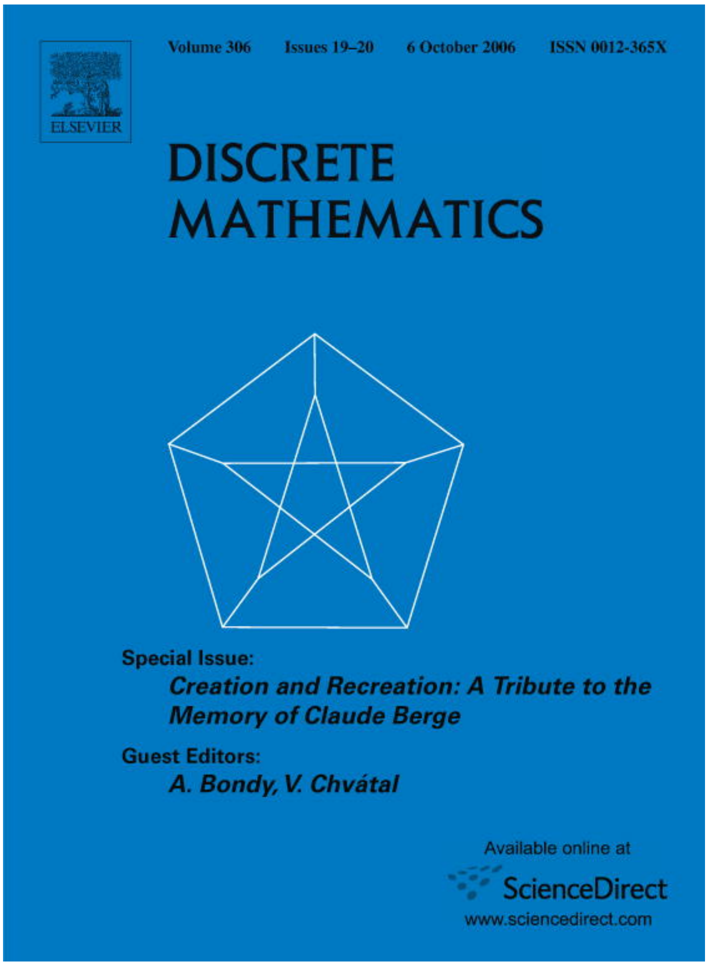
Volume 306 Issues 19–20 6 October 2006 ISSN 0012-365X
ELSEVIER
DISCRETE MATHEMATICS
Special Issue:
Creation and Recreation: A Tribute to the Memory of Claude Berge
Guest Editors:
A. Bondy, V. Chvátal
Available online at
ScienceDirect
www.sciencedirect.com

Available online at www.sciencedirect.com
ScienceDirect
Discrete Mathematics 306 (2006) 2582–2592
DISCRETE MATHEMATICS
www.elsevier.com/locate/disc

A Berge-keeping operation for graphs☆

András Sebő
CNRS, Leibniz-IMAG, Grenoble, France

Received 3 March 2003; received in revised form 1 August 2005; accepted 12 December 2005
Available online 20 July 2006

Abstract

We prove that a certain simple operation does not create odd holes or odd antiholes in a graph unless there are already some. In order to apply it, we need a vertex whose neighborhood has a coloring where the union of any two color classes is a connected graph; the operation is the shrinking of each of the color classes. Odd holes and antiholes do have such a vertex, and this property of minimal imperfect graphs implies the strong perfect graph theorem through the results of the paper. Conceivably, this property may be a target in the search for a proof of the strong perfect graph theorem different from the monumental achievement of Chudnovsky, Robertson, Seymour, and Thomas.

MSC: 05C17; 05C15; 05C69

Keywords: Coloration; Property-keeping operation (minor); Perfect graphs; Berge-graphs; Odd hole; Odd antihole

1. Introduction, definitions, results

Does anyone still remember *monsters*? According to Vašek Chvátal it was the nickname given by Pierre Duchet to minimal counterexamples to Claude Berge's strong perfect graph conjecture.[1] Study of monsters used to be a line of attack on the conjecture: the hope was that by establishing more and more properties of monsters, one would eventually arrive at a list so restrictive that it would demonstrate their nonexistence. Another line of attack was following the hope that every Berge graph either has a relatively transparent reason for perfection or there is a property—preferably one that allows recursion with algorithmic advantages, for instance decomposition—absent from all minimal imperfect graphs, or at least from those of minimum cardinality (or smallest in some fixed good ordering). The proof of the conjecture by Chudnovsky et al. [3] came from the latter direction.

The fact that this proof is long and difficult stimulates alternative proofs or simplifications. In particular, even though we now know that there are no monsters, interest in arguments independent of [3] that establish some relevant properties of monsters persists. In the present paper, we show by such an argument that in a smallest monster G,

every vertex has degree at least $2\omega(G) - 1$,

☆ Research partially supported by the "ADONET" Marie Curie training network of the European Community.
E-mail address: Andras.Sebo@imag.fr.

[1] Through a sequence mutations the preliminary title of my talk at the Princeton meeting on perfect graphs in September 2001 was "Berge Monsters". This is nonsense: by definition, every Monster is Berge. Discussing this with Claude we finally arrived at a related nonmathematical but at least true statement: no Berge is a monster.

0012-365X/$ - see front matter
doi:10.1016/j.disc.2005.12.041

immediately implying the strong perfect graph theorem (SPGT) for claw-free graphs [12]. In the direction of simplifying the other line of attack, any property of Berge graphs that does not follow trivially from the nonexistence of odd holes or odd antiholes as *induced subgraphs* may help. We show

an operation that does not spoil the Berge property,

that decreases the size of the graph, that changes the clique number in a way closely related to antiholes, that does not change the stability number, and that may turn out to be more efficient than taking induced subgraphs.

Graphs in this paper will always be undirected. Parallel edges and loops are allowed, but they are irrelevant to our inquiry, and so they may just as well be deleted. The set of all vertices of G will be denoted by $V(G)$ and the set of all edges of G will be denoted by $E(G)$. A *subgraph* arises by deleting edges and vertices from a graph; it is not necessarily an induced subgraph. The subgraph of G induced by a subset X of $V(G)$ will be denoted by $G(X)$.

A *chord* of a circuit or a path is an edge that joins two vertices at a distance greater than one on the circuit or the path. A *hole* in a graph is its induced subgraph on at least five vertices which is a chordless circuit. An *antihole* in G is a hole in $\overline{G}$. A hole on k vertices will be denoted by C_k. The *length* of a path is the number of its edges. A chordless path on m vertices (whose length is $m-1$) will be denoted by P_m.

A *stable* set is a set of vertices without induced edges and a *clique* is a stable set of the complementary graph $\overline{G}$. The *stability number* $\alpha(G)$ is the largest cardinality of a stable set in G, the *clique number* $\omega(G)$ of G is the largest cardinality of a clique in G, and the *chromatic number* $\chi(G)$ is the smallest cardinality of a partition of $V(G)$ into stable sets (which are called *color classes*). Every G satisfies $\chi(G) \geqslant \omega(G)$; a graph is called *perfect* if $\chi(H)=\omega(H)$ for every induced subgraph H of G. Berge's *strong perfect graph conjecture* states that a graph is perfect if and only if it contains neither an odd hole nor an odd antihole; following Chvátal and Sbihi [5], a graph is called a *Berge graph* if it contains neither an odd hole nor an odd antihole.

If $xy \in E(G)$, then we say that x and y are *adjacent*. The *neighborhood* of a vertex v in a graph G is defined as the set of all vertices of G adjacent to v; it is denoted by $N_G(v)$ or simply by $N(v)$. The *closed neighborhood* of v is defined as the union of $N(v)$ and $\{v\}$; it is denoted by $N[v]$. We write $\overline{N}(v)$ for $V(G)\setminus N[v]$.

A coloring of a graph G by $\omega(G)$ colors will be called *optimal*; G is called *uniquely colorable* if $\chi(G)=\omega(G)$ and G has a unique optimal coloring. A coloring of G will be called *connected* if the union of any two color classes induces a connected graph. Clearly, the optimal coloring of a uniquely colorable graph is connected. (This is not reversible though, even not for perfect graphs: $\overline{C}_6$ has two distinct optimal colorings and both of them are connected. Also, a perfect graph can have a connected coloring that is not optimal: the line graph of $K_{3,4}$ has a natural partition into 3 cliques and also into 4 cliques; the former is an optimal coloration, but the latter is also connected. Note also that a graph does not necessarily have a connected coloring even if it is perfect: for instance $\overline{P}_5$ does not have any.) We will say that a *subset X of $V(G)$ is uniquely colorable* if $G(X)$ is uniquely colorable.

The *shrinking* of a subset T of $V(G)$ means the replacement of T by just one vertex, which is then joined to all vertices adjacent to at least one vertex in T. We consider the edge-set of the new graph to be the same as that of the original graph, except that some edges become loops and the parallel edges may appear; these are irrelevant and may be immediately deleted. When G is a graph, v is a vertex of G, and f is a coloring of $G(N(v))$, we let $G\neg(v,f)$ denote the graph arising from G by shrinking the color classes of f one by one and by deleting v (along with all edges adjacent to v); when we write $G\neg(v,f)$ we always suppose that $v\in V(G)$ and that f is a connected coloration of $N(v)$; when $G(N(v))$ is uniquely colorable, we write simply $G\neg v$ for $G\neg(v,f)$ with f the optimal coloring of $G(N(v))$. If G is the odd antihole $\overline{C}_{2\omega+1}$ and if v is any vertex of G, then $G\neg v = K_{\omega+1}$ ($\omega\in\mathbb{N}$).

In this paper we show the following results:

Lemma 1.1. *If G is a graph, $\omega(G)<\chi(G)$, $v\in V(G)$ and f is a coloring of $N(v)$ with $\omega(G)-1$ colors such that $\omega(G\neg(v,f))\leqslant\omega(G)$, then $G\neg(v,f)$ is also imperfect.*

Theorem 1.1. *If G is a minimal imperfect graph different from an odd antihole and v is a vertex of degree $2\omega(G)-2$ in G, then the subgraph of G induced by $N_G(v)$ is uniquely colorable and $G\neg v$ is imperfect.*

Theorem 1.2. *If G is a Berge graph, v is a vertex of G, and f is a connected coloring of $N_G(v)$, then $G\neg(v,f)$ is also Berge.*

It is easy to see by examples that this theorem is no more valid shrinking only one of the color classes (see Section 4), or if the operation is defined on a subset of the neighborhood.

It is well known [10,9,13,15] that every vertex of every minimal imperfect graph G has degree at least $2\omega(G)-2$; the conjunction of our two theorems guarantees that the smallest monster G cannot have a vertex of degree $2\omega(G)-2$. In other words, *the SPGT is equivalent to the existence of a vertex of degree* $2\omega-2$ *in every minimal imperfect graph G or its complement.* Comparing the Lemma and Theorem 1.2 one gets the following corollary that will turn out to be sharper (since the conditions of the lemma can be easily derived from those of Theorem 1.1, see the proof of the latter):

The SPGT can be derived from the following:

> In every minimal imperfect graph G other than an odd antihole, there are a vertex v and an optimal coloring f of $N(v)$ such that $\omega(G\neg(v, f))\leqslant\omega(G)$.

2. Proof of Theorem 1.1

Let G be an imperfect graph, $v \in V(G)$, and let us write $\alpha=\alpha(G)$ and $\omega=\omega(G)$. □

Proof of the Lemma. Suppose for a contradiction that $G\neg(v, f)$ is perfect. Then by the condition $\omega(G\neg(v, f))\leqslant \omega(G)$ it has an optimal coloration with $\omega(G)$ colors. Since f is an optimal coloration of $N(v)$, it has $\omega-1$ classes becoming vertices of $G\neg(v, f)$; therefore, there is a color class S in this optimal coloration of $G\neg(v, f)$ that does not intersect $N(v)$. But then $S \cup \{v\}$ together with the other $\omega-1$ color classes of $G\neg(v, f)$ provides a coloration of G with ω colors, contradicting the condition $\omega(G)<\chi(G)$. □

Proof of Theorem 1.1. Let now G be a minimal imperfect graph different from an odd antihole and v of degree $2\omega-2$ in G and let N denote the subgraph of G induced by $N_G(v)$.

Since G is minimal imperfect, $\chi(N)=\omega(N)=\omega-1$.

Claim 1. *N is uniquely colorable.*

Let us assume the contrary: N admits two distinct $\omega-1$-colorings, one with color classes S_i $(i=1,2,\ldots,\omega-1)$ and the other with color classes T_i $(i=1,2,\ldots,\omega-1)$; since the two colorings are distinct, we may assume that T_1 is distinct from all S_i.

Consider now the incidence vectors of the sets $S_i \subset N$ $(i=1,2,\ldots,\omega-1)$ and $T_1 \subset N$ in $\mathbb{R}^N$, where $|N|=2\omega-2$. They all meet every ω-clique in exactly one element and therefore the difference of the $\omega-1$ incident vectors in $\mathbb{R}^N$ of $S_i \subset N$ $(i=1,2,\ldots,\omega-1)$ and $T_1 \subset N$ is a set of $\omega-1$ linearly independent vectors all orthogonal to every ω-clique in $\mathbb{R}^N$.

Therefore, there are no more than $2(\omega-1)-(\omega-1)=\omega-1$ linearly independent $\omega-1$-cliques in $N(v)$, contradicting Padberg's [11] result which guarantees that G contains ω linearly independent cliques of size ω which contain v. (The algebraic argument that we have used just now is subsumed in a result of [7].)

Claim 2. $\omega(G\neg v)\leqslant\omega$.

Let Q be an arbitrary clique in $G\neg v$, let Q_0 denote the set of all vertices in Q that are vertices of G (rather than shrunk color classes of $N(v)$). Now $Q_0 \cup N[v]$ *does not contain all the vertices of* G, since then v would not be contained in any stable set larger than 2, so (since G is minimal imperfect) $\alpha\leqslant 2$ would follow implying that G is an odd antihole, contradicting the assumption.

So $Q_0 \cup N[v]$ induces a perfect graph, and therefore it is ω-colorable in G; such an ω-coloring uses exactly $\omega-1$ colors on $N(v)$, and then by Claim 1, the classes of this coloring are those of the unique $\omega-1$-coloring of $N(v)$, which, after shrinking, contain $Q\setminus Q_0$. Therefore, Q is ω-colored, whence $|Q|\leqslant\omega$ and the claim is proved.

Since G is minimal imperfect $\omega(G)<\chi(G)$; according to Claim 2 the other condition of the lemma is also satisfied, and therefore its assertion is true as well. □

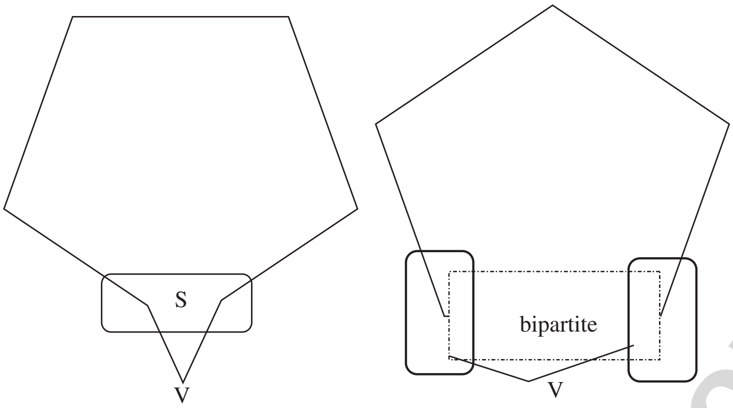

Fig. 1. No odd hole after the operation.

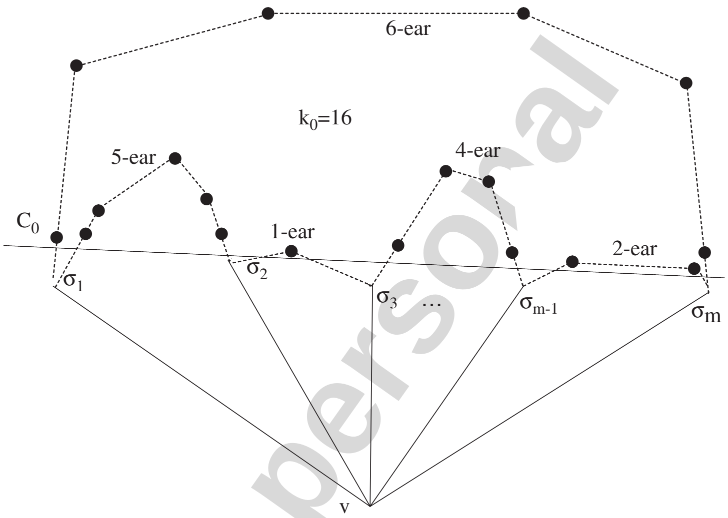

Fig. 2. C_0 and its ears.

3. Proof of Theorem 1.2

Let $\{S_1, \ldots, S_k\}$ ($k \in \mathbb{N}$) be the color classes of f (partitioning $N(v)$), and $N^* := \{\sigma_1, \ldots, \sigma_k\}$ the vertices that they become after shrinking; N^* is a clique of $G^* := G\neg(v, f)$.

We first prove that G^* does not contain any odd hole (Fig. 1). Indeed, an odd hole would contain one or two vertices from N^*:

- If it contains only one vertex $\sigma \in N^*$ (Fig. 1 left), which is the shrunk color-class $S \subseteq N(v)$, then the corresponding edges of G form either an odd hole or an odd chordless (s_1, s_2)-path ($s_1, s_2 \in S$), all vertices of which but s_1, s_2 are in $\overline{N}(v)$. In this latter case add v and the edges s_1v, s_2v to get anyway an odd hole in G, a contradiction.
- If it contains two vertices, σ_i and σ_j (Fig. 1 right), then replace the edges $a\sigma_i$, $\sigma_i\sigma_j$, $\sigma_j b$ of the hole by corresponding edges and a path: $a\sigma_i$, and $\sigma_j b$ can be replaced by as_i, s_jb, ($s_i \in S_i$, $s_j \in S_j$) and $\sigma_i\sigma_j$ by a chordless (necessarily odd) path in the bipartite graph $G(S_i \cup S_j)$ using that this graph is connected by assumption. We finally get an odd hole in G, a contradiction.

We prove now that G^* *does not contain an odd antihole.* We have already proved that it does not contain any C_5.

For a contradiction let C be an odd antihole of G^*, $|C| \geqslant 7$, and C has $m \in \mathbb{N}$ common vertices with N^*; let these be $\sigma_1, \ldots, \sigma_m$ in one of the two cyclic orders of C that we fix, and let $C_0 := C \backslash \{\sigma_1, \ldots, \sigma_m\} \subseteq V(G)$, and $m_0 = |C_0| - 2$ (Fig. 2). Clearly, since N^* is a clique, the predecessor and the successor of σ_i on C are not in N^*, for any $i \in \{1, \ldots, m\}$,

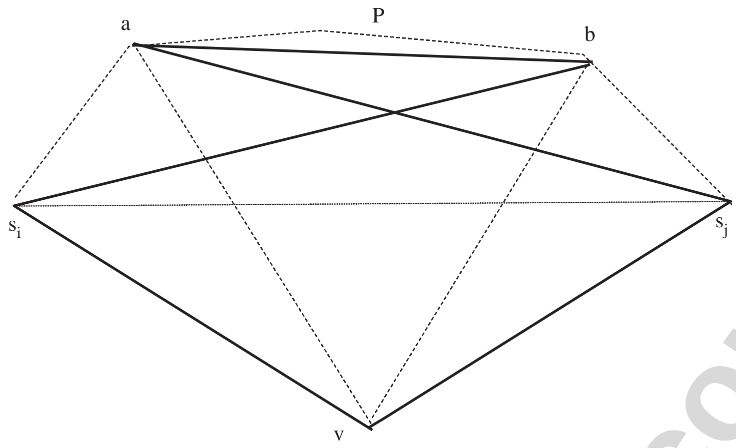

Fig. 3. There is an edge between s_i and s_j.

and all other vertices of C are neighbors of σ_i in G^*. Therefore, $\{\sigma_1, \ldots, \sigma_m\}$ is a clique, and $|N_{G^*}(\sigma_i) \cap C_0| = m_0$. The set $C_0 \subseteq V(G)$ induces a subgraph of $\overline{G}$ with m components, all paths. Such a path of $\overline{G}$ (the entire component) will be called an *ear* (of C); if it has p vertices, we will also say that it is a *p-ear*, and if the preceding and succeeding vertex on C are σ_i and σ_{i+1} ($i = 1, \ldots, m$, if $i = m$,then $i + 1 := 1$), then we say the ear is attached to i and $i + 1$.

As you see, we do not include the vertices σ_i, σ_{i+1} ($i \in \{1, \ldots, m\}$) to the ear attached to them; this ear will be referred to as the ear $(i, i + 1)$ or $(i + 1, i)$. The ear $(i - 1, i)$ is its *predecessor* and the ear $(i + 1, i + 2)$ its *successor*, both are ears *consecutive* to it. The ears $(i - 1, i)$ and $(i, i + 1)$ are consecutive ears; we will also speak about second consecutive ears, for instance $(i - 1, i)$ and $(i + 1, i + 2)$. We will use similar terms for the relative place of vertices on C with respect to the fixed cyclic order. (A vertex of $V(G^*)$ is consecutive to another on C if and only if they are neighbors in $\overline{G}$. We speak about 'neighbors' in G and consecutive vertices on C.)

If every S_i has a representative s_i that has the same neighbors in C_0 as σ_i, then replacing σ_i by s_i ($i = 1, \ldots, k$) we get an odd antihole in G: $s_i s_j \in E(G)$ also follows, for otherwise let P be the longer of the two paths of $\overline{G}$ that are the components of $C \setminus \{\sigma_i, \sigma_j\}$ in $\overline{G}$, and a, b the two endpoints of P. Since $|C| \geqslant 7$ we have $|P| \geqslant 4$, and it follows that v, s_i, b, a, s_j, v is a hole on five vertices (Fig. 3). In general, we will be able to find an odd antihole without these conditions, but will need some case-checking.

Unfortunately, such a representative does not always exist. (For P_5-free graphs it can be easily found though: supposing that there is none, either a C_5 or a P_5 can be found easily using a few lines of the proof of Claim 1 below, similarly to Fig. 3.) However, something weaker does exist, and similar, but somewhat refined arguments will turn out to be sufficient for proving the theorem from these weaker statements.

The weakening consists in allowing that from among the k_0 vertices that are not consecutive to σ_i on C at most one vertex can be missing from the neighbors of the representing vertex. The following definition squeezes out some useful particularity for this missing vertex:

Let us say that s_i ($i \in \{1, \ldots, k\}$) *represents* S_i (for C_0) if $s_i \in S_i$, and

$$N_G(s_i) \cap C_0 = N_{G^*}(\sigma_i) \cap C_0, \tag{1}$$

or if there exists $p = p(s_i) \in N_{G^*}(\sigma_i) \cap C_0$ such that

$$N_G(s_i) \cap C_0 = N_{G^*}(\sigma_i) \cap C_0 \setminus p \tag{2}$$

holds where p has all of the following properties:

(i) p is not second consecutive to σ_i,
(ii) p is not the endpoint of an ear, unless there exists $t_i \in S_i$ so that $p(t_i) \in C_0$ is consecutive to $p(s_i)$ on C, and then $p = p(s_i)$ is either the endpoint
- of a 3-ear attached to σ_i, or
- of a 2-ear not attached to σ_i, in which case p is the first vertex of the 2-ear (in the fixed cyclic order of C).

We define $p(s)$ for $s \in \cup_{i=1}^{k} S_i$ exactly if (2) holds otherwise we leave $p(s)$ undefined. Keep in mind that if s_i represents S_i and $s_i p \notin E(G)$ for some $p \in C_0$ not consecutive to s_i, then (2) holds and $p = p(s_i)$.

We first present the intuition behind the rest of the proof: we want to arrive at an odd circuit in $\overline{G}$ that does not induce any triangle. (Such an odd circuit contains an odd hole as induced subgraph.) We start with C. We will first of all make sure that $p(s_i)$ is not a vertex 'second consecutive' to σ_i. Secondly, we will push $p(s_i)$ 'towards the middle' of the ear that contains it (this will be necessary for proving that s_i is adjacent to s_j with some exceptions that are easy to control). 'Pushing' will be possible between two consecutive points of C_0; it follows that we are obliged to take for $p(s_i)$ the endpoint of an ear only if the other choice is a vertex second consecutive to σ_i or the other choice is also an endpoint of an ear, and in these cases we will be able to finish as well.

Claim 1. *For every $i \in \{1, \dots, k\}$ there exists s_i that represents S_i.*

This relies on the following simple fact on families of sets:

Fact. *Let $\mathscr{E}$ be a family of subsets of a set S none of which contains the other (a clutter). If the components of the graph*

$$H := \{ab\colon \text{ there exist } A, B \in \mathscr{E}, a \in A\backslash B, b \in B\backslash A\}$$

are chordless paths then $|\mathscr{E}| \leqslant 2$, and for the $A \neq B \in \mathscr{E}$, $|(A\backslash B) \cup (B\backslash A)| \leqslant 3$.

It is not hard to imagine how the Fact will be applied: the situation of Fig. 3 does also occur if s_i and s_j are in the same color-class. In that case $ab \in E(G)$ is not possible (we would have a C_5-like on the figure), so $ab \in E(\overline{G})$. Therefore, a color class S of $N_G(v)$ and the family $\mathscr{E}$ defined as the inclusionwise maximal elements of the form $N_G(s) \cap C_0$, $(s \in S)$ define a graph H (see the fact) which is a proper subgraph of C in $\overline{G}$. So all components of this H are induced paths! □

Proof of the Fact. Suppose $\mathscr{E}$ is a clutter and the components of H are chordless paths. For $A \neq B \in \mathscr{E}$ we have: $(|A\backslash B|, |B\backslash A|)$ is (1, 1), (1, 2) or (2, 1), otherwise we get circuits of size four or vertices of degree at least three. Now if for a contradiction $C \in \mathscr{E}$ is a third set different from these, then either all three sets have an element which is in neither of the two others, and these form a triangle in H, a contradiction; or say C has no such an element, that is, we can suppose $C \subseteq A \cup B$.

Since C is contained neither in A nor in B, $C \cap A\backslash B \neq \emptyset$, $C \cap B\backslash A \neq \emptyset$. Finally, $A \cap B\backslash C \neq \emptyset$ is also true, because $C \supseteq A \cap B$ would imply $C \supseteq A$ (if $|A\backslash C| = 1$), or $C \supseteq B$ (if $|B\backslash C| = 1$). Choosing an element in each of the three sets $C \cap A\backslash B \neq \emptyset$, $C \cap B\backslash A \neq \emptyset$, $A \cap B\backslash C \neq \emptyset$, the three chosen elements form a triangle, contradicting again the condition on H, and this contradiction proves the fact.

To prove the claim note that the graph H_i defined from the family $\mathscr{E}_i$ on the set $S := C_0$ like in the Fact, where

$$\mathscr{E}_i := \{X \subseteq C_0 : X = N_G(s) \cap C_0 \text{ for some } s \in S_i, \text{ and } X \text{ is maximal}\}, \tag{3}$$

is a subgraph of $\overline{G}(C_0)$ for all $i = 1, \dots, k$. This is true, because for $s, t \in S_i$, $a \in N_G(s) \cap C_0\backslash N_G(t)$, $b \in N_G(t) \cap C_0\backslash N_G(s)$ imply $ab \in E(\overline{G})$. *For otherwise, v, s, a, b, t, v would be a hole on five vertices*! (We just repeated formally the argument before the proof of the Fact—all this can be followed in Fig. 3 replacing s_i, s_j by t and s.)

But $\overline{G}(C_0)$ is a proper subgraph of a hole, so its components are chordless paths. Then by the Fact, $\mathscr{E}_i$ has at most two members, and their union obviously covers all the $k_0 > 0$ neighbors of σ_i in C_0. In particular, $\mathscr{E}_i = \emptyset$ is not possible.

If $\mathscr{E}_i$ has only one member, then let it be $N_G(s_i) \cap C_0$, $s_i \in S_i$: (1) holds so s_i represents S_i.

If $\mathscr{E}_i$ has two different members, $N_G(s_i) \cap C_0$, $N_G(t_i) \cap C_0$, then assume $|N_G(s_i) \cap C_0| \geqslant |N_G(t_i) \cap C_0|$. Since neither of these sets contains the other, and the symmetric difference of the two sets is at most three, we have

$$|N_G(s_i) \cap C_0\backslash N_G(t_i)| \in \{1, 2\}, \quad |N_G(t_i) \cap C_0\backslash N_G(s_i)| = 1. \tag{4}$$

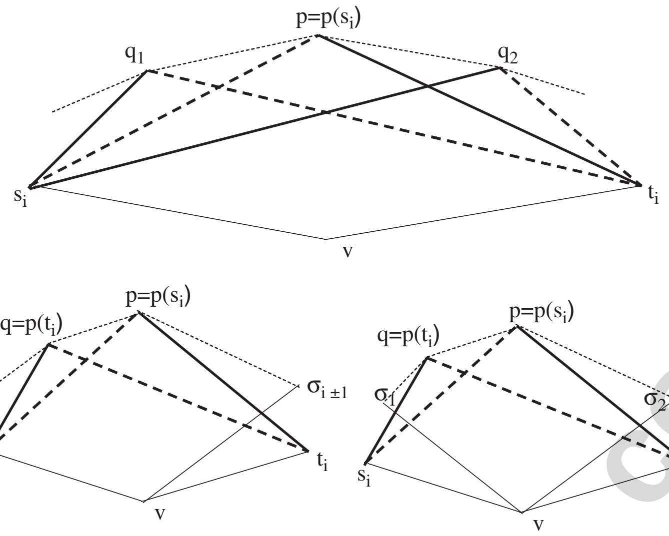

Fig. 4. Finding a representing vertex.

For all $s \in S_i$, $N_G(s) \cap C_0$ is a subset of one of the two elements of $\mathscr{E}_i$, and therefore

$$N_{G^*}(\sigma_i) \cap C_0 = \left(\bigcup_{s \in S_i} N_G(s) \right) \cap C_0 = (N_G(s_i) \cup N_G(t_i)) \cap C_0. \tag{5}$$

Let $p := p(s_i)$ be the unique element of $N_G(t_i) \cap C_0 \backslash N_G(s_i)$ (see (4)). From (5) we see that the neighbors of s_i and those of t_i together cover $N_{G^*}(\sigma_i) \cap C_0$, and then (4) tells us that t_i misses at most two points of $N_{G^*}(\sigma_i) \cap C_0$, and s_i misses one of them: p is the unique element of $N_{G^*}(\sigma_i) \cap C_0$ which is not contained in $N_G(s_i) \cap C_0$, that is, (2) holds. We still have to check that p satisfies (i) and (ii):

Comparing $\{p\} = N_G(t_i) \cap C_0 \backslash N_G(s_i)$ with the fact that the graph defined by $\mathscr{E}_i$ in (3) is a subgraph of $\overline{G}(C_0)$ (whose components are subpaths of the complement of C) we get that the (at most two) vertices of $N_G(s_i) \cap C_0 \backslash N_G(t_i)$ are consecutive to p on C (Fig. 4).

Case 1: $N_G(s_i) \cap C_0 \backslash N_G(t_i) =: \{q_1, q_2\}$. Follow the argument on the first (upper) drawing of Fig. 4. Since q_1, q_2 are consecutive to p on C and neither q_1 nor q_2 are consecutive to σ_i on C (they are both in $N_G(s_i)$), p is not second consecutive to σ_i (satisfies (i)); clearly, the three consecutive vertices q_1, p, q_2 have to be on the same ear, and therefore p is not the endpoint of an ear (satisfies (ii)). We proved that s_i represents S_i in this case.

Case 2: $N_G(s_i) \cap C_0 \backslash N_G(t_i) =: \{q\}$. Follow the argument on the two lower drawings of Fig. 4. We know $q = p(t_i)$, $p = p(s_i)$, and $pq \in E(\overline{G})$, and we will prove that either s_i or t_i represents S_i. Indeed, if one of p and q is neither second consecutive to σ_i nor the endpoint of an ear, then we are done. Otherwise,

- either one of p, q, say q, is consecutive to a vertex r itself consecutive to σ_i, and p is already an endpoint of the ear $(i, i \pm 1)$ (left drawing). Then r, q, p are the three vertices of a 3-ear attached to σ_i and $p = p(s_i)$ satisfies both (i) and (ii) of the definition, that is, s_i represents S_i;
- or both p and q are endpoints of an ear (right drawing). Since p and q are consecutive, they are then the two vertices of a 2-ear. Since both p and q are adjacent to σ_i (p is adjacent to t_i and q to s_i see the drawing) they cannot be consecutive to σ_i on C, so the 2-ear is not attached to σ_i. We can interchange the role of p and q if necessary, so that $p(s_i)$ is the first vertex of the 2-ear, as required.

We proved the claim, since in every case we found a vertex s_i that represents S_i.

Fix now s_i that represents S_i $(i = 1, \ldots, k)$.

Claim 2. *If $i, j \in \{1, \ldots, k\}$, $i \neq j$, then $s_i s_j \in E(G)$ unless all of the following hold*:

- $j = i \pm 1$ *(cyclically)—that is, after renumbering, for the simplicity of the notation,* $i = 1, \ \ j = 2$.
- *The ear* $(1, 2)$ *is a* 1*-ear.* (*Denote its only vertex by* a_1.)
- *The ear* $(2, 3)$ *is a* 2*-ear.* (*Denote its first vertex* a_2; $a_2 := p(s_1)$.)

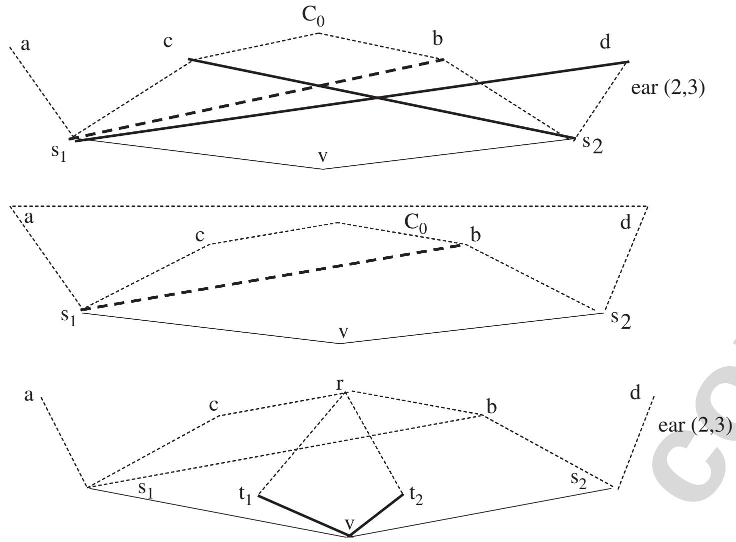

Fig. 5. Is s_1s_2 an edge?

Indeed, suppose $s_is_j \notin E(G)$, and prove the three stated facts:

Let a and b be the predecessor (in the cyclic order of C) of s_i and s_j, respectively. Since σ_i and σ_j are adjacent, they are not consecutive on C and therefore a and b are also not consecutive on C. So $ab \in E(G)$. By the choice of a and b, we have $s_ia \notin E(G)$, $s_jb \notin E(G)$.

Case 1: There is no 1-ear attached to σ_i and σ_j. We prove in this case that after possibly interchanging the notations, $i = 1$, $j = 2$, the ear $(1, 2)$ is a 3-ear and $p(s_1) = b$ (see Fig. 5 upper drawing); then we arrive at a contradiction.

Indeed, if $s_ib, s_ja \in E(G)$, then v, s_i, b, a, s_j, v is a hole on five vertices (similarly to Fig. 3); $s_ib \notin E(G)$ means that $p(s_i) = b$, and $s_ja \notin E(G)$ means $p(s_j) = a$.

Since s_i, s_j represent S_i, S_j, respectively—by (ii) of the definition of a representing vertex—$p(s_i) = b$ or $p(s_j) = a$ are possible only if $j = i + 1$ or $i - 1$, respectively, and the ear attached to σ_i and σ_j is a 3-ear. The cases $j = i + 1$ or $i - 1$ are symmetric, so indeed, we can suppose: $i = 1$, $j = 2$; the ear $(1, 2)$ is a 3-ear; $b = p(s_1)$ (see Fig. 5 upper drawing).

Let now c and d be the successor of s_1 and s_2, respectively. Again, $cd \in E(G)$. Since $d \neq b = p(s_1)$, we have $d \neq a$, and $s_1d \in E(G)$. (Indeed, $d = a$ would mean that there is also a 1-ear attached to σ_1 and σ_2 contradicting Case 1.) If $p(s_2) \neq c$ we arrive at a contradiction again, since then v, s_1, d, c, s_2, v is an odd hole on five vertices (upper drawing).

So $b = p(s_1)$ and $c = p(s_2)$. But then s_1d, $s_2a \in E(G)$, and as before v, s_1, d, a, s_2, v is an odd hole unless $ad \notin E(G)$. In other words, unless a and d are consecutive on C; then we have $|C| = 7$ (Fig. 5 drawing in the middle).

Consider now $t_1 \in S_1, t_2 \in S_2$; by (ii) of the definition of a representing vertex $p(t_1)$, $p(t_2)$ exist and are consecutive to $b = p(s_1)$, $c = p(s_2)$, respectively. If r denotes the middle vertex of the 3-ear $(1, 2)$, we get $p(t_1) = p(t_2) = r$.

If now $t_1t_2 \notin E(G)$, then v, t_1, b, c, t_2, v is a hole on five vertices like many times before. If on the contrary $t_1t_2 \in E(G)$, then we get a new kind of hole on five vertices, that does not include v: r, d, t_1, t_2, a, r (Fig. 5 lower drawing).

We have arrived at a contradiction in all the subcases of Case 1.

Case 2: There is a 1-ear attached to σ_i and σ_j. In this case we do not arrive at a contradiction, but at the claimed particular structure.

Suppose without loss of generality $i = 1$, $j = 2$; let a_1 be the unique vertex of the 1-ear $(1, 2)$, where a_1 is at the same time the successor of σ_1 $(a_1 = b = c)$; let a_2 be the successor of σ_2 on C; a denotes the predecessor of σ_1, as before (Fig. 6).

Now $a_2a \in E(G)$, because otherwise a_2 and a are consecutive on C, and $|C| = 5$; we also have $s_2a \in E(G)$ for otherwise $a = p(s_2)$, and since a is the endpoint of an ear, by (ii) of the definition of what 's_2 represents S_2 means' we have that there is also a 3-ear attached to s_1 and s_2: $|C| \leqslant 6$ follows, a contradiction.

Therefore, if $p(s_1) \neq a_2$, then $s_1a_2 \in E(G)$, and v, s_1, a_2, a, s_2, v is a hole on five vertices. So $a_2 = p(s_1)$, and the ear $(2, 3)$ is a 2-ear by (ii) of the definition of representing vertices. Claim 2 is proved.

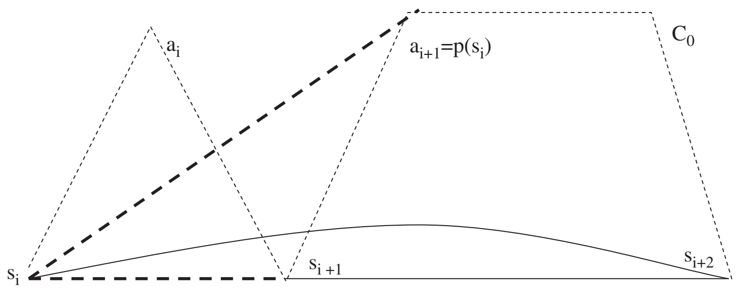

Fig. 6. Fault and its correction.

If $s_i s_{i+1} \notin E(G)$ holds ($s_i \in S_i$, $s_{i+1} \in S_{i+1}$), then we will say that (s_i, s_{i+1}) is a *fault*; Claim 2 provides then a_i, a_{i+1} and we will say that (a_i, a_{i+1}) is the *correction* of the fault (s_i, s_{i+1}). We see from Claim 2 that for different faults (s_i, s_{i+1}) and (s_j, s_{j+1}) $(i, j = 1, \dots, k)$, the vertex-sets $\{s_i, a_i, s_{i+1}, a_{i+1}\}$ and $\{s_j, a_j, s_{j+1}, a_{j+1}\}$ are disjoint. (If $(i, i+1)$ is a fault, $(i-1, i)$ and $(i+1, i+2)$ are not.)

Claim 3. *Let* $\{x, y, z\} \subseteq C_0 \cup \{s_1, \dots, s_k\}$ *be a stable set. Then there exists a fault* (s_i, s_{i+1}), $(i \in \{1, \dots, k\})$ *so that* $\{x, y, z\} = \{s_i, s_{i+1}, a_i\}$ *or* $\{x, y, z\} = \{s_i, s_{i+1}, a_{i+1}\}$, *where* (a_i, a_{i+1})*is the correction of the fault* (s_i, s_{i+1}).

Indeed, let $\{x, y, z\} \subseteq C_0 \cup \{s_1, \dots, s_k\}$ be a stable set. Since $\alpha(G(C_0)) \leqslant \alpha(C) \leqslant 2$ we have $|\{x, y, z\} \cap C_0| \leqslant 2$.

We also know that $\{x, y, z\} \cap C_0 \neq \emptyset$, because otherwise $x = s_1$, $y = s_2$, $z = s_3$ and we noticed that (1, 2), (2, 3) and (3, 1) cannot be all the three faults (faults are vertex-disjoint). Hence we have two cases:

Case 1: $|\{x, y, z\} \cap C_0| = 2$. Suppose say $x = s_1$, $y, z \in C_0$. Since $yz \in E(\overline{G})$, y and z are consecutive on C, in particular they are on the same ear. At most one of them is consecutive to σ_1, and at most one of them is $p(s_1)$; vertices z' for which neither of these hold satisfy $xz' \in E(G)$.

Hence, one of $\{y, z\}$, say y is consecutive to σ_1; then for the other, $z = p(s_1)$. Since y and z are consecutive we can conclude that $p(s_1)$ is second consecutive to σ_1 contradicting (i) of the definition of the representing vertex s_1.

Case 2: $|\{x, y, z\} \cap C_0| = 1$. Then say $z \in C_0$, and (x, y) is a fault. Follow now on (Fig. 6). We can suppose without loss of generality $x = s_1$, $y = s_2$.

If (a_1, a_2) denotes the correction of this fault, then what we have to prove is exactly $z \in \{a_1, a_2\}$.

Since $a_2 = p(s_1)$ (see Claim 2, Fig. 6) the only point in $C_0 \backslash \{a_1, a_2\}$ which is non-adjacent to s_1 is the predecessor a of s_1 on C. Therefore, $z \in \{a_1, a_2, a\}$.

But $z = a$ is not possible because we show $s_2 a \in E(G)$. Indeed, a is not consecutive to s_2 on C, so $s_2 a \notin E(G)$ is possible only if $a = p(s_2)$. But since a is the last vertex of an ear this can hold only if a is contained in a 3-ear attached to s_2 (see (ii) of the definition of representing vertices). If this were true then C would consist only of a 1-ear and a 3-ear, so $|C| = 6$, a contradiction that finishes Case 2 and at the same time the proof of Claim 3.

To finish the proof of the theorem note that the cyclic order of C determines a cyclic order of $C_0 \cup \{s_1, \dots, s_k\}$. Since $\sigma_i x \in E(\overline{G}^*)$ implies that $s_i x \in E(\overline{G})$, $C_0 \cup \{s_1, \dots, s_k\}$ is also an odd circuit in $\overline{G}$. However, edges of the form $s_i p(s_i) \in E(\overline{G})$ are chords in this circuit!

If (s_i, s_{i+1}) is a fault and (a_i, a_{i+1}) is its correction then delete a_i and s_{i+1} and replace the subpath $s_i, a_i, s_{i+1}, a_{i+1}$ by s_i, a_{i+1} (Fig. 6). Let us call this operation the *cutting off* of the fault (s_i, s_{i+1}).

The cutting off of a fault does not change the parity of the number of edges, and since the faults are pairwise disjoint, we can cut off all the faults independently of one another. Let us denote by C' the (uniquely determined) circuit of $\overline{G}$ we get at the end. We have just checked that $|C'|$ is an odd circuit in $\overline{G}$.

Showing $\alpha(C') \leqslant 2$ will finish the proof now. Suppose

$$\{x, y, z\} \subseteq C' \subseteq C_0 \cup \{s_1, \dots, s_k\}$$

is a stable set. By Claim 3 $\{x, y, z\} = \{s_i, s_{i+1}, a_i\}$ or $\{x, y, z\} = \{s_i, s_{i+1}, a_{i+1}\}$, where (s_i, s_{i+1}) is a fault and (a_i, a_{i+1}) is its correction. But this is impossible, since all faults have been cut off! □

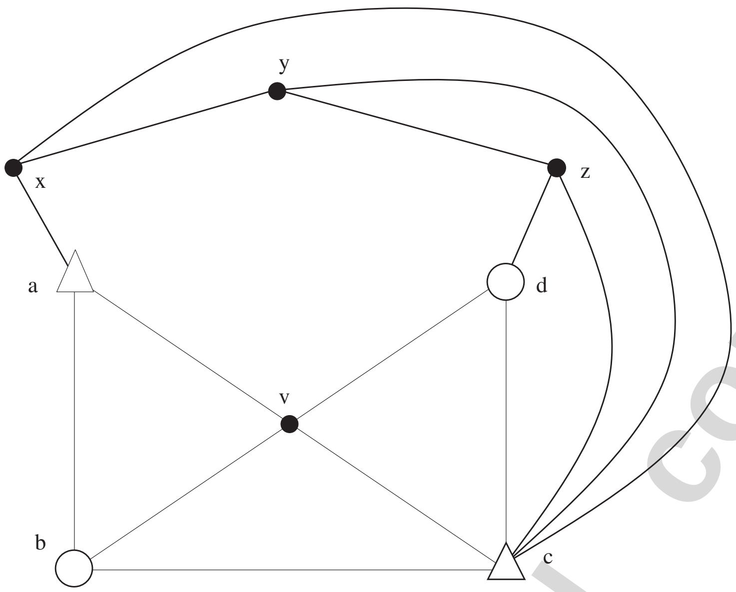

Fig. 7. A stable set alone cannot be shrunk, even if it is a color class.

4. Remarks, questions, variations

In [15], the strong perfect graph conjecture is deduced from the assumption that the intersection graphs of the ω-cliques of every minimal imperfect graph have a vertex of degree $2\omega - 2$ and the challenge of deducing the strong perfect graph conjecture from the same assumption placed on the minimal imperfect graphs themselves (rather than the intersection graphs) is raised. Clearly, this challenge (and some more, see the introduction) is met in the present paper, sharpening several earlier results for instance one arriving at the same conclusion when a graph has two neighboring vertices of degree $2\omega - 2$ [14].

One of the referees asked whether shrinking only one of the color classes in $N_G(v)$ always transforms a Berge graph into a Berge graph. The following example shows that the answer is negative:

On the following figure a connected coloration of $N(v)$ is indicated. The graph is perfect: $\{a, c\}$, $\{b, d, y\}$, $\{v, x, z\}$ is a 3-coloration, and any induced subgraph of it that does not contain a triangle is 2-colorable. (Otherwise it would contain an odd hole, but any odd hole containing the vertex c avoids 4 neighbors of c and therefore it should have at most $8 - 4 = 4$ vertices; on the other hand $G - c$ is perfect, since it is a C_7 with a short chord.) However, the shrinking of $\{b, d\}$ results in a C_5.

It would of course be good to weaken the condition of connected colorations, and to find the right weakening could be decisive for a shorter proof of the SPGT, or a simpler recognition algorithm for perfect graphs. The example of Fig. 7 shows that the operation that we are treating can surely not be decomposed into the natural smaller steps of shrinking color classes of a connected coloration one by one in an arbitrary order.

In Theorem 1.1 little extra work is sufficient to prove that $G\neg v$ is in fact a minimal imperfect graph. Under some tighter conditions this has been known [14] and can be inverted: the "inverse" is constructed by [1] of Boros et al., where for given graph H (with certain precised properties) a graph G is constructed such that $H = G\neg(v, f)$ for some v and f.

We do not know whether the existence of a connected coloring can be decided in polynomial time in a perfect graph.

Adaptations of our proof of Theorem 1.2 yield the following results:

Theorem 4.1 (*Theorem on holes*). *Let k and m be positive integers. If G is a graph such that*

(H) *G has no induced subgraph isomorphic to a* C_i *with* $i \geqslant k$ *or a* P_j *with* $j \geqslant m$,

if v *is a vertex of G, and if f is a* (*not necessarily connected*) *coloring of* $N_G(v)$, *then* (H) *is satisfied with* $G\neg(v, f)$ *in place of G.*

Theorem 4.2 (*Theorem on antiholes*). *If G is a graph such that*

(A) *G has no induced subgraph isomorphic to an odd antihole or to* P_5,

if v *is a vertex of* G, *and if* f *is a* (*not necessarily connected*) *coloring of* $N_G(v)$, *then* (A) *is satisfied with* $G\neg(v, f)$ *in place of* G.

It is essential that the introduced operation increases the clique size of an induced subgraph if and only if the original graph contains an odd antihole: this suggests that the operation may relate the Berge property to the chromatique number or the clique number, and therefore to perfectness. Recently two polynomial algorithms have been found for recognizing Berge-graphs, by [4,6,2], which are both independent of the SPGT and of its proof.

Acknowledgment

This work would not have been completed without the stimulating announcement by Paul Seymour of the workshop on Berge graphs (Princeton, September 2001), and his invitation. I am also indebted to Grigor Gasparyan for a quick and very careful checking of the proof before the meeting, and to two referees for their comments.

Only Vašek Chvátal could help me find the right words for thanking Vašek's pointing out the decisive points of a manuscript that Chvátal has found pointless and might have crossed out with a stroke of his pen. Do not be puzzled if you recognize his particular style: yes, he rewrote the first and last sections reformulating all the results, and providing convincing strength to the presentation of theorems that have been weakened by the proof of the strong perfect graph conjecture. (He has no responsibility in the sloppiness of the proof.)

References

[1] E. Boros, V. Gurvich, S. Hougardy, Recursive generation of partitionable graphs, J. Graph Theory 41 (4) (2002) 259–285.
[2] M. Chudnovsky, G. Cornuéjols, X. Liu, P. Seymour, K. Vušković, Cleaning for Bergeness, November 2002.
[3] M. Chudnovsky, N. Robertson, P. Seymour, R. Thomas, The strong perfect graph theorem, June 2002.
[4] M. Chudnovsky, P. Seymour, Recognizing Berge graphs, November 2002.
[5] V. Chvátal, N. Sbihi, Recognizing claw-free perfect graphs, J. Combin. Theory Ser. B 44 (1988) 154–176.
[6] G. Cornuéjols, X. Liu, K. Vušković, A polynomial algorithm for recognizing perfect graphs, November 2002.
[7] J. Fonlupt, A. Sebő, On the clique rank and the coloration of perfect graphs, in: R. Kannan, W. Pulleyblank (Eds.), Integer Programming and Combinatorial Optimization 1, University of Waterloo Press, 1990.
[9] S.E. Markossian, I.A. Karapetian, On critically imperfect graphs, in: R.N. Tonoian (Ed.), Prikladnaia Matematika, Erevan University, 1984.
[10] E. Olaru, Zur Characterisierung perfekter Graphen, Elektron. Informationsverarbeitung Kybernet. (EIK) 9 (1973) 543–548.
[11] M. Padberg, Perfect zero–one matrices, Math. Programming 6 (1974) 180–196.
[12] K.R. Parthasarathy, G. Ravindra (1961), The strong perfect graph conjecture is true for $K_{1,8}$-free graphs, J. Combin. Theory Ser. B 23 (1976) 272–243.
[13] B. Reed, A semi-strong perfect graph theorem, Ph.D. Thesis, McGill University, Québec, 1986.
[14] A. Sebő, On critical edges in minimal imperfect graphs, J. Combin. Theory 47 (1) (1996) 62–85.
[15] A. Sebő, The connectivity of minimal imperfect graphs, J. Graph Theory 23 (1) (1996) 57–85.