How the proof of the strong perfect graph conjecture was found

Paul Seymour[1]
Princeton University, Princeton NJ 08544

December 15, 2002; revised April 7, 2006

[1]Supported by ONR grant N00014-01-1-0608, and NSF grant DMS-0070912.

Abstract

In 1961, Claude Berge proposed the “strong perfect graph conjecture”, probably the most beautiful open question in graph theory. It was answered just before his death in 2002. This is an overview of the solution, together with an account of some of the ideas that eventually brought us to the answer.

1 Introduction

In the summer of 2002, Maria Chudnovsky, Neil Robertson, Robin Thomas and I announced that we had settled the strong perfect graph conjecture. The proof is about 150 pages, and an overview can only give a vague idea of it. But I gave a talk on it in September 2002, and Vasek Chvátal told me afterwards that if I could put the talk on paper, he would like to have it for the slim and beautiful volume that he and Adrian Bondy had in mind in honour of Claude Berge. So here it is, more or less. What he wanted was an account of the ideas that led us to the approach that finally worked (and some that did not), so I beg the reader's indulgence in this respect; rather a lot of this paper is devoted to what we were thinking and when, rather than to the mathematics itself.

Three of us started on this problem in January 2000, and worked just about full-time on it (joined by Maria as my student from summer 2001 on) for two-and-a-half years, until we announced the result in May 2002. The start date is very clear, for a reason. In 1999 Peter Sarnak (who was chairman of the Princeton math department at that time) came to me and suggested that we try for a grant from AIM, the American Institute of Mathematics. This is a private organization, funded by Fry Electronics, a computer retail chain in California. Every year, AIM supports a group of about three people to work full-time on some project; or more precisely, AIM has an arrangement with some university departments, that AIM supports, say, two people to come to the department as visitors, and in return the department frees some faculty member from teaching. Anyway, this is a great way for a group of three people to work full-time on a project for six months or a year. Three of us, Robertson, Thomas and I, had worked together on several earlier projects, and wanted to do it again, so it seemed an ideal chance for us. The catch is, AIM doesn't want the normal kind of proposal; they like proposals to work on "high-profile" problems, with a correspondingly small chance of success (at least that was what Sarnak told me.) So what should we propose to do?

There was always the Hadwiger conjecture, that graphs not contractible to K_{t+1} are t-colourable. For $t = 4$ that is equivalent to the four-colour theorem, and in [7] we had gone one step further and proved it for $t = 5$. People kept asking us "What about $t = 6$?", and on a grant proposal this would look eminently plausible — except that we felt we had already reached our limits doing $t = 5$, and had no real hope of getting any further.

The other big open question in graph theory was Berge's strong perfect graph conjecture. The group around Gérard Cornuéjols seemed to be making progress with this — they thought that every Berge graph could be built from a few basic classes by some reasonable constructions — and if that could be true, then we ought to be in on it, because that kind of theorem was what we did best, we had had a lot of successes on other problems with the same approach. It seemed a natural topic for us, except for the small difficulty that none of us knew anything about perfect graphs.

Those were two world-class problems, except that on the first we had no hope, and on the second we had no experience. So we decided to put down both problems on the proposal — half a page on perfect graphs and half a page on Hadwiger. (That was a great grant proposal to write, incidentally; it really was about a page long in total.) And AIM accepted the proposal, and gave us the grant. I am dwelling on this at some length, because I am very grateful to AIM; they are just a small private organization, but they gave us a lot of help, and without their grant we might never have had the incentive or the opportunity to get started on perfect graphs. AIM funded us through the essential phase of trying out dozens of bad ideas that don't work, and I'd like them to have some credit for it.

Anyway, in January 2000 Neil and Robin arrived in Princeton for a six-month visit supported by AIM, so now we had to bite the bullet and actually start thinking about perfect graphs.

2 Berge and perfection

The *chromatic number* $\chi(G)$ of a graph G is the minimum number of colours needed to colour the vertices so that every two adjacent vertices get different colours. Determining the chromatic number of a graph is a notoriously hard problem, in practice and in theory, and it's hard even to get decent lower bounds. One obvious lower bound is $\omega(G)$, the size of the biggest clique (= complete subgraph) of G; because if a graph has six pairwise adjacent vertices, you certainly need at least six colours to colour it. And there are some interesting graphs for which $\chi(G) = \omega(G)$, for instance the following:

- **Bipartite graphs.** If G is bipartite then $\chi(G) = \omega(G) = 2$, unless G has no edges. (This is admittedly not a very deep fact.)

- **Complements of bipartite graphs.** If H is bipartite and $G = \overline{H}$, then $\chi(G)$ is the size of the smallest set of vertices and edges of H whose union is $V(H)$; and $\omega(G)$ is the maximum size of a stable set of vertices of H. That these are equal is a theorem of König.

- **Line graphs of bipartite graphs.** If H is bipartite and G is its line graph $L(H)$, then $\chi(G)$ is the edge-chromatic number of H, the minimum number of colours needed to colour the edges so that any two edges that meet get different colours; and $\omega(G)$ is the maximum number of edges that pairwise share an end, which (since H is bipartite) is the same as the biggest vertex-degree in H. That these are equal is the well-known theorem of König on edge-colouring bipartite graphs.

- **Complements of line graphs of bipartite graphs.** If H is bipartite, and $G = \overline{L(H)}$, then $\chi(G)$ is the minimum number of vertices of H that together hit all edges, and $\omega(G)$ is the maximum size of a matching in H. That these are equal is an even more well-known theorem of König, the min-max characterization of maximum matchings in bipartite graphs.

- **Comparability graphs.** If $(P, \leq)$ is a partial order, let G be the graph with $V(G) = P$, in which two vertices are adjacent if one is less than the other in the partial order. Then $\chi(G)$ is the minimum number of antichains in the poset with union P, and $\omega(G)$ is the maximum size of a chain. That these are equal is a rather easy theorem, that can be proved by partitioning the elements of P into antichains by their height. But still, it's another example.

- **Complements of comparability graphs.** If $(P, \leq)$ is a partial order, let G be the complement of the graph defined above. Then $\chi(G)$ is the minimum number of chains with union P, and $\omega(G)$ is the maximum size of an antichain. That these are equal is Dilworth's theorem, and is much less obvious.

And there are many more examples, but that is enough for our purposes. (Note that the fifth class contains the first, and the sixth contains the second, so one might not bother to list the first and second classes. As it turns out, that would be a mistake; the first four classes are the important ones.)

These could perhaps all claim to be interesting, but not every graph satisfying $\chi(G) = \omega(G)$ is interesting. For instance, take any graph with at most 100 vertices, and take the disjoint union of this and K_{100}. The graph we construct satisfies $\chi(G) = \omega(G)$, but it is definitely not interesting in general. We would like to somehow define a class of graphs satisfying $\chi(G) = \omega(G)$, containing the interesting graphs mentioned above, and not the uninteresting ones.

A nice way to do this is Berge's definition [1]; we make the property hereditary. Let us say a graph G is *perfect* if $\chi(H) = \omega(H)$ for every induced subgraph H of G. All the graphs listed above are perfect, except for the one we wanted to exclude, so that works. And this turns out to be an inspired definition. Perfect graphs have pretty theoretical properties; for instance, Lovász proved that the complement of any perfect graph is also perfect (which explains why the examples listed above all come in complementary pairs). And they have important connections with linear and integer programming; for instance, if A is a (0,1)-matrix with n columns, all vertices of the polyhedron $\{\mathbf{x} \in R^n : \mathbf{x} \geq \mathbf{0}, A\mathbf{x} \leq \mathbf{1}\}$ are integral if and only if A arises from a perfect graph (more precisely, the non-dominated rows of A form the incidence matrix of the maximal cliques of a perfect graph). Perfect graphs have come to be recognized as having a natural place in the world.

Not all graphs satisfy $\chi(G) = \omega(G)$. For instance, if G is a cycle of odd length (at least five), then $\chi(G) = 3$ and $\omega(G) = 2$; or if G is the complement of an odd cycle of length $2n + 1 \geq 5$, then $\chi(G) = n + 1$ and $\omega(G) = n$. So if G is perfect then no induced subgraph is an odd hole or antihole (a *hole* means an induced subgraph which is a cycle of length at least four, and an *antihole* is the complement). Let us say G is *Berge* if it contains no odd hole or antihole; so every perfect graph is Berge. In 1961, Berge proposed two excellent conjectures about perfect graphs:

- (The weak perfect graph conjecture, later Lovász's theorem.) *The complement of every perfect graph is perfect.*
- (The strong perfect graph conjecture, SPGC.) *G is perfect if and only if G is Berge.*

We have now proved the SPGC [3].

3 Decomposition theorems

We spent a year or more experimenting with different ideas, trying to get familiar with methods that might work and those that would not. Very early we realized that our only chance of success was to try to prove that every Berge graph is either of some familiar type, or admits some kind of decomposition. There were several reasons why we focussed on this approach — it was what came naturally to us, and we had been quite successful at this kind of theorem, in other contexts; Conforti, Cornuéjols and Vušković had been looking at this approach and thought it was promising; and yet it was a relatively new approach, while the only other approach that we knew (looking at the linear programming implications for a minimum counterexample) had been exhaustively studied, and we didn't think we could contribute any more to that. (Incidentally, Chvátal hotly denies, with references, that trying for a decomposition theorem was a new approach. But at least we thought it was new.)

What do we want to prove? Ideally we would like a set of construction rules such that every Berge graph (and no non-Berge graph) can be constructed via these rules starting from some basic types. Then we could prove the SPGC by assuming there is a smallest counterexample G to the SPGC, and applying this construction theorem to it. We would have to check that basic graphs are perfect, and any graph built from smaller graphs by means of our constructions could not be a minimum counterexample to the SPGC.

As a warm-up, we tried to find a construction for the Berge graphs with no K_4 subgraph. Tucker had already shown that all such graphs are perfect [12], but the structure of such graphs was not known, and we thought this would be a good place to try for construction theorems before we started on the general problem. But it was terrible, we got nowhere — we could find nothing to do with these graphs that we couldn't equally well do with general Berge graphs. Excluding K_4 seemed more like a distraction than a help; it seemed to mask what was really going on, without giving much simplification. (At least that was our excuse for eventually giving up on this problem.) Even now there is no construction known for these graphs.

So we had to reduce our expectations, and try for something weaker. To prove the SPGC it would be enough to show that for every Berge graph, either it lies in some basic class, or it has some feature that proves it is not a minimum counterexample to the conjecture; and this turns out to be a much more fruitful line of research. The advantage is that the "feature" does not have to be a way of constructing the graph from smaller graphs. In fact we used four types of feature; three of them do correspond to such constructions, while the fourth is a special kind of partition of the vertex set (a "balanced skew partition") that cannot occur in a minimum counterexample to the SPGC, but does not correspond to a construction. All four correspond to *decompositions*, ways of breaking up the graph, so from now on let us confine ourselves to features that are decompositions.

At that stage we could still be a little vague about what we wanted to prove. Since we were not trying to do an inductive proof, the precise statement of the theorem didn't really matter; we could go by trial and error, and try to discover the right formulation by experiment. There was a conjecture of Cornuéjols et al. [5], that every Berge graph either is one of the first four types listed above (bipartite graphs, line graphs of bipartite graphs, and their complements), or admits one of a few kinds of decomposition. Certainly we had no objection to calling these four classes basic, and if experience showed that we needed more basic classes we could add them later.

What about the decompositions? Those were more open to debate; it was crucial only that whatever decompositions we used, we had to be able to prove that the smallest counterexample to the SPGC did not admit any such decompositions. Our final list of decompositions was a little different from that proposed in [5]. For aesthetic reasons it would be best to keep the list of decompositions as small as possible, but this was not a prime consideration; all we really wanted to do was to prove the SPGC, so if more decompositions seemed to be useful we would happily throw them into the pot. But we found that all we needed were the same few kinds of decomposition, over and over again. The first is a "2-join". Let A, B be a partition of $V(G)$, let A_1, A_2 be disjoint subsets of A, and let B_1, B_2 be disjoint subsets of B. If for $i = 1, 2$, every vertex of A_i is adjacent to every vertex of B_i, and there are no other edges between A and B, we say that G admits a *2-join.* (Not quite, because we want a 2-join to be a useful thing to have. So we had better insist that A, B are nonempty, and perhaps insist that $|A|, |B| \geq 3$ or something like that. The details are not important here.) 2-joins work well for this problem. There is a corresponding construction (G admits a 2-join if and only if it can be built from two smaller graphs by piecing them together appropriately), and this construction preserves both being Berge and being perfect (that is, the two smaller graphs are perfect if and only if the big graph is); so if there is a smallest counterexample to Berge's conjecture, it cannot admit a 2-join. (Yes, there also exist 1-joins, but it turns out that we don't need them.) In addition, a graph might admit a *complement 2-join*, that is, its complement might admit a 2-join.

In the paper [3] we used — once only — a third kind of decomposition, called an *M-join.* I won't bother to talk about it here, because in her thesis, Maria Chudnovsky showed that we don't need M-joins; the theorem is still true without them. (They are a mild variant of 2-joins, with the same desirable properties.)

But these are not enough; there are Berge graphs that do not admit 2-joins, complement 2-joins or M-joins that don't fall into any familiar classes. We need at least one more kind of decomposition. What seemed the natural candidate was Chvátal's concept of a *skew partition*, that is, a partition (A, B) of $V(G)$ such that A is not connected (that is, the restriction of G to A is not connected) and B is not anticonnected (that is, the restriction of $\overline{G}$ to B is not connected). In a way these are prettier than 2-joins, because a skew partition of G is also a skew partition of $\overline{G}$. And they are useful — for instance, every comparability graph is either bipartite, or a clique, or admits a skew partition (take a point u of the poset which is neither a maximum nor a minimum point, and which is incomparable with some other point v; then the set of its neighbours is not anticonnected, and separates u from v.)

But there was a big difficulty with skew partitions; they didn't satisfy the one crucial condition. It was a heavily-investigated open conjecture (due to Chvátal) that there could be no skew partition in a minimum counterexample to the SPGC. This was a great blight on our hopes, for a long time, because we couldn't prove Chvátal's conjecture, and it does seem that skew partitions are what you get; Berge graphs that do not lie in the basic classes seem to be full of skew partitions, and it's difficult to see anything better in them.

Then we had a good thought — "balanced" skew partitions. A skew partition (A, B) is *balanced* if every induced path (of length at least 2) with ends in B and interior in A has even length, and vice versa for antipaths. These have the same advantages as general skew partitions (that is, they seem to be present in Berge graphs when you need them), and the big disadvantage goes away; we could prove that no minimum counterexample to the SPGC admits a balanced skew partition.

There is another unpleasant feature of skew partitions — unlike 2-joins and M-joins, there is no corresponding composition operation (at least, not one that I would call a composition). This feature does not disappear when we move to balanced skew partitions, but it does not matter for the application to the SPGC. It *does* matter if we want to use the decomposition theorem to get a polynomial-time algorithm to test if a graph is perfect, but that is another story. (Maria and I eventually found such an algorithm, not using the decomposition theorem at all, which was a big surprise [2].)

That was our working conjecture, for more than a year; that for every Berge graph, either it belongs to one of the four basic classes, or it admits a 2-join, complement 2-join, M-join, or balanced skew partition. (Actually, I am not sure where this conjecture originated. It seemed to come to me from nowhere in a moment of inspiration, but apparently Cornuéjols' group made it first, and I had been listening to talks by Cornuéjols, so perhaps I just regurgitated it.) Eventually we worked out what looked like a proof, and for a couple of months last summer we were claiming to have proved it; which was a little embarrassing, because the conjecture is false. (There was an oversight in our proof — we never bothered to verify that one of the skew partitions we found was balanced, and it wasn't.)

We found counterexamples that, for want of a better name, we called *double split graphs*. (They are obtained from split graphs by doubling the vertices.) A *split graph* is a graph whose vertex set can be partitioned into a stable set and a clique. Let H be a split graph, with stable set A and clique B. For each $v \in V(H)$ let v_1, v_2 be two new vertices; the graph G we are constructing will have vertex set these new vertices. For each $a \in A$ let a_1a_2 be an edge of G, so the set of new vertices corresponding to A induces a matching in G. Do the same in the complement for B; that is, all the new vertices corresponding to B are mutually adjacent except for the nonedges b_1b_2 for $b \in B$. For each $a \in A$ and $b \in B$, if $ab \in E(H)$ let a_1b_1, a_2b_2 be edges of G, and otherwise let a_1b_2, a_2b_1 be edges of G. Then the graph G is a double split graph.

It is easy to see that double split graphs are perfect, and in general they do not fall into the four basic classes, and don't admit 2-joins or complement 2-joins or M-joins. They also do not admit balanced skew partitions (they admit skew partitions, but not balanced ones). So we had better call them a fifth basic class, and hope that they are not precursors of anything worse.

Happily, the error in our proof could be fixed now; at the cost of this fifth class, it all works again. We proved the following.

3.1 *Let G be a Berge graph. Then either:*

- *G belongs to a basic class; that is, either*
  - *G or its complement is bipartite, or*
  - *G or its complement is a line graph of a bipartite graph, or*
  - *G is a double split graph,*

  or
- *G admits one of the following:*
  - *a 2-join,*
  - *a complement 2-join,*
  - *an M-join,*
  - *a balanced skew partition.*

If we can prove 3.1, then the SPGC follows; for suppose the SPGC is false. Then there is a smallest counterexample G. But G does not lie in any of the basic classes, since those graphs are all perfect; and G does not admit any of the four types of decomposition listed in 3.1; and so 3.1 is violated, a contradiction.

4 Getting started

How can we set about proving something like 3.1? There are several other theorems in graph theory (if you think a matrix is a bipartite graph with weights on the edges) of the same flavour, for instance:

- To prove that graphs not contractible to K_5 are 4-colourable, Wagner [14] proved that all such graphs could be constructed from planar graphs and one exceptional graph, by glueing them together at low order cutsets.
- A matrix of reals is *totally unimodular* if every square submatrix has determinant $1, 0$ or -1. To get a polynomial algorithm to test if a matrix is totally unimodular, I proved [11] that every such matrix can be constructed from matrices arising from paths in trees, and their transposes, and one exceptional matrix, by assembling them in near-diagonal constructions.

- Tutte [13] conjectured that any 3-connected cubic graph not containing the Petersen graph as a minor was 3-edge-colourable. To prove that, we proved [9] that it was enough to show it for cubic graphs with high cyclic connectivity, and that every such graph could be constructed from cubic graphs that were almost planar, by glueing them together on small edge-cutsets.

- Pólya [6] asked, given a square $(0, 1)$-matrix, when can we change some 1's to -1's so that the determinant of the new matrix equals the permanent of the original? To get a polynomial algorithm to decide this, we proved [8] that every such matrix could be constructed from matrices whose associated bipartite graphs (rows versus columns) were planar, and one exceptional graph, by glueing them together appropriately.

These are just a few examples, but there are two things to be learned from them. One is that sometimes it helps to find an explicit construction for a set of graphs, in order to solve some other question. The second point is, the proofs of most of these theorems follow the same pattern, and perhaps we can follow the pattern again to get decomposition theorems for Berge graphs.

We are trying to prove that all graphs with some property either lie in a few basic classes, or admit one of a few types of decomposition. The paradigm proof is, take some carefully-selected small graph H with the property, that does not fall into any of the basic classes. Then it must admit one of the decompositions. Now examine how H can be contained in a larger graph G with the property, and prove that the decomposition of H extends to a decomposition of G. Then we have proved what we wanted for all graphs G with the property, that contain H; so now we focus on the graphs that don't contain H, pick some new H', and do it again. Eventually the process stops because the graphs that remain can be proved to lie in one of the basic classes. We decided to try to apply the same proof method to Berge graphs.

So what is the right graph H to start with? We want a Berge graph that we can prove induces a decomposition in any bigger Berge graph containing it. And very likely we are going to need several such graphs, not just one. This was a little ambitious, because at that time there were no such theorems known; there were no results at all that said "Any Berge graph containing this particular graph admits a decomposition". And we went most of a year without finding any.

Where should we look? Here we made a big mistake. A *wheel* in G means a subgraph of G consisting of a hole C of length at least six, plus a *hub* X, a nonempty anticonnected set that has several common neighbours on the hole. (It is negotiable exactly how many; but at least three, and at least one edge of common neighbours, and preferably more.) Wheels are very interesting, for several reasons:

- They do not lie in any of the basic classes (with a couple of exceptions).

- Conforti, Cornuéjols, Vušković and Zambelli were proving a sequence of steadily improving theorems (extensions of [4]) that said that any counterexample to the strong perfect graph conjecture contains a wheel or something similar. So if we could prove that wheels induce decompositions, that would come close to proving the SPGC.

- Wheels almost do what we want, in the following sense. One can show that in any Berge wheel (C, X), there is an even number of edges of C with both ends X-*complete* (that is, adjacent to all vertices in X) – call such an edge an X-*complete edge*. So the X-complete edges make an edge-cut of C, dividing $V(C)$ naturally into two parts (say left and right). This in general gives a skew partition of the wheel; the set of left and right vertices that are not incident with X-complete edges is not connected, and its complement (X and the ends of the X-complete edges) is not anticonnected. We need to show that however we enlarge H to a larger Berge graph G, this skew partition of H extends to a skew partition of G. This is not true, but it's almost true; enough of the cases work that one feels there is a theorem hiding here somewhere.

We put a huge amount of effort (about a year of full-time work) into trying to prove that wheels in Berge graphs induce skew partitions, but it was a mistake, it was the wrong way to go; this is not the correct stage of the proof to try to handle wheels.

Nevertheless, thinking about wheels was not a complete waste of time, because it led us to a result that became a fundamental lemma used many times in the final proof, although it does not concern wheels. Suppose that we have an anticonnected set X of vertices in a Berge graph, and we are trying to show that X and its common neighbours separate two given vertices, say L and R (thereby producing a skew partition). We need to prove that every path from L to R either meets X or contains an X-complete vertex. To prove something like this, we need a lemma that says that in the appropriate circumstances, a path can be guaranteed to contain an X-complete vertex. There is a lemma of this kind, and very often it was exactly what we needed, the following. (The *interior* of a path means the subpath obtained by deleting its ends.)

4.1 *Let G be Berge, and let $X \subseteq V(G)$ be anticonnected. Let P be an induced path in $G \setminus X$, of odd length at least five, with both ends X-complete. Then either some internal vertex of P is X-complete, or there is an induced path between two members of X with the same interior as P.*

This was the only decent result about Berge graphs that we proved in the first six months, and we were rather proud of it (it's actually a very elegant result, though one might need a little time to appreciate it). But then we saw an unpublished manuscript of Roussel and Rubio (later published as [10]) with our pet result in it; they had found it before us. So much for the fruits of our six months of AIM-supported labour! We were starting to think we should have worked on the Hadwiger conjecture after all.

Fortunately the AIM people had more faith than we did, and they were kind enough to support Robin to stay for another six months (supported jointly by Princeton). Neil had to go back home, but Robin and I kept on battling with wheels, trying to show that they induce skew partitions. We still couldn't do it, but that autumn we did hit on a significant fact. Let (C, X) be a wheel H say, where C is a hole of length six, with vertices $c_1, \ldots, c_6$ in order; and c_1, c_2, c_4, c_5 are X-complete and c_3, c_6 are not. Suppose H is contained in some larger Berge graph G. (This is the simplest interesting example of a wheel, so we thought about it a lot.) We may assume that X is maximal; so any vertex that is adjacent to c_1, c_2, c_4, c_5 either already belongs to X, or is X-complete (because otherwise we could add it to X, keeping the set anticonnected). Say c_2, c_3, c_4 are the *right* vertices, and the other three are the *left* vertices. To show that the natural skew partition of H extends to a skew partition of G, we would like to show that for any connected subset F disjoint from $V(H)$, if no vertices in F are X-complete, then not both c_3, c_6 have neighbours in F. If this is false, then there is an induced path between c_3, c_6 with interior in F, and there is a minimal subpath P of this path so that both left and right vertices have neighbours in P. This subpath seems to be where the action is, so we investigated it closely. There are three possible kinds of such paths P, and all three are interesting. Let P have vertices $p_1\text{-}\cdots\text{-}p_k$ say. The difference between the cases lies in the edges between $V(C)$ and $V(P)$.

- p_1 is adjacent to c_1, c_6; p_k is adjacent to c_3, c_4; and there are no other edges between C and P. And some $x \in X$ has no neighbours in $C \cup P$ except c_1, c_2, c_5, c_6. In this case, the subgraph of G induced on $V(C \cup P) \cup \{x\}$ is a line graph, the line graph of a bipartite subdivision of K_4. We get this problem because the wheel we started with was still in a basic class, it wasn't complicated enough; and we thought we could get around it by starting with a better wheel. (A mistake! Producing the line graph of a bipartite subdivision of K_4 is a good thing, not a bad thing; we should have valued it.)

- p_1 is adjacent to c_1, c_5 and possibly c_6; p_k is adjacent to c_2, c_4 and possibly c_3; and there are no other edges between C and P. This was the big problem; the same kind of thing happened for more complicated wheels as well. We wrestled with this for ages, but couldn't get around it.

- p_1 is adjacent to c_1; p_k is adjacent to c_2; and there are no other edges between C and P. This was just an annoyance, we thought; it's not really violating the skew partition, it's just making our lives difficult. But here is a nice result: if this happens, then G admits a skew partition, in a different place! One can show that X and its common neighbours separate $V(P)$ from $\{c_3, c_6\}$.

This last was a great theorem for us, because it was the first of its kind; the first result that said that if a Berge graph contains a certain subgraph, then it admits a skew partition. For me at least, this was a great morale booster; it showed that there were theorems waiting to be discovered, of the kind we needed for our strategy to work. (This result is in the paper [3], though only its parents would recognize it now.)

5 A better approach

Anyway, wheels were a learning experience, but not as profitable as we had hoped. The AIM largesse finally dried up, Robin went home at the end of 2000, and we continued wheel-battling separately, generating huge files of notes on myriads of cases, but not really getting anywhere. Neil came back to Princeton for a month around the end of August 2001, and we were joined by my graduate student Maria Chudnovsky, who was ready to start research for her PhD. In desperation we finally gave up on wheels and started trying other things.

Earlier that spring, we had had an interesting conversation with Kristina Vušković, during the Baton Rouge conference. (Nearly everyone else was off on the Mardi Gras excursion, and they had a great time viewing the attractions there, so I'm glad the conversation with Kristina was worthwhile.) She was convinced that there was something going on, to do with line graphs of bipartite graphs (say H) contained in bigger Berge graphs G; that if you grow H inside G, keeping it decently connected, then when it stops being a line graph, it induces a decomposition of G. The seeds of that conversation took some time to germinate, but finally in August we started thinking about maximal line graphs. More precisely, let G be Berge, and let H be an induced subgraph of G, maximal with the property that H is a line graph of a bipartite subdivision of a 3-connected graph; then how can the remainder of G attach to H? This turned out to be a great question — suddenly there was a theorem that wanted to be proved, which made a big change from our wheel experiences. This is what we showed. Let G be Berge, and let J be a bipartite subdivision of a 3-connected graph J_0, with J_0 as large as possible such that the line graph H of J is an induced subgraph of G. So J has *branch-vertices* of degree ≥ 3, and others vertices of degree 2, the latter falling into paths between branch-vertices, called *branches*. The branch-vertices of J become cliques (called *potatoes*) in the line graph H, and the branches become paths of H between potatoes. The first question is, suppose v is a vertex of G not in $V(H)$; what can its set of neighbours in H look like? It turns out that there are only three possibilities for v. Say v is *minor* if all its neighbours in H lie in one potato, or in one of the paths between potatoes; and *major* if it has at most one non-neighbour in every potato. It would be nice if every v was either minor or major, but there is a third possibility; for instance, v could have exactly the same neighbours as some vertex of a potato. To get rid of this third kind of vertex, let us add any such vertex to H. When we do this, the paths between potatoes thicken up into *strips* between potatoes. Consequently H stops being a line graph, but this does not matter much; H still has an overall shape determined by J_0, and we can argue using this thickened-up object just as we did when H was a genuine line graph. We make the union of these strips maximal, and again look at how other vertices can attach to it; and now we can prove that every vertex is either minor or major.

But the really nice thing is that for any connected set F of minor vertices, the set of vertices of H with a neighbour in F looks like the set of neighbours of a single minor vertex (that is, it is contained in a potato or in a strip between potatoes); and for every anticonnected set X of major vertices, the set of vertices of H with a non-neighbour in X looks like the set of non-neighbours of a single major vertex (for each potato, all these non-neighbours in the potato belong to one strip incident with the potato). That's not quite true, but we can analyze all the ways in which it can fail, and handle them all. (The problems only arise when J_0 is very small, and mostly when J is also very small.) And given that, now we can decompose G. If there is a major vertex, take a maximal anticonnected set of them; then it and its common neighbours break H up into its constituent strips, so it is easy to find a skew partition. Now assume there are no major vertices. If there is a component of minor vertices attaching onto a potato, again it is easy to find a skew partition; and if not, and either some component of minor vertices is attaching on a strip, or some strip has more than a couple of vertices, then G admits a 2-join. If none of these happen, then G itself equals H and is a line graph of a bipartite graph.

So we win if we can find a large, decently-connected line graph of a bipartite graph contained in G. We said earlier that we hoped to use the paradigm proof technique that was used for most of the other graph decomposition theorems, which involved finding a small subgraph H of G that did *not* lie in any basic classes, and proving it induced a decomposition of G. Note that we haven't done that, and we are doing something like the opposite. We are finding a big, indeed maximal, subgraph H of G that *does* lie in one of the basic classes; and proving that either G itself lies in the same basic class, or H induces a decomposition of G. This method actually seems more powerful, and I don't know why it took us so long to think of applying it here.

The division of the remaining vertices of G into major and minor, and the phenomenon that connected sets of minor vertices attach like one minor vertex, and anticonnected sets of major vertices attach like one major vertex, might seem amazingly lucky; but a similar thing happened several times in the course of the rest of the proof of 3.1, so often that we came to expect it. I don't know if there is a deep reason behind it.

What happens when J is small? Very small line graphs of bipartite graphs can be complements of line graphs. (We care about two such graphs: the line graph of $K_{3,3}$, and the line graph of $K_{3,3} \setminus e$, the graph obtained by deleting one edge from $K_{3,3}$. These are isomorphic to their own complements).

Suppose G is a big Berge graph, containing a subgraph H which is the line graph of $K_{3,3}$, and this is maximal in the sense we considered. We hope to deduce that G has a decomposition. But it is possible that G is the complement of a line graph of a bipartite graph, with no useful decomposition, so the proof sketched above cannot work when J is this small. The best we can hope is to prove that either G or its complement is a line graph of a bipartite graph, or G has a decomposition, and in fact that works when $J = K_{3,3}$ (we need to use the fact that H is a maximal line graph both in G and in the complement). For the line graph of $K_{3,3} \setminus e$ things are even worse, because this is basic in *three* ways; it is also a double split graph. The proof is therefore going to become even more complicated, and it is not worth sketching here.

But in the end, we were able to prove that 3.1 was true for any Berge graph G containing the line graph of a bipartite subdivision of K_4. Fantastic timing too; I was organizing a conference on perfect graphs at Princeton, and we held the conference and announced the theorem, just a couple of days after we found the result. Fortunately the proof still more or less held together when we had time to examine it later. (The timing of the end of the meeting was not so great, however — the final day was September 11, the day of the World Trade Center attack, and not a good day for travelling.)

Now at last we were getting somewhere; our decomposition approach was working, and if there was any justice then there was a proof of the SPGC to be discovered along these lines. Neil had to go home, but Maria and I just needed to prove more of the same.

6 The rest of the proof

Now we can concentrate on Berge graphs that do not contain the line graph of a bipartite subdivision of K_4 (and if it helps, we can also assume the same for the complement of G). What should be the next step? A *prism* means an induced subgraph formed by two disjoint triangles, joined by three vertex-disjoint paths, and prisms suddenly look much more inviting. Suppose that our Berge graph G contains a prism. A prism is a line graph of a bipartite graph, and since we can assume that G contains no richer line graphs of bipartite graphs, perhaps we can now understand how the remainder of G can attach onto the prism. There are two cases, depending whether the three paths of the prism have even length or odd (they clearly all have even length or all have odd length). For the even case, things work as we would hope; the arguments that we used for analyzing maximal line graphs (major vertices and minor ones, strips and potatoes) can be used in this context as well, and it all works nicely. For the odd prism, it doesn't work nicely at all, it becomes very difficult — we were stuck here for three months, until December 2001. Indeed, we were persuaded that the odd prism was hopeless, and we had better look elsewhere, and we tried a whole lot of other things, without success. Until finally we worked on another graph (three disjoint K_4's, joined by three sets of two paths; this is much bigger than a prism, of course, but it's easier, because not so many things can go wrong at once) and had a good idea that still had some merit when transferred back to a prism. I don't want to talk about the proof for the odd prism here any more, but we proved that 3.1 was true for graphs that contained any prism, except for the prism in which all three paths have length 1. (Call all prisms except this one *long* prisms.)

So now we only need think about Berge graphs that do not contain long prisms or line graphs of bipartite subdivisions of K_4, and nor do their complements. Next there was a cascade of easy, pretty theorems, that we proved in a great flurry of excitement in a matter of a few days; for instance

- A *double diamond* means a graph formed by two copies of $K_4 \setminus e$, joined by a perfect matching joining corresponding vertices. We proved 3.1 for graphs containing a double diamond. (Again, minor vertices and major ones, and it's very pretty this time, because taking complements exchanges the two.)

- We proved that if G is Berge, admitting a skew partition, then either it contains a long prism, or a line graph of a bipartite subdivision of K_4, or a double diamond, or it admits a balanced skew partition. Consequently we proved Chvátal's conjecture that no minimum counterexample to the SPGC contains a skew partition.

- An *odd wheel* is a wheel (C, X) such that there are vertices $u, v \in V(C)$, not X-complete, joined by a path in C containing an odd number of X-complete edges. We proved that any Berge graph containing an odd wheel satisfies 3.1 (minor vertices and major ones again).

That was the end of December 2001. I wrote in my Princeton summary of research for the past year — my claim for a pay-raise — that we had settled the SPGC, because I really thought we would have it in a few days. But it turned out that I cut that corner a little too closely. We only have to consider Berge graphs that contain no long prisms, no line graphs of bipartite subdivisions of K_4, no double diamonds, and no odd wheels, and nor do their complements; and we thought that surely we were over the worst of the proof now, that Berge graphs must be easier to handle when we can exclude so much. But that was not so, and in fact we were just coming to the most difficult part. We made absolutely no further progress, despite working full-time on it, until May 2002.

What one would like to do next is the "hole with a hat" — a hole C of length at least six, and a vertex with exactly two neighbours in C, adjacent. This type of subgraph comes up all the time now, and one can almost handle it. We battled with it for a long time, but failed. Maybe this can be done, and if so it would improve our proof a great deal.

We switched instead to trying to handle general wheels (C, X). Since we have already handled odd wheels, we can assume that every maximal subpath of C consisting only of X-complete vertices has even length; but what can we do with such a wheel? The trick is to choose a wheel (C, X) with X maximal *over all wheels*. That gives more power than just assuming that X is maximal for the given wheel. We prove then that X and its common neighbours disconnect the graph and therefore make a skew partition (unless other good things happen instead); but the proof is tremendously long and complicated, and I'll skip it here.

Now we can exclude wheels in general, as well as everything else, both in G and in its complement, and it does get easier from here on. There are just four more steps. If there are a hole and an antihole in G both of length at least six, then either some vertex of the antihole would be adjacent to more than half the hole, or some vertex of the hole would be nonadjacent to more than half the antihole, and so either G or its complement contains a wheel, a contradiction. (There is a case missed here – perhaps every vertex of the hole is adjacent to exactly half of the antihole, and vice versa – but then G contains a double diamond.) So we can assume there are no antiholes of length at least six. Second, suppose there is a hole of length at least six and a vertex with three consecutive neighbours in the hole. Then we can produce a skew partition, using basically the same machinery that we used for wheels, with a little coaxing; having done wheels this was not difficult. So we can exclude this too. Third, with these excluded, we did the hole with a hat; and fourth we could exclude this too, and do all the graphs that were left. These last two steps together just took a page, and it is not worth describing them in any more detail. That completes the proof.

7 What's left?

Having worked in Berge graphs for three years now, we have developed intuitions and skills that took a long time to grow, and also a great fondness for the graphs themselves. Unfortunately the main problem is solved, and there is a cold wind blowing, almost as if it's time to go and work in a new area. Isn't there something else we can work on in the perfect world instead?

Of course, one might still hope for a short proof of the SPGC itself, without proving 3.1, though I have no idea how to do that. There was one other really nice question — what about a polynomial-time recognition algorithm? Can one decide in polynomial time whether a graph is Berge? Is the question in NP? These were still open. One would expect that the decomposition theorem 3.1 would immediately lead to such an algorithm, because such structure theorems usually do; one simply tests whether the graph has the structure described by the theorem. But our decomposition theorem would not cooperate, and in particular the skew partition part of it gave us major difficulties. We thought it would last us for another three happy years, but sadly its resistance collapsed after just a couple of months, and Maria and I managed to twist it into an algorithm. It didn't feel natural, though, and eventually, to our great surprise, we noticed that we were not really using the decomposition theorem at all; in all the places that we used it, we could do without it, and the algorithm became simpler. Then Cornuéjols, Liu and Vušković found a big simplification to one of the steps (they had also discovered some of our algorithm independently), and we wrote it all up jointly with them [2]. So that's done — what else?

As far as I can see there are only two other good questions left. One is, what about constructions? We have not yet found what one could reasonably call a construction for Berge graphs; the algorithm above is far from a construction. It's not clear whether this exists at all, and certainly we seem to be a long way from it still. I have a few half-baked ideas, but they are not worth detailing here. The other question is, what about a polynomial time algorithm to find an optimal colouring of a perfect graph? This can be done using the ellipsoid method, but that uses polyhedral methods and real numbers, and there should be a "combinatorial" algorithm; it seems a little paradoxical that one can test in polynomial time if a graph is perfect, and still not be able to colour it. So that's a nice question.

There is one other avenue — what about a better proof of 3.1? Does the proof really have to be so long? I don't think there will be a short proof; the theorem itself is too complicated. Over the years I have authored or co-authored a number of similar decomposition theorems, all with long and complicated proofs (though none were as long as this one), and so far no-one has come up with a significant shortening. So I predict that we are stuck permanently with this proof, or something like it.

On the other hand, possibly some parts of the proof could be improved. The first half (up to excluding prisms) of the proof feels right to me; I think that here we are following the bones of the mathematics, and it can't be made much better. But for the second half, there could be some big improvements, and conceivably this half could become quite short, with the right new idea. It needs to be looked at more.

References

[1] C. Berge, “Färbung von Graphen, deren sämtliche bzw. deren ungerade Kreise starr sind”, *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe* 10 (1961), 114.

[2] M. Chudnovsky, G. Cornuéjols, X. Liu, P. Seymour, and K. Vušković, “Recognizing Berge graphs”, *Combinatorica*, 25 (2005), 143-186.

[3] M. Chudnovsky, N. Robertson, P. Seymour and R. Thomas, “The strong perfect graph theorem”, manuscript, June 2002.

[4] M. Conforti and G. Cornuéjols, “Graphs without odd holes, parachutes or proper wheels: a generalization of Meyniel graphs and of line graphs of bipartite graphs”, *J. Combinatorial Theory, Ser. B*, 87 (2003), 300-330.

[5] M. Conforti, G. Cornuéjols and K. Vušković, “Square-free perfect graphs”, manuscript, February 2001.

[6] G. Pólya, Aufgabe 424, *Arch. Math. Phys.* 20 (1913), 271.

[7] N. Robertson, P. Seymour and R. Thomas, “Hadwiger’s conjecture for K_6-free graphs”, *Combinatorica* 13 (1993), 279-361.

[8] N. Robertson, P. Seymour and R. Thomas, “Permanents, Pfaffian orientations, and even directed circuits”, *Annals of Math.*, 150 (1999), 929-975.

[9] N. Robertson, P. Seymour and R. Thomas, “Tutte’s edge-colouring conjecture”, *J. Combinatorial Theory, Ser. B* 70 (1997), 166 - 183.

[10] F. Roussel and P. Rubio, “About skew partitions in minimal imperfect graphs”, *J. Combinatorial Theory, Ser. B* 83 (2001), 171-190.

[11] P. Seymour, “Decomposition of regular matroids”, *J. Combinatorial Theory, Ser. B* 28 (1980), 305-359.

[12] A.C. Tucker, “Critical perfect graphs and perfect 3-chromatic graphs”, *J. Combinatorial Theory, Ser. B* 23, 143-149.

[13] W.T. Tutte, “On the algebraic theory of graph colorings”, *J. Combinatorial Theory* 1 (1966), 15-50.

[14] K. Wagner, “Über eine Eigenschaft der ebenen Komplexe”, *Math. Ann.* 114 (1937), 570-590.