reduction used in the proof of Claim 6.

Now, because of Claim 1, a and b and c have each, exactly one neighbor in W besides x, two (but not three because of Claim 6) of which may coincide. So there exists $y \in W$ which is adjacent to exactly one of a, b, c, say a.

For each triangle $\{a, b, c\}$ of B for which there is exactly one $x \in V$ adjacent to a, b and c (independently of whether $x \in D_1$ or $x \in D_0$) we design a *support* which is one of the vertices y incident to exactly one of a, b and c. (Again, because of Claim 6, at least one support exists; if there are several, we choose arbitrary one of them.) We say then that x is the *flag* (of the triangle or of its support). Because of Claim 2 the vertices of different triangles of B are nonadjacent in $E \setminus B$, and they are obviously nonadjacent in B. By Claim 3 and Claim 4, y can be the support of at most two different triangles, and in case of equality $y \in D_0$.

We will call the flag and the support of a bad triangle a *bad* flag or support. By definition a bad flag x of a triangle is in D_1 and adjacent to the three vertices of the triangle and has of course no other neighbor in H. In the following we make it apparent that (*) still holds, because the incorrect bounds are compensated by some strict inequalities. Each flag is included in a group of two or three vertices for which the sum of the degrees still satisfies the bounds.

We proved that a vertex x for which the bound does not hold is a bad flag. Choosing the notation appropriately the support y of x is adjacent to a vertex denoted a of the triangle and nonadjacent to the other two vertices b and c.

Let $D := (W', A)$ be the following directed graph: W' is the set of the flags and supports of all triangles of B, and $xy \in A$ if there exists a triangle so that x is the flag of the triangle and y is its support.

The *weak components* of a digraph are the components of the underlying undirected graph (that is, of the undirected graph whose edges are the arcs of the digraph without their orientation). (We use the terminology of [17].) We show now:

The weak components of D are paths and circuits. The circuits are directed circuits of D entirely contained in D_0. The edges of each path-component P are oriented so that it is the union of two directed paths, one from the endpoint $x_1 = x_1(P)$ to $y = y(P)$, the other from the other endpoint $x_2 = x_2(P)$ to the same y. Furthermore, $P \setminus \{x_1, x_2\} \subseteq D_0$.

It is possible to have $y = x_1$ or $y = x_2$, for instance in the latter case P is a path directed from x_1 to x_2 (see Fig. 5). The vertex y will be called the *sink* of the component P.

Indeed, by Claim 1 the triangles are disjoint and since all degrees are at most 4 a vertex cannot be the flag of two triangles, so the outdegrees are at most 1. We show now that every $v \in V(D)$ is incident to at most two arcs (that is, the undirected degrees are at most 2):

If $y \in V$ is the support of two triangles, then by Claim 2 and Claim 3 it is in D_0; it cannot be the support of three triangles by Claim 4. If it was the flag of a third triangle it would have $1 + 1 + 3 = 5$ neighbors in G.

So in D no more than one arc can leave a vertex, no more than two arcs can enter a vertex, and if there are two entering arcs, then there is no leaving arc, finishing the proof of the assertion in italic.

Now we define for *every bad flag* $x_1 \in D_1$ a group of 2 or 3 vertices of $V(D)$ consisting of x_1, of the sink y of its component P and of the other endpoint x_2 of P, if $x_2 \neq y$. Such a definition is possible since if $x \in D_1$, it has no other neighbor in H and hence it is not the support of any triangle. We finish the proof by checking that our estimations in (*) are still valid for each such group.

Suppose x_1 is a bad flag and y is the sink of its component denoted by P. Recall that $d_H(x_1) = 3$.

The in-neighbor of y in D (on the path between x_1 and y) is the flag of a triangle, and y is adjacent to exactly one vertex of that triangle that we denote by a.

Case 1. $y \in D_1$.

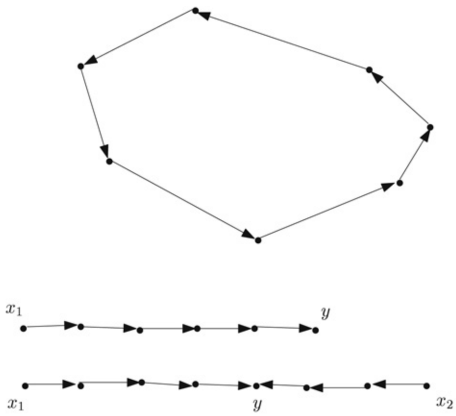

Fig. 5. The three kinds of weak components of D.

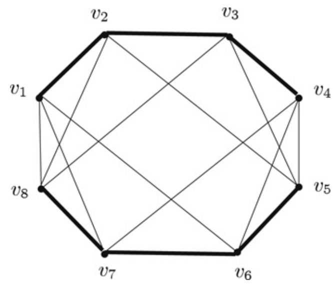

Fig. 6. The implicit algorithm provided by the proof of Theorem 6 can not improve the bound of this example, 6 edges $= (3/4)8$, although the best uncut has 4 edges (v_1v_7, v_2v_8, v_3v_5, v_4v_6).

Then by Claim 3 the only neighbor of y in H is a, in particular the group of x_1 is $\{x_1, y\}$, $d_H(y) = 1$, so our estimation

$$d_H(x_1) + d_H(y) = 3 + 1 = 2 + 2$$

in (*) for this group is correct.

Case 2. $y \in D_0$.

Case 2a: The group of x_1 is of size 2: $\{x_1, y\}$.

Then we show $d_H(y) \leq 3$: For a contradiction, suppose y is adjacent to 3 vertices of D_2 besides a. If 2 of these vertices were nonadjacent, together with a they would contradict Claim 4. So they form a triangle. Because of Claim 6 (as it has already occurred), this triangle has a support contradicting that y is a sink.

Therefore

$$d_H(x_1) + d_H(y) \leq 3 + 3 = 2 + 4,$$

and the bound of (*) is valid for the sum $d_H(x_1) + d_H(y)$.

Case 2b: The group of x_1 is of size 3: $\{x_1, y, x_2\}$.

Then y has two in-neighbors (both on the undirected path P) corresponding to neighbors a and a' in two different triangles whose common support is y. Then by Claim 4 y can no more have other neighbors in $D_2 = U$, so $d_H(y) = 2$. We have then

$$d_H(x_1) + d_H(y) + d_H(x_2) \leq 3 + 2 + 3 = 2 + 4 + 2$$

provided $x_2 \in D_1$, since then x_2 is a bad flag and therefore $d_H(x_2) = 3$;

$$d_H(x_1) + d_H(y) + d_H(x_2) \leq 3 + 2 + 4 < 2 + 4 + 4,$$

provided $x_2 \in D_0$. In both cases (*) is correct, finishing the proof for Case 2b, and at the same time of Claim 7.

Hence, we know that if G is a counterexample with a minimum number of vertices, one has $|D_2| \leq 2|D_0| + |D_1|$. Expressing D_0 in the relation $|D_0| + |D_1| + |D_2| = n$, one gets $|D_1| + 3|D_2| \leq 2n$. Summing up this last inequality and $|D_1| + |D_2| \leq n$ and dividing by 4 we get $|B| = (2|D_2| + |D_1|)/2 \leq \frac{3}{4}n$. □

In some cases, the implicit algorithm provided by the proof of Theorem 6 cannot find the best uncut, of course, since the problem of finding the best uncut is NP-hard (see Fig. 6).

For the paintshop problem, we will say that the p defined at the beginning of the present subsection is the number of "obliged" color changes. Note that p is always a lower bound of the optimum and can be computed by a greedy algorithm.

Corollary 1. *A* PPW(2, 1) *problem has a solution with at most* $p + 3/4(n - p)$ *color changes, where p is the number of "obliged" color changes, and can be found in polynomial time.*

Note that this improves the straightforward bound n proposed in [7] ($1/4p + 3/4n \leq n$, with equality only if $p = n$, that is only if all color changes are "obliged"), and if p tends to infinity the optimum is asymptotically equal to p.

Proof. Construct the associated almost 4-regular graph when all obliged edges are deleted, note that the constructed binary clutter cannot be $\mathcal{O}_5$ and apply Theorem 6. □

5. Conclusion

The main messages of this article are the following:

- PPW(2, 1) is polynomially equivalent to the minimum transversal problem for odd circuits in graphs, that are furthermore almost 4-regular. The reduction of the latter problem (which is polynomially equivalent to minimum uncut, thus the exact solution to max cut as well) to PPW(2, 1) establishes the NP-hardness of PPW(2, 1) in a natural way; the reduction of PPW(2, 1) to the minimum blocker of odd circuits establishes the polynomial solvability of PPW(2, 1) in some special cases. In some of these cases the greedy algorithm already solves the problem, in some others a stronger min–max relation holds.
- The conjectures stated in [7] are consequences of the Necklace theorem. This relation puts the complexity of finding a solution for regular instances in complexity classes different from the usual NP-hard or polynomial solvable problems. For PPW(2, k) there are constructions (see [13,20]) but not likely to be efficient, even though the corresponding decision problem is solvable. For PPW(c, k), $c > 2$ there is no constructive proof, even if of course a complete search of all possible cutting points leads always to one of the existing solutions.
- Among positive results concerning the minimum transversals of these binary clutters or the minimum uncut problem, the most difficult result states that the minima never exceed the 3/8th of the total number of edges (ignoring those that are forced to be in the transversal or minimum uncut for some obvious reason), and a solution satisfying this bound can be found in polynomial time.

Acknowledgments

We thank W. Hochstättler, an anonymous but highly efficient referee and Endre Boros for the quick and very thorough refereeing procedure. The second author's research was partially supported by the Marie Curie network "ADONET" of the European Community.

References

[1] P. Alimonti, V. Kann, Hardness of approximating problems on cubic graphs, in: CIAC97, in: LNCS, vol. 1203, 1997, pp. 288–298.
[2] N. Alon, D. West, The Borsuk–Ulam Theorem and the bisection of necklaces, Proceedings of the American Mathematical Society 98 (1986) 623–628.
[3] N. Alon, Splitting necklaces, Advances in Mathematics 63 (1987) 247–253.
[4] N. Alon, Non-constructive proofs in Combinatorics, in: Proc. of the International Congress of Mathematicians, Kyoto 1990, Japan, Springer-Verlag, Tokyo, 1991, pp. 1421–1429.
[5] P. Bonsma, Th. Epping, W. Hochstättler, Complexity results for restricted instances of a paint shop problem for words, Discrete Applied Mathematics 154 (2006) 1335–1343.
[6] T. Epping, W. Hochstättler, M. Lübbecke, MaxFlow-MinCut Duality for a Paint Shop Problem, Operations Research Proceedings, Klagenfurt, September 2002.
[7] T. Epping, W. Hochstättler, P. Oertel, Complexity results on a paint shop problem, Discrete Applied Mathematics 136 (2004) 217–226.
[8] M. Garey, D. Johnson, Computers and Intractability, Bell Telephone Laboratories, Incorporated, 1979.
[9] A.M.H. Gerards, Graphs and polyhedra — Binary spaces and cutting planes. CWI Tracts 73, 1990, CWI, Amsterdam, Ph.D. Thesis.
[10] C. Goldberg, D. West, Bisection of circle colorings, SIAM Journal of Algebraic and Discrete Mathematics 6 (1985) 93–106.
[11] B. Guenin, A characterization of weakly bipartite graphs, Journal of Combinatorial Theory B 83 (2001) 112–168.
[12] M. Grötschel, L. Lovász, A. Schrijver, Geometric Algorithms and Combinatorial Optimization, Springer Verlag, Berlin, Heidelberg, 1988.
[13] F. Meunier, Discrete Splittings of the Necklace, Mathematics of Operations Research (in press).
[14] C. Papadimitriou, On the complexity of the parity argument and other inefficient proofs of existence, Journal of Computer and System Sciences 48 (1994) 498–532.
[15] C. Papadimitriou, Computational Complexity, Addison Wesley Inc., 1994.
[16] C. Papadimitriou, M. Yannakakis, Optimization, approximation and complexity classes, Journal of Computer and System Sciences 43 (1991) 425–440.
[17] A. Schrijver, Combinatorial Optimization, vol. 1–3, Springer-Verlag, Berlin, Heidelberg, New York, 2003.
[18] P.D. Seymour, The matroids with the max-flow min-cut property, Journal of Combinatorial Theory B 23 (1977) 189–222.
[19] P.D. Seymour, Matroids and multicommodity flows, European Journal of Combinatorics 2 (1981) 257–290.
[20] F. Simmons, F. Su, Consensus-halving via theorems of Borsuk–Ulam and tucker, Mathematical Social Sciences 45 (2003) 15–25.
[21] E. Tannier, personnal communication, 2002.