

Towards a characterization of bipartite switching classes by means of forbidden subgraphs

Jurriaan Hage (jur@cs.uu.nl) and Tero Harju

institute of information and computing sciences,
utrecht university

technical report UU-CS-2006-028

www.cs.uu.nl

Abstract. We investigate which switching classes do not contain a bipartite graph. Our final aim is a characterization by means of a set of critically non-bipartite graphs: they do not have a bipartite switch, but every induced proper subgraph does. In addition to the odd cycles, we list a number of exceptional cases and prove that these are indeed critically non-bipartite. Finally, we give a number of structural results towards proving the fact that we have indeed found them all. The search for critically non-bipartite graphs was done using software written in C and Scheme. We report on our experiences in coping with the combinatorial explosion.

AMS classification: 05C22, 05C75

Keywords: switching classes, bipartite graphs, forbidden subgraphs, combinatorial search

1 Introduction

For a finite undirected graph $G = (V, E)$ and a set $\sigma \subseteq V$, the *switch* of G by σ is defined as the graph $G^\sigma = (V, E')$, which is obtained from G by removing all edges between σ and its complement $\bar{\sigma}$ and adding as edges all nonedges between σ and $\bar{\sigma}$. The switching class $[G]$ determined by G consists of all switches G^σ for subsets $\sigma \subseteq V$.

A switching class is an equivalence class of graphs under switching. The initiators of the theory of switching classes of graphs were Van Lint and Seidel [12]. They used the model in their investigation of elliptic geometry. The book on 2-structures by Ehrenfeucht, Harju and Rozenberg has a number of chapters on switching classes of graphs and their generalizations [3]. A book completely devoted to the subject of switching is the first author's thesis [4]. Part of the motivation for the general model treated in the latter two books is that they constitute a way in which to model the semantics of a certain type of networks of processors. Switching classes have also been found useful in the fields of psychosociology, and in the investigation of Ising models in statistical physics. For more information on applications of switching classes consult the preface of the dynamic bibliography of signed and gain graphs and allied areas, compiled by Zaslavsky [13].

A graph $G = (V, E)$ is *bipartite* if V can be partitioned into two sets A and $V - A$, such that all edges in E are between A and $V - A$. An equivalent characterization is to say that G is bipartite if it contains no cycles of odd length. Deciding whether a graph is bipartite is easily done by visiting each node at most once.

The bipartiteness problem for switching classes is to determine, given the generator G of the switching class, whether this class contains a bipartite graph:

$\mathcal{B}(G)$ iff there exists $H \in [G]$ such that H is bipartite.

Easy examples are when G is complete bipartite (this includes the discrete graph), since all graphs in $[G]$ are (complete) bipartite. If G is K_3 , then the class contains a non-bipartite graph, K_3 , but all other graphs in this switching class are bipartite. A last example here is when G is K_5 : then $[G]$ contains no bipartite graphs (since this class contains only K_5 , $K_4 \cup x$ and $K_3 \cup K_2$). Hage, Harju and Welzl give an algorithm to decide \mathcal{B} in time quadratic in the number of vertices in the graph [7]. In this paper, we try to characterize the switching classes that do not contain a bipartite graph by means of a set of forbidden subgraphs. In other words, we intend to find a minimal set of graphs \mathcal{F} such that

$\mathcal{B}(G)$ iff for all $F \in \mathcal{F}$, F is not an induced subgraph of any $H \in [G]$.

Obviously, if $F_1 \in [F_2]$, then at most of one these need be in \mathcal{F} . Furthermore, we can restrict \mathcal{F} to so-called critically non-bipartite graphs: they do not have a bipartite switch, but every induced subgraph does.

For the case that the predicate \mathcal{B} is 'acyclic' instead of 'bipartite', Hage and Harju showed that besides the infinite family of simple cycles C_n with $n \geq 7$, there are only finitely many of such critically cyclic graphs [6] (a computer program discovered 905 such graphs, divided into 24 switching classes). They proved that among the graphs of order at least 10 there are no more exceptional cases: only simple cycles. Since the computer program was not proven to be correct, it is possible that critically cyclic graphs of less than 10 vertices were missed.

By means of a computer program, we have managed to find six exceptional critically non-bipartite switching classes, besides the ones generated by the simple cycles C_n of odd order $n \geq 7$. We suspect that these are all of them, but we have not managed to prove that fact. What we do show in this paper is that the graphs we found are indeed critically non-bipartite graphs. In addition, we give a number of results about switching classes that do contain bipartite graphs in Section 3, and structural properties of the critically non-bipartite graphs. The first of these properties help show that a graph has no bipartite switches, the latter restricts us in what kind of critically non-bipartite we may expect to find in the future. In the final section, we explain how we tackled the computational problem of finding the critically non-bipartite graph, especially how to deal with the combinatorial explosion in such a search. There is quite a bit of programming effort and computer time involved here, and the fact that we have performed this search (completely up to 12 and partly for 13 vertices) and report on the outcomes, is one of the contributions of this paper.

2 Preliminaries

For a (finite) set V , let $|V|$ be the *cardinality* of V . We shall often identify a subset $A \subseteq V$ with its characteristic function $A : V \rightarrow \mathbf{Z}_2$, where $\mathbf{Z}_2 = \{0, 1\}$ is the cyclic group of order two. We use the convention that for $x \in V$, $A(x) = 1$ if and only if $x \in A$. The restriction of a function $f : V \rightarrow W$ to a subset $A \subseteq V$ is denoted by $f|_A$. We denote set difference by $A - B$. It contains the elements in A which are not in B . If B is a singleton $\{b\}$, then we may write $A - b$ for brevity.

The set $E(V) = \{\{x, y\} \mid x, y \in V, x \neq y\}$ denotes the set of all unordered pairs of distinct elements of V . We write xy or yx for the

undirected pair $\{x, y\}$. The graphs of this paper will be finite, undirected and simple, i.e., they contain no loops or multiple edges. We use $E(G)$ and $V(G)$ to denote the set of edges E and the set of vertices V , respectively, and $|V|$ and $|E|$ are called the *order*, respectively, *size* of G . Analogously to sets, a graph $G = (V, E)$ will be identified with the characteristic function $G : E(V) \rightarrow \mathbf{Z}_2$ of its set of edges so that $G(xy) = 1$ for $xy \in E$, and $G(xy) = 0$ for $xy \notin E$. Later we shall use both notations, $G = (V, E)$ and $G : E(V) \rightarrow \mathbf{Z}_2$, for graphs.

Let $G = (V, E)$ be a graph. A vertex $x \in V$ is *adjacent* to $y \in V$ if $xy \in E$. The *degree* of x in G is the number of vertices it is adjacent to. The *neighbours* of u in G , denoted $N_G(u)$, or $N(u)$ if G is clear from the context, is the set of vertices adjacent to u in G . A vertex which is not adjacent to any other vertex in a graph is called *isolated*.

For a graph $G = (V, E)$ and $X \subseteq V$, let $G|_X$ denote the *subgraph* of G *induced* by X . Hence, $G|_X : E(X) \rightarrow \mathbf{Z}_2$. For two graphs G and H on V we define $G + H$ to be the graph such that $(G + H)(xy) = G(xy) + H(xy)$ for all $xy \in E(V)$. The disjoint union of two graphs G and H , is denoted $G \cup H$.

A graph $G = (V, E)$ is *bipartite* if V can be partitioned into two sets A and $V - A$, such that all edges in E are between A and $V - A$. An equivalent characterization is to say that G is bipartite if and only if it contains no cycles of odd length.

Some graphs we will encounter in the sequel are K_V , the clique on the set of vertices V , and \bar{K}_V , the complement of K_V which is the discrete graph on V ; the complete bipartite graph on A and $V - A$ is denoted by $K_{A, V-A}$. If the choice of vertices is unimportant we can write K_n , \bar{K}_n and $K_{m, n-m}$ for $n = |V|$ and $m = |A|$.

We continue now with definitions for the switching of graphs.

A *selector* for G is a subset $\sigma \subseteq V(G)$, or alternatively a function $\sigma : V(G) \rightarrow \mathbf{Z}_2$. We reserve lower case σ for selectors (subsets). A *switch* of a graph G by σ is the graph G^σ such that for all $xy \in E(V)$,

$$G^\sigma(xy) = \sigma(x) + G(xy) + \sigma(y) .$$

Clearly, this definition of switching is equivalent to the one given at the beginning of the introduction. The set $[G] = \{G^\sigma \mid \sigma \subseteq V\}$ is called the *switching class* of G . The set of graphs $[G]$ is called a switching class, because switching is a reflexive, symmetric and transitive operation: composition of two selectors amounts to taking the symmetric difference. This result can be used to prove the following.

Lemma 1.

It holds that $G^{\sigma_1} = G^{\sigma_2}$ if and only if $\sigma_1 = \sigma_2$ or $\sigma_1 = V(G) - \sigma_2$.

A selector σ is *constant* on $X \subseteq V$ if $X \subseteq \sigma$, or $X \cap \sigma = \emptyset$. The name arises from the fact that, in these cases, $G|_X = G^\sigma|_X$.

We now give a few (standard) results from the literature that will be used in this paper, see e.g. Hage [4].

Lemma 2.

The switching class $[\overline{K}_V]$ equals the set of all complete bipartite graphs on V .

Proof. Given any complete bipartite graph $K_{\{\sigma, V-\sigma\}}$, we can obtain it from \overline{K}_V by switching with respect to σ . Also, every switch of \overline{K}_V can divide its partition at most in two, so it must be a bipartite graph. It is obviously complete bipartite, because \overline{K}_V has no edges.

From the observation that computing G^σ amounts to computing $G + K_{\sigma, V(G)-\sigma}$ we obtain the following result.

Lemma 3.

It holds that $G \in [H]$ if and only if $G + H \in [\overline{K}_V]$.

Lemma 4.

Let $G = (V, E)$ be a graph, $u \in V$ and $A \subseteq V - \{u\}$. There exists a unique graph $H \in [G]$ such that the neighbours of u in H are the vertices in A .

Proof. The vertex u has no adjacent vertices in $G_u = G^{N(u)}$, where $N(u)$ is the set of neighbours of u in G . Switching G_u with respect to A connects u to every vertex in A (and no others) yielding H .

To show that H is unique: let H' be such that $N(u) = A$ in H' . Since H and H' are in the same switching class $H + H'$ is a complete bipartite graph (Lemma 3), say $G_{B, V-B}$. Because u has the same neighbours in both, u is isolated in $G_{B, V-B}$. Hence, $G_{B, V-B}$ is a discrete graph and, consequently, $H = H'$.

As a corollary we find that for every vertex $x \in V(G)$, there is a unique graph in $[G]$ where x is isolated.

3 Switching classes that do contain a bipartite graph

The first important thing to realize is that if a switching class contains a bipartite graph, then all graphs in the switching class are of the form

shown in Figure 1(a): two arbitrarily connected complete bipartite graphs, with bipartitions $(A, V - A)$ and $(\sigma, V - \sigma)$ for some $\sigma \subseteq V$. In this picture $A, V - A$ is the supposed bipartition and σ is any given selector, yielding at most a four-partition, $A \cap \sigma, (V - A) \cap \sigma, A \cap (V - \sigma)$ and $V - (A \cup \sigma)$. In other words the graph is 4-colourable. In the following we shall call an element of this partition a block.

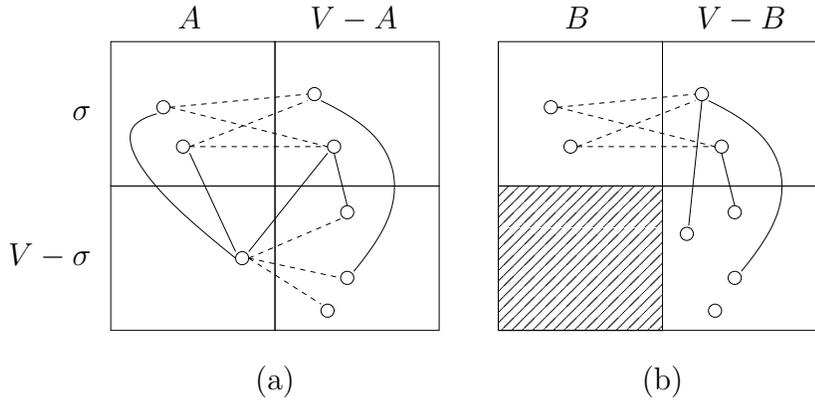


Fig. 1. Partitioned into (a) four and (b) three blocks

This property is an instance of a more general one, which says that if G has chromatic number χ , then every graph in $[G]$ has a chromatic number between $\chi/2$ and 2χ [4].

If three of the blocks are empty, then the graph is discrete, and its switching class contains exactly all complete bipartite graphs (of the same order). If exactly two are empty, then the graph itself is bipartite. Note that by Lemma 4 we can turn a four-partition into a three-partition by making the set of neighbours of one of the vertices empty, as in Figure 1(b).

The following result was already proved in Hage, Harju, Welzl [7]. We give a proof in a different style, here.

Lemma 5.

If G contains (induced) C_{2k+1} for $k \geq 3$, then every $H \in [G]$ contains an induced odd cycle.

Proof. We prove the result by showing that C_{2k+1} with $k \geq 3$ cannot be embedded in Figure 1. As a result, any graph that contains C_{2k+1}

induced, cannot be embedded in Figure 1. This implies that it does not have a bipartite switch and thus that all its switches have an induced odd cycle.

First, observe that if one of the four blocks has more than two vertices, then the juxtaposed block (the opposite part of the corresponding complete bipartite graph) must be empty. This is a consequence of the fact that no vertex in C_{2k+1} can have degree more than two.

Consider the case for $k \geq 4$. Without loss of generality, σ contains at least $k + 1 \geq 5$ vertices. Since two of these must be adjacent, both $A \cap \sigma$ and $(V - A) \cap \sigma$ are non-empty, and one of them has at least three vertices.

We conclude with the case for $k = 3$: C_7 is not bipartite so at least three blocks are nonempty. If one block is empty, then at least one of the blocks has at least three vertices. Because these form an independent set, the remaining vertices of C_7 are K_2 plus two isolated vertices, which is not a complete bipartite graph. If every block is non-empty, then no block can have more than two vertices. Hence, three blocks have two vertices, and one has one vertex. But since two of these doubletons are completely connected, we have an induced cycle C_4 .

Since omitting any vertex from C_{2k+1} gives a bipartite graph, we have the following.

Corollary 1. *The graphs C_{2k+1} are all critically non-bipartite.*

Lemma 2 showed that there is a switching class consisting of all complete bipartite graphs. The following result shows that there is only one switching class containing only bipartite graphs.

Theorem 1.

The switching class $[G]$ contains only bipartite graphs if and only if $[G]$ consists of the complete bipartite graphs on the domain of G .

Proof. The if-part is clear. Now, let $H \in [G]$ be a bipartite graph on A and $V - A$; if both sets contain at most one node, then H is complete bipartite.

Since H is not discrete there is an edge uv in H . If H is not connected we can switch a node x in another component and get a triangle $\{u, v, x\}$. If H is connected, then let $u \in A$ and $v \in V - A$ be such that they are not adjacent in H (they exist because H is not complete bipartite). But since the graph is connected, u has a neighbour, say x . Clearly, $x \in V - A$ and thus x and v are not adjacent. Again $\{u, v, x\}$ can be switched into a triangle.

The above result can also be viewed as an example of characterizing the switching classes that contain only bipartite graphs by means of forbidden subgraphs, the forbidden graphs in this case being K_3 and $K_1 \cup K_2$.

Corollary 2. *If a graph in a switching class avoids both K_3 and $K_1 \cup K_2$, then the switching class contains only bipartite graphs.*

A similar result was obtained by Hertz [8] for switching classes that contain only perfect graphs.

4 The known exceptional cases

In addition to the cycles C_{2k+1} , we have found a small number of critically non-bipartite switching classes. Representatives of these are given in Figure 2 (we shall refer to these graphs by (n) where n is the order of the graph). Thus we have graphs (5) for the K_5 , (6) for $C_5 \cup x$, (7) for $G \cup x$, where G is the “antenna graph”, and (8) (which switches to $\text{co-}C_7 \cup x$), (9) and (11) for the others. The only graph among in the set which is not part of the set of critically cyclic graphs of [6] is (11).

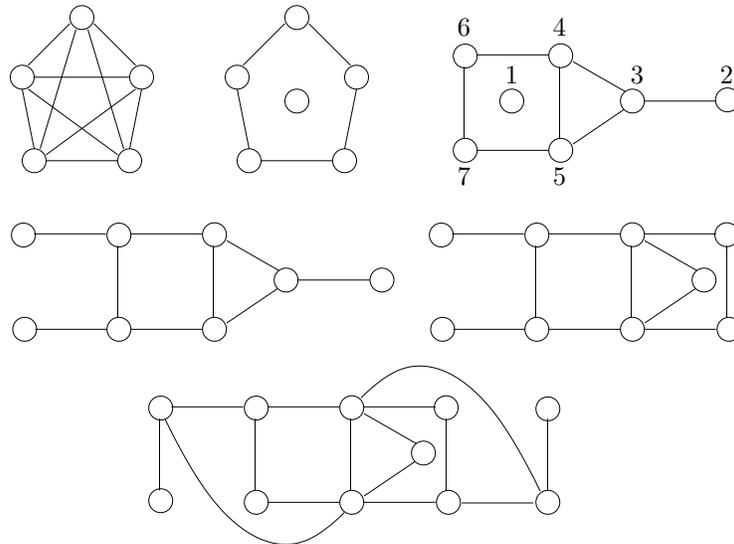


Fig. 2. The exceptional critically non-bipartite graphs

Lemma 6.

Each of the graphs in Figure 2 is critically non-bipartite.

Proof. We have to prove for each graph that all their switches are non-bipartite, and if we remove a vertex, then the resulting graph can be switched to a bipartite graph. The latter follows if for every vertex, we can find a switch so that all odd cycles go through that single vertex.

The graph K_5 switches only to $K_4 \cup x$ and $K_3 \cup K_2$, and omitting from the latter one of the vertices of the K_3 gives the bipartite $K_2 \cup K_2$. We consider next (6), the $C_5 \cup x$. Omitting any vertex on the C_5 which is part of (6), removes the odd cycle. Switching a single vertex on the C_5 yields the *net* (a K_3 with an edge attached to each vertex of the triangle). The vertex x is one of the vertices of the K_3 . The other graphs can be handled similarly, and these cases will be omitted here.

Proving that a graph has no bipartite switches can be done in two ways: compute all switches and verify that none of them is bipartite, or show that the graph cannot be embedded in Figure 1(a) (or Figure 1(b) if it has an isolated vertex). We illustrate the former using (7) as an example. Since (7) has an isolated node, we use Figure 1(b). Each of the vertices of the triangle goes into a different part of the partition, so one of them ends up in the same class as the isolated vertex. There are two cases: 1 goes with 3 or 1 goes with 4 (5 is exactly the same). In the first case, 2 has to go with either 4 or 5, but then it must be connected to the other, which it is not. In the second case, we cannot put 6 anywhere: it cannot go with 5 since it is not connected to 3, it cannot go with 3 since it is not connected to 5, and it cannot go with 4, since it is connected to 4.

5 Properties of critically non-bipartite graphs

Some simple but essential properties of critically non-bipartite graphs are the following.

Lemma 7.

Let G be critically non-bipartite and let $v \in V(G)$.

- i. If G has an isolated vertex, then the largest induced cycle of G is C_3 .*
- ii. There is a selector σ on $V - \{v\}$ such that $(G - v)^\sigma$ is bipartite.*
- iii. $[G]$ contains a graph H in which all odd cycles go through v , and there is at least one such cycle. Furthermore, the same holds for $H^{\{v\}}$.*

Proof. For the first case, it suffices to observe that $C_5 \cup x$, and C_{2k+1} for $k \geq 3$ are already critically non-bipartite. The second case follows from the definition of critically non-bipartite. The final case follows from the second case: if $(G - v)^\sigma$ is bipartite, then if we set $\sigma(v) = 0$, then $(G^\sigma) - v$ is bipartite, so all odd cycles in G^σ go through v . The same holds if we take $\sigma(v) = 1$.

Because every graph can be switched to a graph with an isolated vertex, the following bound is as tight as possible.

Lemma 8.

Every critically non-bipartite graph has at most one isolated vertex.

Proof. Let G be a critically non-bipartite graph with isolated vertices $I = \{v_1, \dots, v_m\}$ and assume $m \geq 2$. First of all, because G has isolated vertices and G is not bipartite, there is at least one induced C_3 by Lemma 7(i), say on $U = \{u_1, u_2, u_3\}$.

Let σ be the selector of Lemma 7(ii) such that $H = (G - v_1)^\sigma$ is bipartite. Then σ is not constant on U . Assume without loss of generality that $\sigma(u_1) = 1$ and $\sigma(u_2) = \sigma(u_3) = 0$. Then $\sigma(v_2) = \dots = \sigma(v_m) = 0$, because if $\sigma(v_i) = 1$ for $2 \leq i \leq m$, then $\{v_i, u_2, u_3\}$ is a triangle in H .

We extend σ to $V(G)$ by $\sigma(v_1) = 0$, and we prove that G^σ is bipartite or H is not bipartite. By Lemma 7(iii) all odd cycles go through v_1 in G^σ , but v_2 is on no odd cycle at all. This leads to a contradiction because v_2 and v_1 are both not selected by σ , which implies that they have exactly the same neighbours in G^σ .

This leads to the following corollary which restricts us somewhat in our search for new critically non-bipartite graphs.

Corollary 3. *Every critically non-bipartite graph G of order at least six has at most two components. If G has two components, then one of these is a single vertex and if G contains an induced cycle, then this cycle is a C_3 .*

6 Conjectures

Although we have investigated the matter at some length, we have not been able to prove the following result which we leave as conjecture:

Conjecture 1.

Besides the odd simple cycles C_{2k+1} for $k \geq 6$, there are no critically non-bipartite graphs of order at least 12.

We leave the reader with some observations that might lead to such a proof. The main problem we have compared to the situation of [6] is that do not have the equivalent of Lemma 5.7 of that paper. It basically gives a normal form for critically cyclic graphs which implies that we have a subgraph of a limited number of types. The brunt of the work was to consider these one by one, but it was this lemma that reduced the work to a finite amount of cases.

What we have observed is that each of the special critically non-bipartite graphs we found has a switch in which there is a single C_3 . For K_5 this is $K_3 \cup K_2$, for $C_5 \cup x$ this is C_3 with an edge attached to each of its vertices (the net in the naming of ISCGI), and the others are given in Figure 2. Although these are not all connected, most of them are. In the case of (7), it does have a connected switch with a single induced cycle, but it is a C_5 (of which two adjacent vertices have a edge attached to them) .

Some conjectures that might be useful towards proving the main conjecture:

- i. Every critically non-bipartite graph has a switch that has a single odd cycle which is a C_3 .
- ii. Such a graph is unique (up to isomorphism).
- iii. Such a graph is planar.

The program that we used in our search for critically non-bipartite graphs has been applied fully up to 12 vertices. About one third of the switching classes on 13 vertices have been considered. Among these we found quite a few critically non-bipartite graphs, all switchable to the cycle C_{13} .

7 The software

To determine the critically non-bipartite graphs, we used a program written in C++. As the number of vertices n increases, the number of switching classes increases as $n^{(n-1)(n-2)/2}$ in the worst case (depending on how well we can avoid looking at isomorphic switching classes). Given a number n the program will generate a list of non-bipartite graphs of order n , from which graphs can be omitted which have certain induced subgraphs (i.e., a critically non-bipartite graph on fewer vertices). Thus we obtain the critically non-bipartite graphs on n vertices. There is a separate program that can remove from such a list all isomorphic graphs, and if need be, all graphs that switch to isomorphic graphs. This is an important tool,

because the brute force algorithm is not guaranteed to generate only non-isomorphic graphs.

This approach is not particularly fast for small numbers of vertices, because the number of isomorphic graphs is relatively high. We used here the files from Spence [11] which list representatives for the switching classes up to isomorphism and up to complementation for up to 10 vertices. This means that up to that number of vertices, we shall (almost) never generate duplicates. The only exceptional cases are those in which the complement of a graph G has a switch isomorphic to G .

In general, it is impossible to extend beyond ten vertices without doing a lot of duplicate work. The reason is that starting from 11 vertices, the sizes of the file in the line of Spence's are simply too large. However, if we are investigating a fixed predicate, such as bipartite, and we already have quite a few graphs that we can forbid, then it may be worthwhile to compute a list of switching classes on a given number of vertices that already exclude the known critical non-bipartite graphs. For instance, for the file on 10 vertices and known critically non-bipartite graphs on at most 10 vertices, this reduces the number of graphs in the file by a factor of about sixteen.

The representation of graphs is as simple as possible, by means of a fixed size adjacency matrix. This makes updating and accessing very efficient, and given that we only investigate relatively small graphs, size is not a problem. We chose to represent each edge by an `int`. Since we consider only undirected graphs, a boolean would be more suitable, but it turned out that since many operation in `C++` work on integers, an inordinate amount of computation time goes into coercing booleans to integers (in some cases as much as seventy percent of execution time). Additional speed-ups were obtained using a fast graph copying operation, that works directly on memory using `memcpy`.

The trustworthiness of our results is enhanced by the existence of a program written in `Scheme` with similar, but limited functionality. Computation in `Scheme` is too slow to perform a brute force search, but it can verify that the graphs found by the `C++` program are indeed critically non-bipartite. The added trustworthiness is a consequence of the fact that the two implementations differ markedly in their approach to the problem.

7.1 Improving efficiency

Profiling was used to determine bottlenecks in the program. Thus we discovered the loss of efficiency due to the encoding undirected edges using booleans, when these are used with for instance bitwise or arithmetic

operators. Another lesson learned from profiling is that it sometimes pays not to be too declarative: a surprising fact, was that at some point the test for bipartiteness was faster than the test whether a graph was a tree or a forest. The reason, as it turned out, was that the latter were implemented by counting the number of edges and comparing that with the number of components, while the former was implemented by means of a low-level, depth-first search.

A less trivial exercise was to speed up the process of switching. Performing a switch takes time quadratic in the number of vertices in the switch. In the original set-up, we took a graph G as a basis, and computed the graphs in the switching class, by performing every possible switch (up to complementation) to G . In the average case, about a quarter of the vertices take part in a switch. But there is a better way: starting from any graph G , there is a sequence of singleton switches that takes one to each graph in the switching class in turn, without visiting a graph twice (until you return to G at the end). There are strong connections here with Gray codes, the Game of Hanoi, and the number of 2-divisors of an integer. See [5] for more details.

Since the graphs K_5 and $C_5 \cup x$ occurred very often and early on in the lists of forbidden graphs, it paid off to have special functions to determine whether they occurred induced in the switch of a graph. These implementations exploit the symmetry of these graphs, avoiding unnecessary injections, and we do not need to explicitly construct the switches of K_5 or G . Here we use Seidel's elementary result that a switching class is determined uniquely by the parities of the numbers of edges in triangles that contain one arbitrary but fixed vertex of, in this case, K_5 . When we extend graphs vertex by vertex, we can use an incremental version that only considers injections which involve the last added vertex. This worked well only for K_5 and $C_5 \cup x$, not for the other graphs.

8 Closing remarks

We have used a computer program to search for critically non-bipartite switching classes. Besides those generated by the simple odd cycles C_{2k+1} , $k \geq 3$, we found six exceptional ones. We have some preliminary results toward proving that we have in fact found them all, but a proof still eludes us. We have applied the same method to other types of graphs, such as chordal graphs, but in those cases, the number of critical graphs soon explodes.

References

1. P. J. Cameron. Cohomological aspects of two-graphs. *Math. Z.*, 157:101 – 119, 1977.
2. D. G. Corneil and R. A. Mathon, editors. *Geometry and Combinatorics: Selected Works of J. J. Seidel*. Academic Press, Boston, 1991.
3. A. Ehrenfeucht, T. Harju, and G. Rozenberg. *The Theory of 2-Structures*. World Scientific, Singapore, 1999.
4. J. Hage. *Structural Aspects Of Switching Classes*. PhD thesis, Leiden Institute of Advanced Computer Science, 2001. <http://www.cs.uu.nl/people/jur/2s.html>.
5. J. Hage. Enumerating submultisets of multisets. *Inf. Proc. Letters*, 85(4):221 – 226, 2003.
6. J. Hage and T. Harju. A characterization of acyclic switching classes using forbidden subgraphs. *SIAM Journal on Discrete Mathematics*, 18(1):159 – 176, 2004.
7. J. Hage, T. Harju, and E. Welzl. Euler graphs, triangle-free graphs and bipartite graphs in switching classes. *Fundamenta Informaticae*, 58(1):23–37, November 2003.
8. A. Hertz. On perfect switching classes. *Discrete Applied Math.*, 89:263–267, 1998.
9. J. J. Seidel. A survey of two-graphs. In *Intern. Coll. Teorie Combinatorie (Roma, 1973)*, volume I, pages 481–511, Rome, 1976. Accad. Naz. Lincei. Reprinted in [2].
10. J. J. Seidel and D. E. Taylor. Two-graphs, a second survey. In L. Lovasz and V.T. Sós, editors, *Algebraic Methods in Graph Theory (Proc. Internat. Colloq., Szeged, 1978)*, volume II, pages 689–711, Amsterdam, 1981. North-Holland. Reprinted in [2].
11. E. Spence. Tables of two-graphs. <http://gauss.maths.gla.ac.uk/~ted/>.
12. J. H. van Lint and J. J. Seidel. Equilateral points in elliptic geometry. *Proc. Kon. Nederl. Acad. Wetensch., Ser. A*, 69:335 – 348, 1966. Reprinted in [2].
13. T. Zaslavsky. A mathematical bibliography of signed and gain graphs and allied areas. *Electronic J. Combin.*, 1999. Dynamic Survey No. DS8.