7

Perfect graphs

NICOLAS TROTIGNON

1. Introduction
2. Lovász's perfect graph theorem
3. Basic graphs
4. Decompositions
5. The strategy of the proof
6. Book from the Proof
7. Recognizing perfect graphs
8. Berge trigraphs
9. Even pairs: a shorter proof of the SPGT
10. Colouring perfect graphs
References

Perfect graphs were defined by Claude Berge in the 1960s. They are important objects for graph theory, linear programming and combinatorial optimization. Berge made a conjecture about them (now called the strong perfect graph theorem or SPGT) which was proved by Chudnovsky, Robertson, Seymour and Thomas in 2002. This survey about perfect graphs mostly focuses on the SPGT.

1. Introduction

Every graph G clearly satisfies $\chi(G) \geq \omega(G)$, where $\omega(G)$ is the clique number of G, because the vertices of a clique must receive different colours. A graph G is *perfect* if every induced subgraph H of G satisfies $\chi(H) = \omega(H)$. A chordless cycle of length $2k+1$, for $k \geq 2$, satisfies $3 = \chi > \omega = 2$, and its complement satisfies $k+1 = \chi > \omega = k$; these graphs are therefore *imperfect*. Since perfect graphs are closed under taking induced subgraphs, they must be defined by excluding a family $\mathcal{F}$ of graphs as induced subgraphs. The *strong perfect graph theorem* (SPGT for short) states that the two examples just given are the only members of $\mathcal{F}$.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Let us make this more formal. A *hole* in a graph G is an induced cycle of length at least 4, and an *antihole* is a hole of $\overline{G}$. A graph is *Berge* if it does not contain an odd hole or an odd antihole. The following was conjectured by Berge [3] in the 1960s and was the object of much research until it was finally proved in 2002 by Chudnovsky, Robertson, Seymour and Thomas [13].

Theorem 1.1 (Strong perfect graph theorem) *A graph is perfect if and only if it is Berge.*

One direction of the proof is easy: every perfect graph is Berge since, as we observed above, odd holes and antiholes satisfy $\chi = \omega + 1$. The proof of the converse statement is very long and relies on structural graph theory. The main step is a *decomposition theorem* (Theorem 5.1), which states that every Berge graph is either in a well-understood *basic* class of perfect graphs (see Section 3), or has some *decomposition* (see Section 4). The goal of this survey is to make this result and its meaning understandable to a non-specialist mathematician. For this purpose, not only the proof is surveyed, but also some results that were seminal to it, and some that were proved after it. Another goal is to present perfect graphs as a lively subject for researchers, by mentioning several attractive open questions.

We start with an open question that has the same flavour as the SPGT. In 1987 Gyárfás [35] proposed the following generalization of a perfect graph: a graph G is *χ-bounded by a function f* if every induced subgraph H of G satisfies $\chi(H) \leq f(\omega(H))$; thus, perfect graphs are χ-bounded by the identity function. There might exist a short proof that, for some function f (possibly fastly increasing), all Berge graphs are χ-bounded by f, but so far the proof of the SPGT is the only known proof of this fact. In [35] the following conjecture is stated.

Conjecture A *There exists a function f such that, for all graphs G, if G has no odd hole, then G is χ-bounded by f.*

The SPGT can be rephrased as follows: perfect graphs and Berge graphs form the same class. We still keep two names for the same class, and use *Berge* when we concentrate on excluding holes and antiholes, and *perfect* when we focus on colourings. Note that the distinction between *Berge* and *perfect* has to be kept when we sketch the proof of the SPGT, since in the proof the two classes are supposed to be potentially different. A *minimally imperfect graph* is an imperfect graph, every proper induced subgraph of which is perfect. A restatement of the SPGT is thus that minimally imperfect graphs are precisely the odd holes and antiholes. Many statements about minimal imperfect graphs are therefore trivial to check by using the SPGT, but when proving the SPGT it is essential to prove them by other means. Observe that a minimal (and therefore minimum) counter-example to the SPGT has to be a Berge minimally imperfect graph.

We mostly follow the terminology from [13], which is not fully standard. Since the whole theory deals with induced subgraphs, we write *G contains H* to mean that H is an induced subgraph of G and write *path* instead of chordless path or induced path.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

When a and b are vertices of a path P, we denote by aPb the subpath of P whose ends are a and b. A subset A of $V(G)$ is *complete* to a subset B of $V(G)$ if A and B are disjoint and every vertex of A is adjacent to every vertex of B; we also say that B is *A-complete*.

We use the prefix *anti* to mean a property or a structure of the complement (as for holes and antiholes). For instance, an *antipath* in G is a path in $\overline{G}$, A is *anticomplete* to B means that no edge of G has an end in A and the other one in B, and G is *anticonnected* if its complement is connected. By $C(A)$ we mean the set of all A-complete vertices in G, and by $\overline{C}(A)$ the set of all A-anticomplete vertices. By a *colouring* of a graph, we mean an optimal colouring of the vertices.

Since we focus on the SPGT and its proof, most of this survey is on the structure of Berge graphs, and we omit many other important aspects of the theory, such as the ellipsoid method used by Gröstchel, Lovász and Schrijver [34] to colour any input perfect graph in polynomial time. The theory of semi-definite programming started there.

Several earlier surveys exist. The survey of Lovász [45] from the 1980s is still good, while two books are completely devoted to perfect graphs [4], [50], and contain a lot of material cited below. Part VI of Schrijver [55] is the most comprehensive survey on perfect graphs. Chudnovsky, Robertson, Seymour and Thomas [12] wrote a good survey just after their proof, while Roussel, Rusu and Thuillier [53] wrote a good survey about the long sequence of results, attempts and conjectures that finally led to the successful approach.

For more historical aspects, see Section 67.4g of Schrijver [55]. Berge and Ramírez Alfonsín ([50, pp.1–12]) wrote an article about the origin of perfect graphs, and Seymour [56] wrote an article, interesting even to a non-mathematician, telling the story of how the SPGT was proved.

2. Lovász's perfect graph theorem

As pointed out by Preissmann and Sebő (see [50, pp. 185–214]), the following theorem conveniently gives a weak (1) and a strong (2) characterization of perfect graphs. Hence, to prove perfection checking the weak one is enough, while to use perfection one may rely on the strong one. This allows a kind of leverage that is used often in what follows.

Theorem 2.1 *Let G be a graph. The following three statements are equivalent.*

(1) *G is perfect.*
(2) *For each induced subgraph H of G and each $v \in V(H)$, H contains a stable set that contains v and intersects all maximum cliques of G.*
(3) *For each induced subgraph H of G, H contains a stable set that intersects all maximum cliques of G.*

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Proof To prove that (1) implies (2), consider an optimal colouring of H and the colour class S that contains v. Then S is a stable set that intersects all maximum cliques of H, for otherwise $\chi(H \setminus S) \geq \omega(H)$, which is a contradiction.

Trivially, (2) implies (3).

To prove that (3) implies (1), consider the following greedy colouring algorithm:

Step 1. Set $i \leftarrow 1$ and $H \leftarrow G$.
Step 2. While H is non-empty, consider a stable set of H, as in (3), give it colour i, and set $H \leftarrow H \setminus S$ and $i \leftarrow i + 1$.

Since ω decreases at each iteration in Step 2, this algorithm produces a colouring of G with $\omega(G)$ colours, and it can be processed for any induced subgraph of G. Thus, G is perfect. ■

Replicating a vertex v of a graph G means adding a new vertex v' adjacent to v and to all neighbours of v. Note that the property $\chi = \omega$ is not preserved by replication. As an example, consider the non-perfect graph G obtained by replicating one vertex of C_5. Clearly, $\chi(G) = \omega(G) = 3$. However, by replicating any vertex of degree 2 in G, a graph G' such that $\chi(G') = 4 > \omega(G') = 3$ is obtained. Thus, the following lemma due to Lovász [44] is more surprising than it may look at first glance.

Theorem 2.2 (Replication lemma) *Perfection is preserved by replication.*

Proof Let G be a perfect graph and let G' be obtained from G by replicating a vertex v. We show that G' satisfies Theorem 2.1(3). Let H be an induced subgraph of G'. We look for a stable set S that intersects all maximum cliques of H. If H contains at most one vertex from $\{v, v'\}$, then it is isomorphic to an induced subgraph of G, so clearly S exists. Otherwise, $H \setminus v'$ is perfect, so, by Theorem 2.1(2), there exists a stable set S that contains v and intersect all maximum cliques of $H \setminus v'$. In fact, S intersects all maximum cliques of H, because a maximum clique of H contains v if and only if it contains v'. ■

The following result was conjectured by Berge as the *weak perfect graph conjecture* and is now called the *perfect graph theorem*. We give the proof that, among the available ones, is most closely related to the rest of this chapter.

Theorem 2.3 *If a graph is perfect, then its complement is perfect.*

Proof Let G be a perfect graph. Construct a graph G' as follows: start from G and replicate every vertex v $\alpha_v - 1$ times, where α_v is the number of maximum stable sets of G that contain v; note that replicating -1 times means deleting the vertex, and replicating 0 times means doing nothing. From its construction, G' can be covered by k disjoint maximum stable sets that therefore form an optimal colouring of G'. Since G' is perfect, by Theorem 2.2, it follows that G' has a clique K' of size k. Since a clique and a stable set intersect in at most one vertex, K' intersects all maximum stable sets of G'. Now construct a clique K of G by taking, for each vertex of K',

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

the vertex of G from which it is replicated. The clique K that we obtain intersects all maximum stable sets of G. By the same argument, a clique intersecting all maximum stable sets can be found in any induced subgraph of G. Thus the complement of G satisfies Theorem 2.1(3) and is therefore perfect. ■

In [35], the question of generalizing Theorem 2.3 is discussed, but since then, little progress has occurred in this direction (see [36]). In particular, the following neat generalization of Theorem 2.3 (due to Gyárfás) is still open.

Conjecture B *There exists a function f such that, for all graphs G, if G is χ-bounded by $x \mapsto x + 1$, then $\overline{G}$ is χ-bounded by f.*

The perfect graph theorem has a polyhedral proof found by Fulkerson [30], related to polyhedral characterizations of perfect graphs discovered by him and Chvátal [18]. Lovász [43] proved a deep characterization of perfect graphs suggested by Hajnal.

Theorem 2.4 *A graph G is perfect if and only if every induced subgraph H of G satisfies $\alpha(H)\,\omega(H) \geq |V(H)|$.*

This characterization implies that deciding the perfection of an input graph is a CoNP problem (see [50, pp. 185–214]). It can be proved by a simple argument, relying on linear algebra, discovered by Gasparian [32]. Since the characterization is self-complementary, it gives another proof of the perfect graph theorem. This characterization is the starting point of many developments of great signifiance, such as the theory of partitionable graphs surveyed by Preissmann and Sebő (see [50, pp. 185–214]). It has deep connections with combinatorial optimization, as explained in Cornuéjols [25].

3. Basic graphs

In this section, we survey the five basic classes that are used in the proof of the SPGT: bipartite graphs, their complements, line graphs of bipartite graphs, their complements, and double split graphs. Bipartite graphs are easily checked to be perfect, and so, by Theorem 2.3, their complements are perfect; this can be restated as: 'every bipartite graph G satisfies $\theta(G) = \alpha(G)$'. Since $|V(G)| = \theta(G) + \nu(G) = \alpha(G) + \tau(G)$ for any triangle-free graph G, we deduce that every bipartite graph G satisfies $\nu(G) = \tau(G)$; this can be rephrased as: 'the complements of line graphs of bipartite graphs are perfect'. Again applying Theorem 2.3, we deduce that line graphs of bipartite graphs are perfect; this can be restated as: 'every bipartite graph G satisfies $\Delta(G) = \chi'(G)$'. Theorem 2.3 thus implies the perfection of three of the four *historical basic classes* of perfect graphs: bipartite graphs, their complements, their line graphs and the complements of their line graphs.

We now turn our attention to the fifth class that is first presented in the proof of the SPGT: *double split graphs*. As presented in [13], this class is not closed under taking

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

induced subgraphs, so we prefer to use *doubled graphs*, that are easily seen to form the class of induced subgraphs of double split graphs.

A *good partition* of a graph G is a partition (X, Y) of $V(G)$ (possibly, X or $Y = \varnothing$) such that:

- Every component of $G[X]$ has at most two vertices, and every anticomponent of $G[Y]$ has at most two vertices.
- For every component C_X of $G[X]$, every anticomponent C_Y of $G[Y]$, and every vertex v in $C_X \cup C_Y$, there exist at most one edge and at most one antiedge between C_X and C_Y that are incident to v.

A graph is *doubled* if it has a good partition. Doubled graphs are easily seen to be perfect, by a direct colouring argument. They are also closed under taking induced subgraphs and complements. A graph is *basic* if it belongs to at least one of the five classes defined here. For each basic class, the characterization by excluding induced subgraphs is known (see Beineke [2] for line graphs, and Alexeev, Fradkin and Kim [1] for doubled graphs); the recognition can be performed in polynomial time (see Lehot [40] or Roussopoulos [54] for line graphs and Maffray [46] for doubled graphs). Also the colouring and maximum clique problems can be solved in polynomial time (see Schrijver [55] for the historical classes and Maffray [46] for doubled graphs).

In this section, we have focused on the basic graphs that play an important role in the proof of the SPGT. But any class of graphs whose perfection is simple to prove can potentially serve as a basic class of a decomposition theorem, so all classes are potentially of interest. The book of Brandstädt, Le and Spinrad [6] is on general graph classes, but contains a lot of material on perfect graphs. The most complete catalogue of classes of perfect graphs seems to be that written by Hougardy [38], which describes 120 classes. Chapter 66 in Schrijver [55] also contains an extensive survey about classes of perfect graphs. Golumbic's book [33] surveys algorithmic aspects of several classes of perfect graphs.

The most important class that we omit from here is the seminal class of hole-free graphs; it is the first class with a decomposition theorem, has some connections with graph minor theory and tree-width, and is also important in fast graph-searching algorithms (see Sections 9.7–9.8 in Bondy and Murty [5]). Another important class that has connections with ordered sets is the class of comparability graphs, introduced by Gallai [31]; the English translation by Maffray and Preissmann ([50, pp. 25–66]) contains a short survey.

4. Decompositions

By a *decomposition* of a graph we mean a way to partition its vertices with some prescribed adjacencies. A decomposition is *useful* if it can be proved that a minimum counter-example to the SPGT cannot admit the decomposition. Indeed, suppose that

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

we can prove a statement such as: 'every Berge graph is either basic or has a useful decomposition'. The SPGT can then be proved as follows: consider a minimum counter-example, which is thus a Berge graph, imperfect, and of minimum size. It does not admit a decomposition (because the decomposition is useful), and since it is imperfect, it cannot be basic, contradicting the statement.

A 2-*join* of a graph G, first defined by Cornuéjols and Cunningham [26], is a partition (X_1, X_2) of $V(G)$ for which there exist disjoint non-empty sets $A_1, B_1 \subseteq X_1$, $A_2, B_2 \subseteq X_2$ satisfying:

- A_1 is complete to A_2, B_1 is complete to B_2, and these edges are the only ones between X_1 and X_2.
- $|X_1| \geq 3$ and $|X_2| \geq 3$.
- every component of $G[X_i]$ intersects A_i and B_i, for $i = 1$ and 2.
- if $|A_i| = |B_i| = 1$, then $G[X_i]$ is not a path of length 2 joining the members of A_i and B_i, for $i = 1$ and 2.

Cornuéjols and Cunningham [26] proved that a minimally imperfect graph admitting a 2-join must be an odd hole (so 2-joins are useful). Note that here, *path* 2-*joins* are allowed: these are 2-joins such that, for $i = 1$ and 2, $|A_i| = |B_i| = 1$ and $G[X_i]$ is a path from the unique vertex in A_i to the unique vertex in B_i. Some papers (mostly those involving Conforti, Cornuéjols or Vušković) restrict the notion of 2-joins to *non-path* 2-*joins*. Section 8 discusses the relevance of excluding path 2-joins. When X_i, A_i and B_i are as above, it is customary to set $C_i = X_i \setminus (A_i \cup B_i)$. It is easy to prove that in a Berge graph, all paths from A_i to B_i with interior in C_i have the same parity, since otherwise an odd hole exists. There are thus two kinds of 2-joins: *odd* 2-joins and *even* ones. If (X_1, X_2) is a 2-join of $\overline{G}$, then it is a *complement* 2-*join* of G.

When G is a graph and $A \subseteq V(G)$, we denote by $C(A)$ the set of vertices of G complete to A, and by $\overline{C}(A)$ the set of vertices of G anticomplete to A. A *homogeneous pair* (first defined by Chvátal and Sbihi [20], who proved that they are useful) is a pair of disjoint sets $A, B \subseteq V(G)$ for which $|A|, |B| \geq 2$, every vertex of A has a neighbour and a non-neighbour in B, every vertex of B has a neighbour and a non-neighbour in A, and the sets

$$A,\ B,\ C(A) \cap \overline{C}(B),\ \overline{C}(A) \cap C(B),\ C(A) \cap C(B) \text{ and } \overline{C}(A) \cap \overline{C}(B)$$

are all non-empty and partition $V(G)$.

The decompositions presented so far are nice in the following sense. When applied recursively, they yield decomposition trees of polynomial size that allow the solution of several problems; the machinery is too heavy to be presented here (see Section 10).

We now turn our attention to other cutsets that do not have this nice property. A *star cutset* in a graph G (first defined by Chvátal [19]) is a set S of vertices for which $G \setminus S$ is disconnected and S contains a vertex v called the *centre*, complete to $S \setminus v$. The following is a result of Chvátal [19].

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Theorem 4.1 (Star cutset lemma) *A minimally imperfect graph has no star cutset.*

Proof Let G be a minimally imperfect graph, and suppose for a contradiction that G has a star cutset S centred at v. Let (X, Y) be a partition of $V(G \setminus S)$ for which $|X|$ and $|Y| \geq 1$ and X is anticomplete to Y. We now prove that G satisfies Theorem 2.1(3); this implies that G is perfect, giving the required contradiction. Since every proper induced subgraph of G is perfect, it remains to find the desired stable set in G. By Theorem 2.1(2), there exists a stable set A_X in $G[S \cup X]$ that contains v and intersects all maximum cliques of $G[S \cup X]$. A similar stable set A_Y exists in $G[S \cup Y]$. We see that $A_X \cup A_Y$ is then a stable set of G that intersects all maximum cliques of G. ■

It is quite easy to turn the above proof into a colouring algorithm that would, for instance, colour any perfect graph every induced subgraph of which is either basic or decomposable by a star cutset. The algorithm would output a stable set that intersects all maximum cliques and, following the lines of the proof of Theorem 2.1, would yield a colouring algorithm. Unfortunately, this algorithm does not run in polynomial time, because the star cutset can be very big (for instance, all the graph except for two vertices). In this case, the complexity analysis of the recursive calls leads to an exponential number of calls. There is a similar problem with the generalizations that we consider next. This is the main reason why the decomposition of Berge graphs does not lead to a polynomial-time colouring algorithm.

A *skew partition* of a graph G (first defined by Chvátal [19]) is a partition (A, B) of $V(G)$ for which $G[A]$ is not connected and $G[B]$ is not anticonnected; in this case, we say that B is a *skew cutset*. Following a prophetic insight that some *self-complementary* decomposition generalizing the star cutset should play a role, Chvátal conjectured that a minimally imperfect graph has no skew partition and (a less formal statement) that skew partitions should appear in the decomposition of Berge graphs. Observe that if a graph with at least five vertices and at least one edge has a star cutset, then it has a skew partition.

The proof of Chvátal's conjectures escaped their investigators, but several fruitful attempts were made. In particular, many special kinds of skew partitions were proved not to be in minimal imperfect graphs (see Reed [51] for a survey). In the opposite direction, a generalization of skew partitions was proved to decompose all Berge graphs in the following theorem of Conforti, Cornuéjols and Vušković [23]. Here, a *double star cutset* in a graph G is a set $S \subseteq V(G)$ such that $G \setminus S$ is disconnected and S contains an edge vw for which every vertex of S is adjacent to at least one of v and w. Observe that antiholes of length at least 6 have double star cutsets. Note that, for a Berge graph G, the following yields information for both G and $\overline{G}$.

Theorem 4.2 *A graph with no odd hole is either basic, or has a* 2*-join or a double star cutset.*

One of the breakthroughs made in the proof of the SPGT is the concept of a *balanced skew partition*. For a graph G, a partition (skew or not) (A, B) of $V(G)$ is *balanced* if every path of length at least 3 with ends in B and interior in A, and

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

every antipath of length at least 3 with ends in A and interior in B have even length. It is straightforward to check that a partition (A, B) of a Berge graph is balanced if and only if adding a vertex complete to B and anticomplete to A yields a Berge graph.

As we will see, the notion of a balanced skew partition is sufficiently particular to allow a short proof that a minimum counter-example to the SPGT cannot contain it, and is sufficiently general to be found in all non-basic Berge graphs that cannot be decomposed otherwise.

Theorem 4.3 *If (A, B) is a balanced partition of a perfect graph G, and if every Berge graph of order at most $|V(G)|+1$ is perfect, then $G[B]$ admits a colouring that can be extended to a colouring of G.*

Proof Consider the graph G' that is obtained by adding to G a clique of size $k = \omega(G) - \omega(G[B])$ complete to B and anticomplete to A. It is Berge because (A, B) is balanced. So it is perfect when $k \leq 1$ from our assumption, and it is also perfect when $k \geq 2$ by several applications of Theorem 2.2. Observe that $\omega(G') = \omega(G)$. An $\omega(G)$-colouring of G' then yields a colouring of $G[B]$ that extends to a colouring of G. ■

The following theorem is due to Chudnovsky, Robertson, Seymour and Thomas [13].

Theorem 4.4 *A minimum imperfect Berge graph admits no balanced skew partition.*

Proof Let G be a minimum imperfect Berge graph, so $\chi(G) > \omega(G)$. By Theorem 2.3, $\overline{G}$ is also a minimum imperfect Berge graph. Let (A, B) be a balanced skew partition in G, so A partitions into two sets A_1 and A_2 anticomplete to one another, and B partitions into two sets X and Y complete to one another. By Theorem 4.1 $|A_1|$ and $|A_2| \geq 2$, for otherwise the vertex in A_1 or A_2 would be the centre of a star cutset in $\overline{G}$. It follows from the minimality of G that every Berge graph of size $|B \cup A_i|+1$ is perfect for $i = 1$ and 2. By Theorem 4.3, consider an $\omega(G[B])$ colouring C_i of $G[B]$ that extends to a colouring of $G[B \cup A_i]$. Let X_i be the set of vertices of $G[B \cup A_i]$ whose colour in the colouring C_i is present in X, and let $Y_i = (B \cup A_i) \setminus X_i$. Because of the colouring C_i, $\omega(G[X_i]) = \omega(G[X])$, so $\omega(G[X_1 \cup X_2]) = \omega(G[X])$. By the minimality of G, it follows that $G[X_1 \cup X_2]$ has an $\omega(G[X])$-colouring. Because of the colouring C_i, $\omega(G[Y_i]) = \omega(G) - \omega(G[X_i]) = \omega(G) - \omega(G[X])$, so $\omega(G[Y_1 \cup Y_2] = \omega(G) - \omega(G[X])$. By the minimality of G, it follows that $G[Y_1 \cup Y_2]$ has an $(\omega(G) - \omega(G[X]))$-colouring, and hence that G has an $\omega(G)$-colouring, which is a contradiction. ■

The following question may be hopeless to answer, because a proof would imply a direct argument for the skew partition conjecture (no such argument exists at present). Also, a proof of the following, together with Theorem 4.2, would yield a new proof of the SPGT.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Problem *Find a direct proof of the following: if G is a minimum Berge imperfect graph, then at least one of G or $\overline{G}$ admits no double star cutset.*

Rusu (see [50, pp. 167–183]) wrote a survey about cutsets in perfect graphs, and Reed [51] wrote a survey about skew partitions (on which this section is mostly based).

An important question about decompositions is their detection in polynomial time. For a graph with n vertices and m edges, the following decompositions can be detected in polynomial time: 2-join (in time $O(n^2m)$, see Charbit, Habib, Trotignon and Vušković [8]), homogeneous pair (in time $O(n^2m)$, see Habib, Mamcarz and de Montgolfier [37]) and skew partitions (in time $O(n^4m)$, see Kennedy and Reed [39]). Trotignon [57] showed that balanced skew partitions are NP-hard to detect, but devised an $O(n^9)$-time non-constructive algorithm that certifies whether an input Berge graph has a balanced skew partition.

5. The strategy of the proof

The main result in [13] is Theorem 5.1 below and, as we know from the earlier sections, it implies the SPGT. Its statement is the result of a long sequence of attempts by many researchers, as explained in the introduction of [13] or in [53]. A slight variant on what seems now to be the right statement was first conjectured by Conforti, Cornuéjols and Vušković [22]. They proved it in the square-free case, and some of the arguments that they discovered are essential in the strategy described below – in particular, the attachments to prisms, and the use of Truemper configurations.

Theorem 5.1 *Every Berge graph is basic, or has a* 2*-join, a complement* 2*-join, a homogeneous pair or a balanced skew partition.*

The strategy used by Chudnovsky, Robertson, Seymour and Thomas to prove Theorem 5.1 is classical in structural graph theory. It consists of identifying a 'dense' basic class and a 'sparse' basic class, as we now explain. The 'dense' basic class does not contain the obstruction (here an odd hole or antihole), but 'almost' contains it, so that if a graph G contains an induced dense subgraph H, then any vertex exterior to H must attach in a very specific way to H, either enlarging the basic graph to a bigger basic graph, or entailing a decomposition. For the sake of proving the decomposition theorem, it can therefore be assumed that this particular class of basic graphs is excluded. The process can then be iterated with a new kind of dense basic graphs. The sparse class is what remains when all dense substructures are excluded. Observe that this method is in some sense safer than proofs by induction. Finding the right induction hypothesis is time-consuming because, for each failure, one has to restart from scratch, whereas a result asserting that some dense substructure entails a decomposition is a true statement that can be used in the future, even if the strategy of the proof changes.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

For the proof of Theorem 5.1 the sparse class is formed by bipartite graphs and their complements. The dense class is more complicated. At the beginning of the proof it is formed by 'sufficiently' connected line graphs, their complements and doubled graphs. The simplest line graphs in this context are the prisms, where a *prism* is a graph made of three vertex-disjoint paths $P_1 = a_1 \dots b_1$, $P_2 = a_2 \dots b_2$ and $P_3 = a_3 \dots b_3$ of length 1 or more, for which $a_1a_2a_3$ and $b_1b_2b_3$ are triangles and no edges exist between the paths except those of the two triangles. Note that for a prism in a Berge graph, the lengths of the three paths must have the same parity. The prism is *odd* or *even* according to this parity.

To understand why prisms and their generalizations are 'dense', one can check that a Berge graph formed of a prism and one vertex not in the prism is either a larger line graph, or has a 2-join or some kind of skew partition such as a star cutset. For this purpose, it is convenient to know that Berge graphs have no pyramids, where a *pyramid* is a graph made of three paths $P_1 = a \dots b_1$, $P_2 = a \dots b_2$ and $P_3 = a \dots b_3$ of length at least 1, two of which have length at least 2, are vertex-disjoint except at a, and are such that $b_1b_2b_3$ is a triangle and no edges exist between the paths except those of the triangle and the three edges incident with a. What makes this work is that prisms are 'close to' containing pyramids, and a pyramid always contains an odd hole. One can also try to prove a similar statement for the line graph of the 2-subdivision of K_4, or similar statements with the vertex outside the prism replaced by some path with neighbours in exactly two paths of the structure, for instance. All this should lead to variants of Theorem 10.1 from [13], whose proof is easy to read since it does not rely on any technical lemmas. More on attachments to prisms is explained in Section 6 below.

In fact, a dozen dense basic graphs are considered: first, several kinds of line graphs of bipartite subdivisions of K_4; even prisms; then long prisms (*long* means that at least one of the paths has length at least 2); the *double diamond* (the graph obtained from a cycle *abcd* by adding an edge *vw* such that *v* is adjacent to *a* and *b* and *w* to *c*, *d*, and a second edge $v'w'$ with the same attachment to *abcd*); various kinds of wheels (where a *wheel* is a graph formed by a hole H together with a vertex with at least three neighbours in the hole); and antiholes of length at least 6. For each of these dense basic classes, it was proved that containing it leads to being basic or having some decomposition; the Berge graphs handled next are thus supposed not to contain this kind of induced subgraph. At the end of this process, so many induced subgraphs are excluded that the graph under consideration, or its complement, is bipartite. Needless to say, identifying this long sequence of 'dense' graphs is a *tour de force*, especially since for each of them, the technicalities are really involved.

Despite all these technicalities, the objects considered in the proof of Theorem 5.1 are very combinatorial. This leads to the following problem.

Problem *Can the above strategy be transformed into a polynomial-time algorithm whose input is an arbitrary graph G and whose output is either an odd hole, an odd antihole, or a partition of the vertices of G certifying one of the outcomes of Theorem 5.1?*

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

In [13] the global strategy of the proof is well explained at the beginning, and more about the strategy can be found in [56] and [12]. How structural methods can be used generally for classes closed under taking induced subgraphs is discussed by Chudnovsky and Seymour [14]. Prisms, pyramids and wheels are 'Truemper configurations', graphs that play an important role in many decompostion theorems (see the survey of Vušković [60]). To make a start on the above problem, the first step is the detection in polynomial time of the structures that are used in the proof (line graphs of a bipartite subdivision of K_4, even prisms, odd prisms and wheels). Apart from wheels they all can be detected in polynomial time in a Berge graph (see Maffray and Trotignon [47]). Diot, Tavenas and Trotignon proved that detecting a wheel in a Berge graph is an NP-complete problem, but for the sake of the problem above, this problem might be by-passed.

6. Book from the Proof

We are now ready to investigate some technicalities of the proof of the SPGT. A lemma due to Roussel and Rubio [52] is used at many steps in [13]; in fact, the authors of [13] rediscovered it (in joint work with Thomassen) and initially named it the *Wonderful lemma* because of its many applications.

The Roussel–Rubio lemma states that, in a sense, any anticonnected set of vertices of a Berge graph behaves like a single vertex. How does a vertex v 'behave' in a Berge graph? If a path of odd length (at least 3) has both ends adjacent to v, then v must have other neighbours in the path, for otherwise there is an odd hole. The lemma states roughly that an anticonnected set T of vertices behaves similarly: if a path of odd length (at least 3) has both ends complete to T, then at least one internal vertex of the path is also complete to T (there are exceptions, see the statement below). When $P = xx' \dots y'y$ is a path of length at least 3, a *leap* for P is a pair of non-adjacent vertices v, w such that $N(v) \cap V(P) = \{x, x', y\}$ and $N(w) \cap V(P) = \{x, y', y\}$. We denote by P^* the interior of a path P. The result of Roussel and Rubio is as follows.

Theorem 6.1 *Let T be an anticonnected set of vertices in a Berge graph G. If P is a path of odd length at least 3, vertex-disjoint from T, whose ends are T-complete, then one of the following holds:*

(1) *an internal vertex of P is T-complete;*
(2) *there is a leap for P in T;*
(3) *P has length 3, and its two internal vertices are the ends of an antipath of odd length whose interior is in T.*

We now give a corollary of Theorem 6.1 that is used constantly in [13].

Corollary 6.2 *Let T be an anticonnected set of vertices in a Berge graph G. If P is a path with odd length at least 3, vertex-disjoint from T, whose ends are T-complete and such that no internal vertex of P is T-complete, then every T-complete vertex has a neighbour in P^*.*

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Proof Let v be a T-complete vertex, and suppose that v has no neighbour in P^* (so v is not in $V(P) \cup T$). Apply Theorem 6.1 to T and P. If conclusion (2) holds, then a path of odd length with same interior as P joins the members of the leap. Together with v it forms an odd hole, giving a contradiction. If conclusion (3) holds, then the antipath of odd length can be completed to an odd antihole through v, giving a contradiction. ■

We now present a self-contained lemma from [13] in order to give some taste of the technicalities. We place this lemma in context by stating informally how line graphs are handled.

A line graph K of a bipartite graph is formed from a collection of cliques and a collection of paths linking them. If K is contained in a Berge graph G, a vertex $v \in V(G \setminus K)$ is *major* if it has many (at least 2) neighbours in each clique, and *minor* if it has neighbours in at most one of the paths or in at most one of the cliques. An important result is that every vertex is major or minor, or allows one to obtain a larger line graph. (This is not fully true, but vertices that are neither major nor minor can be considered as part of some of the paths in what is called a generalized line graph; therefore, when the generalized line graphs are properly defined, every vertex is major or minor with respect to a maximal generalized line graph.) A next step is to prove that connected (or sometimes anticonnected) components of vertices of $G \setminus K$ behave as a vertex. For more details of these components, see the proof of 10.1 in [13]; we deal with anticomponents below.

If there are no major vertices, then the graph is formed of the line graph and possibly a collection of minor components. If some of these components are attached to a clique then there is a balanced skew partition, and if some are attached to a path there is a 2-join. If there are major vertices, then it can be proved that an anticomponent of these is formed of vertices that are all major *in the same way*, meaning that they are all attached to exactly the same vertices of the cliques. An anticomponent of major vertices is therefore complete to its attachment and forms a skew cutset separating parts of the line graph. The next lemma (a variant of 7.3 of [13]) proves a statement of this form. It illustrates another breakthrough made in [13]: showing how the Roussel–Rubio lemma can be used to find skew partitions.

Theorem 6.3 *In a Berge graph G, let K be a prism with triangles $a_1a_2a_3$ and $b_1b_2b_3$ and paths $P_i = a_i \dots b_i$ for $i = 1, 2, 3$. Suppose that, for each $i = 1, 2, 3$, P_i has length at least 2. Let Y be an anticonnected set of vertices each of which has at least two neighbours in $\{a_1, a_2, a_3\}$ and at least two neighbours in $\{b_1, b_2, b_3\}$ (so Y is disjoint from K). Then at least two members of $\{a_1, a_2, a_3\}$ and at least two members of $\{b_1, b_2, b_3\}$ are Y-complete.*

Proof Suppose not; then there is an antipath with interior in Y joining two vertices both of which are in $\{a_1, a_2, a_3\}$ or in $\{b_1, b_2, b_3\}$. Let Q be a shortest such antipath and suppose up to symmetry that Q is from a_1 to a_2. Every vertex in Y is adjacent to either a_1 or a_2, so Q has length at least 3. From the minimality of Q, a_3 is

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Q^*-complete, and so is at least one of $\{b_1, b_2, b_3\}$ (say b_i). Since Q can be completed to an antihole via $a_1 b_i a_2$, it follows that Q has even length, at least 4. Let $Q = a_1 q_1 \dots q_n a_2$, let a'_i be the neighbour of a_i in P_i for $i = 1, 2$, and $P = a'_1 P_1 b_1 b_2 P_2 a'_2$. We now prove two claims.

(1) *At least one internal vertex of P is $(Q^* \setminus q_1)$-complete, and at least one internal vertex of P is $(Q^* \setminus q_n)$-complete.*

If one of b_1 or b_2 is Q^*-complete, the claim is obviously true. Otherwise, none of b_1, b_2 is Q^*-complete, so there exists a antipath from b_1 to b_2 whose interior is in Q^*, and from the minimality of Q this antipath has the same interior as Q. It follows that one of b_1, b_2 is complete to $Q^* \setminus q_1$ and the other is complete to $Q^* \setminus q_n$. This proves (**1**).

(2) *If an internal vertex of P is $(Q^* \setminus q_1)$-complete or $(Q^* \setminus q_n)$-complete, then it is Q^*-complete. If a'_1 is $(Q^* \setminus q_1)$-complete, then it is Q^*-complete.*

If an internal vertex v of P is $(Q^* \setminus q_n)$-complete then it is Q^*-complete, since otherwise, vq_nQa_1v is an odd antihole. If v is a'_1 is an internal vertex of P, and if v is $(Q^* \setminus q_1)$-complete, then v is Q^*-complete, for otherwise, vq_1Qa_2v is an odd antihole. This proves (**2**).

If both a'_1, a'_2 are Q^*-complete, then Q can be completed to an odd antihole via $a_1 a'_2 a'_1 a_2$, giving a contradiction. So, up to symmetry, we may now suppose from here on that a'_1 is not Q^*-complete. It follows by (**2**) that a'_1 is not $(Q^* \setminus q_1)$-complete.

Let x be the $(Q \setminus q_1)$-complete vertex of P closest to a'_1 along P. By (**1**), x exists and is an internal vertex of P. By (**2**), x is in fact Q^*-complete.

If $a'_1 a_1 Px$ has odd length, then by Corollary 6.2 applied to $a_1 a'_1 Px$ and $Q^* \setminus q_1$, there is a contradiction because a_3 is $(Q^* \setminus q_1)$-complete and has no neighbour in the interior $a_1 Px$. So $a_1 Px$ has even length. It follows that $xPa'_2 a_2$ has odd length. By Corollary 6.2 applied to $xPa'_2 a_2$ and $Q^* \setminus q_n$, and because of a_3, there must be an internal vertex v of $xPa'_2 a_2$ that is $(Q^* \setminus q_n)$-complete. By (**2**), v is in fact Q^*-complete. But then, $a_3 a_1 a'_1 Px$ is odd and its ends are Q^*-complete, but none of its internal vertex is Q^*-complete. This contradicts Corollary 6.2 because of v. ■

The above lemma shows how the Roussel–Rubio lemma 'generates' antiholes to provide contradictions. Maybe it can be used to investigate the structure of several subclasses of Berge graphs, such as in the next problem.

Problem *Is there a structural characterization of Berge graphs with no antihole of length at least 6?*

A proof of the Roussel–Rubio lemma is given in Maffray and Trotignon [48], where some applications are also presented. A shorter proof in the case when T is a stable set is due to Kapoor, Vušković and Zambelli (see [58] where several very simple applications are presented). For applications of the Roussel–Rubio lemma, see [13, 3.1].

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

About technicalities, the best source is of course Chudnovsky, Robertson, Seymour and Thomas [13]. This paper is well organized: the first four sections are devoted to technical lemmas that are used intensively in the sequel, but as a result of this organization, it is hard to extract a meaningful chunk, and it is difficult to appreciate how useful the technical lemmas are without knowing the sequel, and the sequel is difficult to understand without mastering the technicalities of the lemmas. The easiest parts of [13] are perhaps 15.1, Section 15 and Section 16. Section 10 is devoted to even prisms and is a self-contained chunk that has fewer technicalities than the rest of the paper, but still relies greatly on technical lemmas from the previous sections.

7. Recognizing perfect graphs

From here on we investigate results on perfect graphs that were obtained after the proof of the SPGT. Soon after this proof, in 2002, another open question was solved, when Chudnovsky, Cornuéjols, Liu, Seymour and Vušković [11] found a polynomial-time algorithm that decides whether an input graph is Berge. Note that this algorithm is independent of the SPGT: it takes an arbitrary graph G as input and in time $O(n^9)$ outputs an odd hole or antihole of G, if there is one. We give here a brief outline, from the introduction to [11]. In what follows, G is the input graph.

> ... we would like to decide either that G is not Berge, or that G contains no odd hole. (To test Bergeness, we just run this algorithm on G and then again on the complement of G.) If there is an odd hole in G, then there is a shortest one, say C. A vertex of the remainder of G is C-major if its set of neighbours in C is not a subset of the vertex set of any 3-vertex path of C; and C is *clean* (in G) if there are no C-major vertices in G. If there happens to be a clean shortest odd hole in G, then it stands out and can be detected relatively easily; and that essentially is the first step of our algorithm, a routine to test whether there is a clean shortest odd hole. The remainder of the algorithm consists of reducing the general problem to the 'clean' case that was just handled. If C is a shortest odd hole in G, let us say a subset X of $V(G)$ is a *cleaner* for C if $X \cap V(C) = \varnothing$ and every C-major vertex belongs to X. Thus if X is a cleaner for C then C is a clean hole in $G \setminus X$. The idea of the remainder of the algorithm is to generate polynomially many subsets of $V(G)$, such that if there is a shortest odd hole C in G, then one of the subsets will be a cleaner for C. If we can do that, then we delete each of these subsets in turn, thereby generating polynomially many induced subgraphs; and we know that there is an odd hole in G if and only if in one of these subgraphs there is a clean shortest odd hole. Thus we can decide whether G has an odd hole by testing whether any of these subgraphs has a clean shortest odd hole.

Theorem 7.1 *There exists an algorithm that decides whether an input graph is Berge in time* $O(n^9)$.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Let us now explain briefly the two steps from the sketch above.

The first step relies on the *shortest-path detector*, a method designed by Chudnovsky and Seymour [11]. The shortest-path detector enumerates all possible subsets of fixed size of vertices that are potentially particular vertices of some structure to be detected, and computes shortest paths between them to complete the structure until the desired structure is detected. At first glance this should fail because between the paths chords could exist (the paths might fail to be disjoint also). But a clever argument shows that this failure leads to a smaller structure similar to the one to be detected, so that when the enumeration suddenly matches the particular vertices of a *smallest* structure in the graph, the problem does not arise, and the structure is detected. This method is used in [11] to decide in time $O(n^9)$ whether a graph contains a pyramid.

The second step relies on a powerful technique discovered by Conforti and Rao [24], called the *cleaning*, which consists in 'guessing' a cleaner. Here, the way to guess the cleaner comes from some Roussel–Rubio-like lemmas, asserting that the set X of vertices to be 'guessed' that are all major vertices of some possible smallest odd hole, are all common neighbours of some restricted set S of neighbours. One can therefore enumerate all possible sets S by brute force, and for each of them nominate X as the set of S-complete vertices.

The following problem is still open.

Problem *Is there a polynomial-time algorithm to decide whether an input graph has an odd hole?*

In [11], another algorithm for recognizing Berge graphs is given. It relies on decompositions and uses Theorem 4.2. Detecting odd holes can be solved in polynomial time under the assumption that the largest size of a clique in the input graph is bounded by some constant. This was proved by Conforti, Cornuéjols, Liu, Vušković and Zambelli [21], a good source for understanding the main ideas of the recognition of Berge graphs. Simple examples of shortest path detectors can be found in Maffray and Trotignon [47].

8. Berge trigraphs

An obvious question arising from Theorem 5.1 is whether it is best possible. Is each outcome necessary? Are there outcomes that can be made more precise? A careful reader of [13], noticing that the outcome 'homogeneous pair' is obtained only once in the whole proof, might therefore wonder whether it is really necessary. Chudnovsky proved that it is not. Let us explain how.

A way of proceeding is to consider a smallest Berge graph G for which the homogeneous pair is the only outcome satisfied by G in Theorem 5.1, and to look for a contradiction. A natural idea is then to 'contract' a homogeneous pair (A, B)

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

so as to find a smaller Berge graph G', and to apply Theorem 5.1 to G' and prove that any outcome of Theorem 5.1 in G' yields an outcome in G (because G' and G are very similar). From our initial assumption, G' therefore satisfies no outcome of Theorem 5.1, and this provides a contradiction. We call this a *bootstrap method*, because it improves structural theorems somehow 'for free'. The natural way to 'contract' a homogeneous pair (A, B) of G is to replace A by a vertex a complete to $C(A)$ and anticomplete to $\overline{C}(A)$, and to replace B by a vertex b complete to $C(B)$ and anticomplete to $\overline{C}(B)$.

The bootstrap method is hard to implement in this context. The problem is with ab: should it be an edge or an antiedge of G'? Both choices lead to difficult technicalities: if ab is chosen to be a non-edge, then some skew cutset may separate a from b in G' while A and B are linked in G. If a is chosen to be adjacent to b, then the same phenomenon may happen in the complement. No neat example of these bad phenomena can be given, because (as we will see) the outcome 'homogeneous pair' is not necessary in Theorem 5.1. But any attempt of proof must face this issue and is likely to fail.

Chudnovsky's idea was to leave undecided the adjacency between a and b in G'. For this purpose, she defined *trigraphs* as graphs with edges, anti-edges, and a third kind of adjacency: a *switchable pair*. For each pair of distinct vertices v and w, vw is an edge, an antiedge or a switchable pair. A *realization* G of a trigraph T is any graph on $V(T)$ such that all edges of T are edges of G and all antiedges of T are antiedges of G (so every switchable pair is transformed into an edge or an antiedge). A trigraph is *Berge* if every realization is Berge.

We do not give the long list of definitions that extends the vocabulary of graphs to trigraphs. The key point is in the definitions of the decompositions: all the 'important' edges in these definitions are not allowed to be switchable pairs. For instance, if (X, Y) is a 2-join of a trigraph T, no switchable pair of T is from X to Y. If (X, Y) is a skew partition of a trigraph T, then X must be partitioned into sets X_1 and X_2 for which no edge and no switchable pair exist between X_1 and X_2, and Y must be partitioned into sets Y_1 and Y_2 for which no antiedge and no switchable pair exist between Y_1 and Y_2. It is easy to guess how useful this requirement is: for instance, when building G' from G as in the paragraph above, ab is defined to be a switchable pair, and so the problem that we mentioned with the skew cutset separating a from b does not arise any more.

One slight problem with this notion of trigraph is that Theorem 5.1 has to be proved again for Berge trigraphs. Chudnovsky proved several decomposition theorems similar to Theorem 5.1 for Berge trigraphs. The proof of her main one is self-contained (that is, it does not rely on the proof of Theorem 5.1) and runs to more than two hundred pages. We do not give the precise statements here since we do not give the precise definitions of decompositions and basic classes of trigraph. The precise statements of the theorems are in [10], where only parts of the proofs are given, and the complete proof is in [9]. With trigraphs, the bootstrap method works

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

smoothly and yields the following result of Chudnovsky. We translate back to graphs, since a graph is a particular trigraph.

Theorem 8.1 *Every Berge graph is basic, or has a* 2*-join, a complement* 2*-join or a balanced skew partition.*

Another potential improvement of Theorem 5.1 concerns the technical requirements in the definition of a 2-join. Some authors consider only *non-path* 2*-joins*, and in some applications it is essential to use these, because one needs to replace one side of a 2-join by a long path in proofs by induction or in recursive algorithms, and this obviously fails if the side is already a long path. Trotignon [57] investigated this question and obtained the following result starting from Theorem 8.1 and with the bootstrap method applied to graphs. It easily implies the different variants of Theorem 5.1. *Path cobipartite graphs* and *path double-split graphs* are Berge graphs obtained by subdividing edges in complements of bipartite graphs and in double split graphs respectively (a more precise definition can be given, but this one is enough here). Note that subdividing an edge in any graph creates a path 2-join. In [57], several graphs that yield a unique outcome in the following theorem of Trotignon are given.

Theorem 8.2 *If G is a Berge graph, then G is basic or has a balanced skew partition, or one of G, $\overline{G}$ is a path-cobipartite graph, or is a path-double split graph, or has a non-path* 2*-join, or has a homogeneous pair and a path* 2*-join.*

Little work has been devoted to the algorithmic aspects of trigraphs. The following problem is still open.

Problem *What is the complexity of recognizing Berge trigraphs?*

It seems now that trigraphs are an important tool for graph decompositions, as suggested by their use in claw-free and bull-free graphs (see [14]). For algorithms for Berge trigraphs, Chudnovsky, Trotignon, Trunck and Vušković [17] seems to be the only available reference.

9. Even pairs: a shorter proof of the SPGT

An *even pair* in a graph is a pair of distinct vertices such that every path between them has even length. This notion is involved in algorithmic aspects of perfect graphs and in a significantly shorter proof of the SPGT due to Chudnovsky and Seymour [15]. Meyniel [49] proved the following result.

Theorem 9.1 *A minimally imperfect graph has no even pair.*

It follows that even pairs may be an ingredient of a useful decomposition theorem for Berge graphs. Every known example of a Berge graph with no even pair contains either an odd prism or an antihole. *Artemis graphs* are such graphs with no odd holes,

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

no antiholes of length at least 5, and no prisms. The following result was proved by Maffray and Trotignon [48].

Theorem 9.2 *Every Artemis graph is a clique or has an even pair.*

Recall that given two vertices v and w in a graph G, the operation of *contracting* them means removing v and w and adding one vertex with edges to every vertex of $G \setminus \{v, w\}$ that is adjacent in G to at least one of v or w; we denote by G/vw the graph that results from this operation. Fonlupt and Uhry [29] proved that *if G is a perfect graph and if $\{v, w\}$ is an even pair in G, then G/vw is perfect and has the same chromatic number as G*. Using this idea, it is easy to turn the proof of Theorem 9.2 into a polynomial-time algorithm that colours every Artemis graph (see [42] for a faster algorithm).

The proof of the next theorem uses several ideas from the proof of Theorem 9.2, is very readable, and (as we will see) saves about fifty pages in the proof of the SPGT. An *odd wheel* (C, T) in a graph G consists of a hole C of length at least 6, and a non-empty anticonnected subset $T \subseteq V(G) \setminus V(C)$ for which at least three vertices of C are T-complete and there is a path P of C with odd length at least 3, whose ends are not T-complete and whose internal vertices are all T-complete. We say that G is *impoverished* if G is Berge and if G and $\overline{G}$ contain no odd wheel, long prism or double diamond. A *dominant pair* in G is a pair $\{v, w\}$ of non-adjacent vertices for which every other vertex of G is adjacent to at least one of v, w. The following theorem is due to Chudnovsky and Seymour [15].

Theorem 9.3 *If G is impoverished, then either G admits a star cutset, an even pair or a dominant pair, or G is a complete graph.*

It is not difficult to prove that a minimally imperfect graph has no dominant pair (see [15]), and it follows easily by Theorem 9.1 and Theorem 4.1 that every impoverished graph is perfect. Since the last fifty-five pages of the proof of the SPGT are devoted to finding a skew partition in an impoverished graph, they are no longer necessary and can be replaced by the nine pages needed to prove Theorem 9.3.

A good survey (with many conjectures) on even pairs is by Everett, de Figueiredo, Linhares Sales, Maffray Porto and Reed (see [50, pp. 67–92]). It seems that even pairs provide a cornucopia of conjectures: other conjectures on even pairs can be found in Burlet, Maffray and Trotignon [7] and Lévêque and de Werra [41]. Short proofs of the existence of even pairs in classical classes of perfect graphs (Meyniel graphs and weakly chordal graphs) can be found in Trotignon and Vušković [58].

10. Colouring perfect graphs

In the 1980s Gröstchel, Lovász and Schrijver [34] devised a polynomial-time algorithm that colours any input perfect graph. This algorithm relies on the ellipsoid method, and one may wonder whether a more combinatorial algorithm exists. Even if

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

there is no formal definition of what a *combinatorial* algorithm should be, most mathematicians agree that any algorithm that relies on graph searches and decompositions, or even on classical linear programming, can be called 'combinatorial' and that the ellipsoid method cannot be, so this question is considered as open. A more precise question is whether a fast colouring algorithm can be derived from Theorem 5.1. We describe here recent progress in this direction.

Gröstchel, Lovász and Schrijver [34] devised a polynomial-time combinatorial algorithm that colours every Berge graph, under the assumption that a polynomial-time algorithm is available for the maximum stable set (this is also explained in [17]). So our question now is whether Theorem 5.1 (or one of its variants from Section 8) helps us to find a maximum weighted stable set. It turns out that the question is easy for all basic classes – for the historical ones see [55], and for doubled graphs see [17]. Also homogeneous pairs and complement 2-joins are not hard to handle (see [17]).

For 2-joins, the situation is more complicated because a maximum stable set may overlap a 2-join in many ways. In fact, Naves (see [59]) observed that for some class of graphs decomposable by 2-joins into very simple graphs, it is NP-hard to compute a maximum stable set. We now explain how 2-joins can actually help to find stable sets *in Berge graphs*. If (X_1, X_2) is a 2-join and if A_1, B_1, C_1, A_2, B_2, C_2 are as in the definition of a 2-join, we set $\alpha_{AC} = \alpha(G[A_1 \cup C_1])$, $\alpha_{BC} = \alpha(G[B_1 \cup C_1])$, $\alpha_C = \alpha(G[C_1])$ and $\alpha_X = \alpha(G[X_1])$. The next inequalities from Trotignon and Vušković [59] tell us how stable sets and 2-joins overlap in Berge graphs.

Theorem 10.1 *If (X_1, X_2) is an odd 2-join of G, then $\alpha_C + \alpha_X \leq \alpha_{AC} + \alpha_{BC}$, and if (X_1, X_2) is an even 2-join of G, then $\alpha_{AC} + \alpha_{BC} \leq \alpha_C + \alpha_X$.*

The above lemma allows to construct blocks of decomposition of a 2-join that preserve being Berge and allow us to keep track of α (see [17] for a precise definition of the blocks). Interestingly, 2-joins are used to compute α in other classes of graphs, while they seem to be hard to use in general: for claw-free graphs see Faenza, Oriolo and Stauffer [28], and for even-hole-free graphs with no star cutsets see Trotignon and Vušković [59].

In [17], Theorem 10.1 is used to prove the following result. Many technicalities are needed, some of which come from the fact that the decomposition blocks for 2-joins that keep track of α do not preserve having no balanced skew partition. It is also proved in [17] that for Berge graphs with no balanced skew partition there exist *extreme* decompositions, in which one of the blocks is basic; these are very convenient for proofs by induction. To handle all these technicalities, it is convenient (if not essential) to work with trigraphs, but here we state the result for graphs. The following result is due to Chudnovsky, Trotignon, Trunck and Vušković [17].

Theorem 10.2 *There exists an $O(n^7)$-time algorithm whose input is a Berge graph with no balanced skew partition and whose output is a maximum weighted stable set of G and a colouring of G.*

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

Up to now, no-one has discovered how skew partitions could be handled so as to provide a polynomial colouring algorithm. One might think that a Berge graph that is uniquely decomposable with a balanced skew partition must have an even pair. To support this idea, Chudnovsky and Seymour [16] studied the structure of Berge graphs with no K_4 and no even pair. They describe them quite precisely in the next result.

Theorem 10.3 *If G is a 3-connected K_4-free Berge graph with no even pair and with no clique cutset, then one of G and $\overline{G}$ is the line graph of a bipartite graph.*

Unfortunately, it seems that this theorem does not generalize to larger values of ω, as shown by the *WKBGSF*, a graph G (discovered by Chudnovsky and Seymour, unpublished) shown in Fig. 1. Every edge of G is the middle edge of a P_4 so that, in the complement, each pair of non-adjacent vertices can be linked by a P_4, so $\overline{G}$ has no even pairs. However, G is perfect, and the balanced skew partition is the only outcome of Theorem 5.1 satisfied by G.

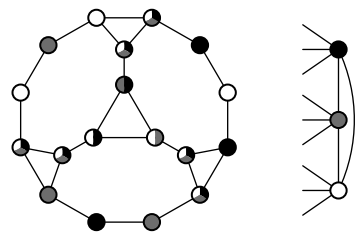


Fig. 1. The WKBGSF (Worst Known Berge Graph So Far): a vertex in the left part is adjacent to one of the three vertices in the right part if they share one of the colours black, white or grey

The following two problems may be the most important ones about perfect graphs.

Problem *Describe the structure of Berge graphs with no even pairs.*

Problem *Find a combinatorial polynomial-time algorithm that colours every Berge graph.*

Acknowledgement We thank Maria Chudnovsky, Frédéric Maffray and Irena Penev for their help.

References

1. N. Alexeev, A. Fradkin and I. Kim, Forbidden induced subgraphs of double-split graphs, *SIAM J. Discrete Math.* **26** (2012), 1–14.
2. L. W. Beineke, Characterizations of derived graphs, *J. Combin. Theory* **9** (1970), 129–135.
3. C. Berge, Färbung von Graphen, deren sämtliche bzw. deren ungerade Kreise starr sind, Technical report, *Wiss. Z. der Martin-Luther-Univ. Halle-Wittenberg, Math.-Natur. Reihe* **10**, 1961.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

4. C. Berge and V. Chvátal (eds.), *Topics on Perfect Graphs*, Ann. Discrete Math. **21**, North-Holland, 1984.
5. J. A. Bondy and U. S. R. Murty, *Graph Theory*, Graduate Texts in Math. **244**, Springer, 2008.
6. A. Brandstädt, V. B. Le and J. P. Spinrad, *Graph Classes: A Survey*, SIAM, 1999.
7. M. Burlet, F. Maffray and N. Trotignon, Odd pairs of cliques, *Graph Theory in Paris, Proc. Conf. in Memory of Claude Berge* (eds. A. Bondy *et al.*), Birkhäuser (2007), 85–95.
8. P. Charbit, M. Habib, N. Trotignon and K. Vušković, Detecting 2-joins faster, *J. Discrete Algorithms* **17** (2012), 60–66.
9. M. Chudnovsky, *Berge Trigraphs and their Applications*, Ph.D. thesis, Princeton University, 2003.
10. M. Chudnovsky, Berge trigraphs, *J. Graph Theory* **53** (2006), 1–55.
11. M. Chudnovsky, G. Cornuéjols, X. Liu, P. Seymour and K. Vušković, Recognizing Berge graphs, *Combinatorica* **25** (2005), 143–186.
12. M. Chudnovsky, N. Robertson, P. Seymour and R. Thomas, Progress on perfect graphs, *Math. Programming (B)* **97** (2003), 405–422.
13. M. Chudnovsky, N. Robertson, P. Seymour and R. Thomas, The strong perfect graph theorem, *Ann. of Math.* **164** (2006), 51–229.
14. M. Chudnovsky and P. Seymour, Excluding induced subgraphs, *Surveys in Combinatorics, 2007*, London Math. Soc. Lecture Notes **346** (2007), 99–119.
15. M. Chudnovsky and P. Seymour, Even pairs in Berge graphs, *J. Combin. Theory (B)* **99** (2009), 370–377.
16. M. Chudnovsky and P. Seymour, Three-colourable perfect graphs without even pairs, *J. Combin. Theory (B)* **102** (2012), 363–394.
17. M. Chudnovsky, N. Trotignon, T. Trunck and K. Vušković, Coloring perfect graphs with no balanced skew-partitions, submitted.
18. V. Chvátal, On certain polytopes associated with graphs, *J. Combin. Theory (B)* **18** (1975), 138–154.
19. V. Chvátal, Star-cutsets and perfect graphs, *J. Combin. Theory (B)* **39** (1985), 189–199.
20. V. Chvátal and N. Sbihi, Bull-free Berge graphs are perfect, *Graphs Combin.* **3** (1987), 127–139.
21. M. Conforti, G. Cornuéjols, X. Liu, K. Vušković and G. Zambelli, Odd hole recognition in graphs of bounded clique size, *SIAM J. Discrete Math.* **20** (2006), 42–48.
22. M. Conforti, G. Cornuéjols and K. Vušković, Square-free perfect graphs, *J. Combin. Theory (B)* **90** (2004), 257–307.
23. M. Conforti, G. Cornuéjols and K. Vušković, Decomposition of odd-hole-free graphs by double star cutsets and 2-joins, *Discrete Appl. Math.* **141** (2004), 41–91.
24. M. Conforti and M. R. Rao, Testing balancedness and perfection of linear matrices, *Math. Programming* **61** (1993), 1–18.
25. G. Cornuéjols, *Combinatorial Optimization: Packing and Covering*, CBMS-NSF Regional Conf. Series in Applied Math. **74**, SIAM, 2001.
26. G. Cornuéjols and W. H. Cunningham, Composition for perfect graphs, *Discrete Math.* **55** (1985), 245–254.
27. E. Diot, S. Tavenas and N. Trotignon, Detecting wheels, *Analysis Discrete Math.* **8** (2014), 111–122.
28. Y. Faenza, G. Oriolo and G. Stauffer, An algorithmic decomposition of claw-free graphs leading to an $O(n^3)$-algorithm for the weighted stable set problem, *SODA* (2011), 630–646.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

29. J. Fonlupt and J. P. Uhry, Transformations which preserve perfectness and h-perfectness of graphs, *Bonn Workshop on Combinatorial Optimization* (eds. A. Bachem, M. Grötschel and B. Korte), *Ann. Discrete Math.* **16** (1982), 83–85.
30. D. R. Fulkerson, Anti-blocking polyhedra, *J. Combin. Theory (B)* **12** (1972), 50–71.
31. T. Gallai, Transitiv orientierbare Graphen, *Acta Math. Acad. Sci. Hungar.* **18** (1967), 25–66.
32. G. S. Gasparian, Minimal imperfect graphs: a simple approach, *Combinatorica* **16** (1996), 209–212.
33. M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Elsevier, 2004.
34. M. Gröstchel, L. Lovász and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, Springer-Verlag, 1988.
35. A. Gyárfás, Problems from the world surrounding perfect graphs, *Zastowania Mat. Appl. Math.* **19** (1987), 413–441.
36. A. Gyárfás, Z. Li, R. C. S. Machado, A. Sebő, S. Thomassé and N. Trotignon, Complements of nearly perfect graphs, *J. Combin.* **4** (2013), 299–310.
37. M. Habib, A. Mamcarz and F. de Montgolfier, Algorithms for some H-join decompositions, *LATIN 2012*, 446–457.
38. S. Hougardy, Classes of perfect graphs, *Discrete Math.* **306** (2006), 2529–2571.
39. W. S. Kennedy and B. Reed, *Fast Skew Partition Recognition*, Lecture Notes in Computer Science **4535** (2008), 101–107.
40. P. G. H. Lehot, An optimal algorithm to detect a line graph and output its root graph, *J. Assoc. Comp. Mach.* **21** (1974), 569–575.
41. B. Lévêque and D. de Werra, Graph transformations preserving the stability number, *Discrete Appl. Math.* **160** (2012), 2752–2759.
42. B. Lévêque, F. Maffray, B. Reed and N. Trotignon, Coloring Artemis graphs, *Theor. Comp. Sci.* **410** (2009), 2234–2240.
43. L. Lovász, A characterization of perfect graphs, *J. Combin. Theory (B)* **13** (1972), 95–98.
44. L. Lovász, Normal hypergraphs and the perfect graph conjecture, *Discrete Math.* **2** (1972), 253–267.
45. L. Lovász, Perfect graphs, *Selected Topics in Graph Theory* (eds. L. W. Beineke and R. J. Wilson), Academic Press (1983), 55–87.
46. F. Maffray, *Fast recognition of doubled graphs*, Technical Report, *Les Cahiers Leibniz* **202**, 2013.
47. F. Maffray and N. Trotignon, Algorithms for perfectly contractile graphs, *SIAM J. Discrete Math.* **19** (2005), 553–574.
48. F. Maffray and N. Trotignon, A class of perfectly contractile graphs, *J. Combin. Theory (B)* **96** (2006), 1–19.
49. H. Meyniel, A new property of critical imperfect graphs and some consequences, *Europ. J. Combin.* **8** (1987), 313–316.
50. J. L. Ramírez Alfonsín and B. A. Reed (eds.), *Perfect Graphs*, Wiley Interscience, 2001.
51. B. A. Reed, Skew partitions in perfect graphs, *Discrete Appl. Math.* **156** (2008), 1150–1156.
52. F. Roussel and P. Rubio, About skew partitions in minimal imperfect graphs, *J. Combin. Theory (B)* **83** (2001), 171–190.
53. F. Roussel, I. Rusu and H. Thuillier, The Strong Perfect Graph Conjecture: 40 years of attempts and its resolution, *Discrete Math.* **309** (2009), 6092–6113.
54. N. D. Roussopoulos, A max $\{m, n\}$ algorithm for determining the graph H from its line graph G, *Inform. Proc. Letters* **2** (1973), 108–112.
55. A. Schrijver, *Combinatorial Optimization, Polyhedra and Efficiency*, Springer, 2003.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press

56. P. Seymour, How the proof of the strong perfect graph conjecture was found, *Gazette des Math.* **109** (2006), 69–83.
57. N. Trotignon, Decomposing Berge graphs and detecting balanced skew partitions, *J. Combin. Theory (B)* **98** (2008), 173–225.
58. N. Trotignon and K. Vušković, On Roussel–Rubio-type lemmas and their consequences, *Discrete Math.* **311** (2011), 684–687.
59. N. Trotignon and K. Vušković, Combinatorial optimization with 2-joins, *J. Combin. Theory (B)* **102** (2012), 153–185.
60. K. Vušković, The world of hereditary graph classes viewed through Truemper configurations, *Surveys in Combinatorics 2013*, London Math. Soc. Lecture Notes **409** (2013), 265–326.

https://doi.org/10.1017/CBO9781139519793.010 Published online by Cambridge University Press