



Positive-instance driven dynamic programming for treewidth

Hisao Tamaki¹ 

Published online: 16 October 2018

© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

Consider a dynamic programming scheme for a decision problem in which all subproblems involved are also decision problems. An implementation of such a scheme is *positive-instance driven* (PID), if it generates positive subproblem instances, but not negative ones, building each on smaller positive instances. We take the dynamic programming scheme due to Bouchitté and Todinca for treewidth computation, which is based on minimal separators and potential maximal cliques, and design a variant (for the decision version of the problem) with a natural PID implementation. The resulting algorithm performs extremely well: it solves a number of standard benchmark instances for which the optimal solutions have not previously been known. Incorporating a new heuristic algorithm for detecting safe separators, it also solves all of the 100 public instances posed by the exact treewidth track in PACE 2017, a competition on algorithm implementation. We describe the algorithm, prove its correctness, and give a running time bound in terms of the number of positive subproblem instances. We perform an experimental analysis which supports the practical importance of such a bound.

Keywords Treewidth · Tree decomposition · Dynamic programming · Positive-instance driven

1 Introduction

Suppose we design a dynamic programming algorithm for some decision problem, formulating subproblems, which are decision problems as well, and recurrences among

A preliminary and abridged version of this paper was presented at the 25th European Symposium on Algorithms.

✉ Hisao Tamaki
tamaki@cs.meiji.ac.jp

¹ Meiji University, 1-1-1 Higashi-Mata, Tama, Kawasaki 214-8571, Japan

those subproblems. A standard approach is to list all subproblem instances and scan the list from “small” ones to “large”, deciding the answer, positive or negative, to each instance by means of these recurrences. When the number of positive subproblem instances are expected to be much smaller than the total number of subproblem instances, a natural alternative is to generate positive instances only, using recurrences to combine positive instances to generate a “larger” positive instance. We call such a mode of dynamic programming execution *positive-instance driven* or *PID* for short. One goal of this paper is to demonstrate that PID is not simply a low-level implementation strategy but can be a paradigm of algorithm design for some problems.

The decision problem we consider is that of deciding, given graph G and positive integer k , if the treewidth of G is at most k . This graph parameter was introduced by Robertson and Seymour (1986) and has had a tremendous impact on graph theory and on the design of graph algorithms (see, for example, a survey Bodlaender and Koster 2008). The treewidth problem is NP-complete (Arnborg et al. 1987) but fixed-parameter tractable: it has an $f(k)n^{O(1)}$ time algorithm for some fixed function $f(k)$ as implied by the graph minor theorem of Robertson and Seymour (2004), and an explicit $O(f(k)n)$ time algorithm was given by Bodlaender (1996). A classical dynamic programming algorithm due to Arnborg et al. (1987) (ACP algorithm) runs in $n^{k+O(1)}$ time. Bouchitté and Todinca (2002) developed a more refined dynamic programming algorithm (BT algorithm) based on the notions of minimal separators and potential maximal cliques, which lead to algorithms running in $O(1.7549^n)$ time or in $O(n^5 \binom{\lceil (2n+k+8)/3 \rceil}{k+2})$ time (Fomin et al. 2008; Fomin and Villanger 2012). Another important approach to treewidth computation is based on the perfect elimination order (PEO) of minimal chordal completions of the given graph. PEO-based dynamic programming algorithms run in $O^*(2^n)$ time with exponential space and in $O^*(4^n)$ time with polynomial space (Bodlaender et al. 2012), where $O^*(f(n))$ means $O(n^c f(n))$ for some constant c .

There has been a considerable amount of effort on implementing treewidth algorithms to be used in practice and, prior to this work, the most successful implementations for exact treewidth computation are all based on PEO. The authors of Bodlaender et al. (2012) implemented the $O^*(2^n)$ time dynamic programming algorithm and experimented on its performance, showing that it works well for small instances. For larger instances, PEO-based branch-and-bound algorithms are known to work well in practice (Gogate and Dechter 2004). Recent proposals for reducing treewidth computation to SAT solving are also based on PEO (Samer and Veith 2009; Berg and Jarvisalo 2014). From the PID perspective, this situation is somewhat surprising, since it can be shown that each positive subproblem instance in the PEO-based dynamic programming scheme corresponds to a combination of an indefinite number of positive subproblem instances in the ACP algorithm, and hence the number of positive subproblem instances can be exponentially larger than that in the ACP algorithm. Indeed, a PID variant of the ACP algorithm was implemented by the present author and has won the first place in the exact treewidth track of PACE 2016 (Dell et al. 2017), a competition on algorithm implementations, outperforming other submissions based on PEO. Given this success, a natural next step is to design a PID variant of the BT algorithm, which is tackled in this paper.

The resulting algorithm performs extremely well, as reported in Sect. 8. It is tested on DIMACS graph-coloring instances (Johnson 1996), which have been used in the literature on treewidth computation as standard benchmark instances (Gogate and Dechter 2004; Bodlaender and Koster 2006; Musliu 2008; Samer and Veith 2009; Bodlaender et al. 2012; Berg and Järvisalo 2014). Our implementation of the algorithm solves all the instances that have been previously solved (that is, with matching upper and lower bounds known) within 10 seconds per instance on a typical desktop computer and solves 13 out of the 42 previously unsolved instances. For nearly half of the instances which it leaves unsolved, it significantly reduces the gap between the lower and upper bounds. It is interesting to note that this is done by improving the lower bound. Since the number of positive subproblem instances are much smaller when $k < \text{tw}(G)$ than when $k = \text{tw}(G)$, the PID approach is particularly good at establishing strong lower bounds.

We also adopt the notion of safe separators due to Bodlaender and Koster (2006) in our preprocessing and design a new heuristic algorithm for detecting safe separators. With this preprocessing, our implementation also solves all of the 100 public instances posed by PACE 2017 (2017), the successor of PACE 2016. It should be noted that these test instances of PACE 2017 are much harder than those of PACE 2016: the winning implementation of PACE 2016 mentioned above, which solved 199 of the 200 instances therein, solves only 62 of these 100 instances of PACE 2017 in the given time of 30 minutes per instance.

Adapting the BT algorithm to work in PID mode has turned out non-trivial. Each subproblem instance in the BT algorithm for given graph G and positive integer k takes the form of a connected set C of G such that $N_G(C)$, the open neighborhood of C in G , is a minimal separator of G with cardinality at most k . For each such C , we ask if C is *feasible*, in the sense that there is a tree decomposition of the subgraph of G induced by $C \cup N_G(C)$ of width at most k that has a bag containing $N_G(C)$ (see Sect. 2 for the definition of a tree decomposition of a graph). The difficulty of making the BT algorithm PID comes from the fact that the recurrence for deciding if C is feasible may involve an indefinite number of connected sets C' such that $C' \subset C$. Thus, even if the number of positive instances is small, there is a possibility that the running time is exponential in that number. We approach this issue by introducing an auxiliary structure we call O-blocks (see Sect. 3) and formulate a recurrences that are binary: a combination of a feasible connected set and a feasible O-block may yield either a larger feasible connected set or a larger feasible O-block. Due to this binary recurrence, we obtain an upper bound on the running time of our algorithm which is sensitive to the number of subproblem instances (Observation 1 in Sect. 5). To support the significance of such a bound, we perform an experimental analysis which shows the existence of huge gaps between the actual number of combinatorial objects corresponding to subproblems and the known theoretical upper bounds.

The rest of this paper is organized as follows. In Sect. 2, we introduce notation, define basic concepts and review facts in the literature. In Sect. 3, we precisely define the subproblems in our dynamic programming algorithm and formulate recurrences. We describe our algorithm and prove its correctness in Sect. 4 and then analyze its running time in Sect. 5. In Sect. 6, we describe our experimental analysis. In Sect. 7,

we describe some implementation details. Finally, in Sect. 8, we give details of the performance results sketched above.

2 Preliminaries

In this paper, all graphs are simple, that is, without self loops or parallel edges. Let G be a graph. We denote by $V(G)$ the vertex set of G and by $E(G)$ the edge set of G . For each $v \in V(G)$, $N_G(v)$ denotes the set of neighbors of v in G : $N_G(v) = \{u \in V(G) \mid \{u, v\} \in E(G)\}$. For $U \subseteq V(G)$, the *open neighborhood* of U in G , denoted by $N_G(U)$, is the set of vertices adjacent to some vertex in U but not belonging to U itself: $N_G(U) = (\bigcup_{v \in U} N_G(v)) \setminus U$. The *closed neighborhood* of U in G , denoted by $N_G[U]$, is defined by $N_G[U] = U \cup N_G(U)$. We also write $N_G[v]$ for $N_G[\{v\}] = N_G(v) \cup \{v\}$. We denote by $G[U]$ the subgraph of G induced by U : $V(G[U]) = U$ and $E(G[U]) = \{\{u, v\} \in E(G) \mid u, v \in U\}$. In the above notation, as well as in the notation further introduced below, we will often drop the subscript G when the graph is clear from the context.

We say that vertex set $C \subseteq V(G)$ is *connected in G* if, for every $u, v \in C$, there is a path in $G[C]$ between u and v . It is a *connected component* of G if it is connected and is inclusion-wise maximal subject to this condition. A vertex set C in G is a *component associated with $S \subseteq G$* , if C is a connected component of $G[V(G) \setminus S]$. For each $S \subseteq V(G)$, we denote by $\mathcal{C}_G(S)$ the set of all components associated with S . A vertex set $S \subseteq V(G)$ is a *separator* of G if $|\mathcal{C}_G(S)| > \mathcal{C}_G(\emptyset)|$, that is, if its removal increases the number of connected components of G . A component C associated with separator S of G is a *full component* if $N_G(C) = S$. A separator S is a *minimal separator* if there are at least two full components associated with S . This term is justified by this fact: if S is a minimal separator and a, b vertices belonging to two distinct full components associated with S , then for every proper subset S' of S , a and b belong to the same component associated with S' ; S is a minimal set of vertices that separates a from b . A *block* is a pair (S, C) , where S is a separator and C is a component associated with S ; it is a *full block* if C is a full component, that is, $S = N(C)$.

Graph H is *chordal* if every induced cycle of H has length exactly three. H is a *minimal chordal completion* of G if it is chordal, $V(H) = V(G)$, $E(G) \subseteq E(H)$, and $E(H)$ is minimal subject to these conditions. A vertex set $\Omega \subseteq V(G)$ is a *potential maximal clique* of G , if Ω is a clique in some minimal chordal completion of G .

A *tree-decomposition* of G is a pair (T, \mathcal{X}) where T is a tree and \mathcal{X} is a family $\{X_i\}_{i \in V(T)}$ of vertex sets of G such that the following three conditions are satisfied. We call members of $V(T)$ *nodes* of T and each X_i the *bag* at node i .

1. $\bigcup_{i \in V(T)} X_i = V(G)$.
2. For each edge $\{u, v\} \in E(G)$, there is some $i \in V(T)$ such that $u, v \in X_i$.
3. The set of nodes $I_v = \{i \in V(T) \mid v \in X_i\}$ of $V(T)$ induces a connected subtree of T .

The *width* of this tree-decomposition is $\max_{i \in V(T)} |X_i| - 1$. The *treewidth* of G , denoted by $\text{tw}(G)$ is the minimum width of all tree-decompositions of G . We may assume that the bags X_i and X_j are distinct from each other for $i \neq j$ and, under this

assumption, we will often regard a tree-decomposition as a tree T in which each node is a bag.

We call a tree-decomposition T of G *canonical* if each bag of T is a potential maximal clique of G and, for every pair X, Y of adjacent bags in T , $X \cap Y$ is a minimal separator of G . The following fact is well-known. It easily follows, for example, from Proposition 2.4 in Bouchitté and Todinca (2001).

Lemma 1 *Let G be an arbitrary graph. There is a tree-decomposition T of G of width $\text{tw}(G)$ that is canonical.*

The following local characterization of a potential maximal clique is crucial. We say that a vertex set $S \subseteq V(G)$ is *cliquish* in G if, for every pair of distinct vertices u and v in S , either u and v are adjacent to each other or there is some $C \in \mathcal{C}(S)$ such that $u, v \in N(C)$. In other words, S is cliquish if completing $N(C)$ for every $C \in \mathcal{C}(S)$ into a clique makes S a clique.

Lemma 2 (Theorem 3.15 in Bouchitté and Todinca 2001) *A separator S of G is a potential maximal clique of G if and only if (1) S has no full-component associated with it and (2) S is cliquish.*

It is also shown in Bouchitté and Todinca (2001) that if Ω is a potential maximal clique of G and S is a minimal separator contained in Ω , then there is a unique component C_S associated with S that contains $\Omega \setminus S$. We need an explicit way of forming C_S from Ω and S .

Let $K \subseteq V(G)$ be an arbitrary vertex set and S an arbitrary proper subset of K . We say that a component $C \in \mathcal{C}(K)$ is *confined to S* if $N(C) \subseteq S$; otherwise it is *unconfined to S* . Let $\text{unconf}(S, K)$ denote the set of components associated with K that are unconfined to S . Define the *crib* of S with respect to K , denoted by $\text{crib}(S, K)$, to be $(K \setminus S) \cup \bigcup_{C \in \text{unconf}(S, K)} C$: it is the union of $K \setminus S$ and all those components associated with K that have neighborhoods intersecting $K \setminus S$.

The following lemma relies only on the second property of potential maximal cliques, namely that they are cliquish, and will be applied not only to potential maximal cliques but also to separators with full components, which are trivially cliquish.

Lemma 3 *Let $K \subseteq V(G)$ be a cliquish vertex set. Let S be an arbitrary proper subset of K . Then, $\text{crib}(S, K)$ is a full component associated with S .*

Proof Let $C = \text{crib}(S, K)$. We first show that $G[C]$ is connected. Suppose $K \setminus S$ has two distinct vertices u and v . Since K is cliquish, either u and v are adjacent to each other or there is some component $C' \in \mathcal{C}(K)$ such that $u, v \in N(C')$. In the latter case, as C' is unconfined to S , we have $C' \subseteq C$. Therefore, u and v belong to the same connected component of $G[C]$. As this applies to every pair of vertices in $K \setminus S$, $K \setminus S$ is contained in a single connected component of $G[C]$. Moreover, each component $C' \in \mathcal{C}(K)$ contained in C is unconfined to S , by the definition of $\text{crib}(S, K)$, and hence has a neighbor in $K \setminus S$. Therefore, we conclude that $G[C]$ is connected. Each vertex v not in $S \cup C$ belongs to some component in $\mathcal{C}(K)$ that is confined to S and hence does not have a neighbor in C . Therefore, C is a component associated with S .

To see that C is a full component, let $u \in S$ and $v \in K \setminus S$ be arbitrary. Since K is cliquish, either u and v are adjacent to each other or there is some $C' \in \mathcal{C}(K)$ such that $u, v \in N(C')$. As such C' is unconfined to S in the latter case, we conclude that $u \in N(C)$ in either case. Since this holds for arbitrary $u \in S$, we conclude that C is a full component associated with S . \square

Remark 1 As $\text{crib}(S, K)$ contains $K \setminus S$, it is clear that it is the only component associated with S that intersects K . Therefore, the above mentioned assertion on potential maximal cliques is a corollary to this Lemma.

3 Recurrences on oriented minimal separators

In this section, we fix graph G and positive integer k that are given in the problem instance: we are to decide if the treewidth of G is at most k . We assume that G is connected.

For connected set $C \subseteq V(G)$, we denote by $G\langle C \rangle$ the graph obtained from $G[N[C]]$ by completing $N(C)$ into a clique: $V(G\langle C \rangle) = N[C]$ and $E(G\langle C \rangle) = E(G[N[C]]) \cup \{\{u, v\} \mid u, v \in N(C), u \neq v\}$. We say C is *feasible* if $\text{tw}(G\langle C \rangle) \leq k$. Equivalently, C is feasible if $G[N[C]]$ has a tree-decomposition of width k or smaller that has a bag containing $N(C)$.

Let us first review the BT algorithm (Bouchitté and Todinca 2001) adapting it to our decision problem. We first list all minimum separators of cardinality k or smaller and all potential maximal cliques of cardinality $k + 1$ or smaller. Then, for each pair of a potential maximal clique Ω and a minimal separator S such that $S \subset \Omega$, place a link from S to Ω . To understand the difficulty of formulating a PID variant of the algorithm, it is important to note that the pair (Ω, S) to be linked is easy to find from the side of Ω , but not the other way round. Then, we scan the full blocks $(N(C), C)$ of minimal separators in the increasing order of $|C|$ to decide if C is feasible, using the following recurrence: C is feasible if and only if there is some potential maximal clique Ω such that $N(C) \subset \Omega$, $C = \text{crib}(N(C), \Omega)$, and every component $D \in \text{unconf}(N(C), \Omega)$ is feasible. Finally, we have $\text{tw}(G) \leq k$ if and only if there is a potential maximal clique Ω with $|\Omega| \leq k + 1$ such that every component associated with Ω is feasible.

To facilitate the PID construction, we orient minimal separators as follows. We assume a total order $<$ on $V(G)$. For each vertex set $U \subseteq V(G)$, the *minimum element* of U , denoted by $\min(U)$, is the smallest element of U under $<$. For vertex sets U and W , we say U *precedes* W and write $U \prec W$ if $\min(U) < \min(W)$.

We say that a connected set C is *inbound* if there is some full block associated with $N(C)$ that precedes C ; otherwise, it is *outbound*. Observe that if C is inbound then $N(C)$ is a minimal separator, since $N(C)$ has another full component associated with it and, contrapositively, if $N(C)$ is not a minimal separator then C is necessarily outbound. We say a full block $(N(C), C)$ is *inbound (outbound)* if C is inbound (outbound, respectively).

Lemma 4 *Let K be a cliquish vertex set and let A_1, A_2 be two components associated with K . Suppose that A_1 and A_2 are outbound. Then, either $N(A_1) \subseteq N(A_2)$ or $N(A_2) \subseteq N(A_1)$.*

Proof Let K , A_1 , and A_2 be as above and suppose neither of $N(A_1)$ and $N(A_2)$ is a subset of the other. For $i = 1, 2$, let $C_i = \text{crib}(N(A_i), K)$. Since $N(A_2) \setminus N(A_1)$ is non-empty and contained in $K \setminus N(A_1)$, A_2 is contained in C_1 . We have $A_1 \prec C_1$ as A_1 is outbound and hence $A_1 \prec A_2$. A contradiction, since similarly we have $A_2 \prec A_1$. \square

Let K be a cliquish vertex set. Based on the above lemma, we define the *outlet* of K , denoted by $\text{outlet}(K)$, as follows. If no non-full component associated with K is outbound, then we let $\text{outlet}(K) = \emptyset$. Otherwise, $\text{outlet}(K) = N(A)$, where A is a non-full component associated with K that is outbound, chosen so that $N(A)$ is maximal. We define $\text{support}(K) = \text{unconf}(\text{outlet}(K), K)$, the set of components associated with K that are not confined to $\text{outlet}(K)$. By Lemma 4, every member of $\text{support}(K)$ is inbound.

We call a full block $(N(C), C)$ an *I-block* if C is inbound and $|N(C)| \leq k$. We call it an *O-block* if C is outbound and $|N(C)| \leq k$.

We say that an I-block $(N(C), C)$ is *feasible* if C is feasible. We say that an O-block $(N(A), A)$ is feasible if $N(A) = \bigcup_{C \in \mathcal{C}} N(C)$ for some set \mathcal{C} of feasible inbound components. Note that this definition of feasibility of an O-block is somewhat weak in the sense that we do not require every inbound component associated with $N(A)$ to be feasible.

We say that a potential maximal clique Ω is *feasible* if $|\Omega| \leq k + 1$ and every $C \in \text{support}(\Omega)$ is feasible.

In order to formulate mutual recurrences among feasible I-blocks, O-blocks, and potential maximal cliques, we need the following auxiliary notion of *buildable* potential maximal cliques.

Let Ω be a potential maximal clique with $|\Omega| \leq k + 1$. For each $C \in \text{support}(\Omega)$, block $(N(C), C)$ is an I-block, since C is inbound as observed above and we have $|N(C)| \leq k$ by our assumption that $|\Omega| \leq k + 1$. We say that Ω is *buildable* if $|\Omega| \leq k + 1$ and either

1. $\Omega = N[v]$ for some $v \in V(G)$,
2. there is some subset \mathcal{C} of $\text{support}(\Omega)$ such that $\Omega = \bigcup_{D \in \mathcal{C}} N(D)$ and every member of \mathcal{C} is feasible, or
3. $\Omega = N(A) \cup (N(v) \cap A)$ for some feasible O-block $(N(A), A)$ and a vertex $v \in N(A)$.

It will turn out that every feasible potential maximal clique is buildable (Lemma 9).

Lemma 5 *We have $\text{tw}(G) \leq k$ if and only if G has a feasible potential maximal clique Ω with $\text{outlet}(\Omega) = \emptyset$.*

Proof Suppose first that G has a feasible potential maximal clique Ω with $\text{outlet}(\Omega) = \emptyset$. Note that $\text{support}(\Omega) = \mathcal{C}(\Omega)$, as every $C \in \mathcal{C}(\Omega)$ is unconfined to an empty set. For each component $C \in \text{support}(\Omega)$, let T_C be the tree-decomposition of $G \langle C \rangle$ of width k or smaller, which exists since C is feasible by the definition of a feasible potential maximal clique. Let X_C be a bag of T_C such that $N(C) \subseteq X_C$. Combine these tree-decompositions into a tree T by adding bag Ω and letting each X_C in T_C be adjacent to Ω . That T satisfies the first two conditions for tree-decomposition is

trivial. The third condition is also satisfied, since, if a vertex v appears in $N[C]$ for two or more members C in $\text{support}(\Omega)$, then v appears in X_C for each such C and in Ω . Therefore, T is a tree-decomposition of G of width k or smaller and hence $\text{tw}(G) \leq k$.

For the converse, suppose the treewidth of G is k or smaller. Let T be a canonical tree-decomposition of G of width k or smaller: each bag of T is a potential maximal clique and the intersection of each pair of adjacent bags of T is a minimal separator. Orient each edge of T as follows. Let X and Y be adjacent bags in T and let $S = X \cap Y$. Let V_X be the union of bags in the maximal subtree of T containing X but not containing Y ; define V_Y similarly. Let C be the outbound full component associated with the minimal separator S . Then, C is contained in exactly one of V_X and V_Y . If C is contained in V_X then we orient the edge between X and Y from Y to X ; otherwise from X to Y . Since T is a tree, the resulting directed tree has a sink X_0 . Then, each component C associated with X_0 is inbound and hence $\text{outlet}(X_0) = \emptyset$. We show that each such C is moreover feasible. Indeed, the required tree-decomposition of $G\langle C \rangle$ may be obtained from T by taking intersection of every bag with $N[C]$: the resulting tree is a tree-decomposition of $G[N(C)]$ and contains the bag $X_0 \cap N[C] \supseteq N(C)$. The width of the tree-decomposition is not greater than that of T and hence is k or smaller. Therefore, I-block $(N(C), C)$ for each component C associated with X_0 is feasible and hence the potential maximal clique X_0 is feasible. \square

Lemma 6 *Let C be a connected set of G such that $N(C)$ is a minimal separator. Let Ω be a potential maximal clique of $G\langle C \rangle$. Then, Ω is a potential maximal clique of G .*

Proof For each component D associated with $N(C)$, let H_D be a minimal chordal completion of $G\langle D \rangle$. In particular, choose H_C so that Ω is a clique in H_C . Let H be the union of these graphs: $V(H) = V(G)$ and $E(H) = \bigcup_{D \in \mathcal{C}(N(C))} E(H_D)$. It is clear that H is chordal. Let H' be a minimal chordal completion of G contained in H . It is well-known that every minimal separator is a clique in every chordal completion and hence $N(C)$ is a clique in H' . Therefore, the minimality of H_D for each D implies that $H' = H$. As Ω is a clique in H_C , it is a clique in H and hence is a potential maximal clique of G . \square

The following is our oriented version of the recurrence in the BT algorithm described in the beginning of this section.

Lemma 7 *An I-block $(N(C), C)$ is feasible if and only if there is some feasible potential maximal clique Ω with $\text{outlet}(\Omega) = N(C)$ and $\bigcup_{D \in \text{support}(\Omega)} D = C$.*

Proof Suppose first that there is a feasible potential maximal clique Ω as in the lemma. For each component $D \in \text{support}(\Omega)$, let T_D be a tree-decomposition of $G\langle D \rangle$ of width k or smaller and X_D be a bag in T_D containing $N(D)$. Combine these tree-decompositions T_D , $D \in \text{support}(\Omega)$, into a tree T by adding bag Ω and let it be adjacent to X_D for each $D \in \text{support}(\Omega)$. We confirm that T is a tree-decomposition of $G[N[C]]$. Every vertex $v \in N[C]$ appears in some bag of T since C is the union of D for all $D \in \text{support}(\Omega)$ and the bag Ω contains $N(C)$. Every edge of $G[N[C]]$ appears in some bag of T for the same reason. The third condition for T being a tree decomposition is also satisfied, since, if a vertex v appears in $N[D]$ for two or more

members D in $\text{support}(\Omega)$, then v appears in X_D for each such D and in Ω . Therefore, T is a tree decomposition of $G[N[C]]$ of width k or smaller and hence the bag Ω in T contains $N(C)$, T attests the feasibility of the I-block $(N(C), C)$.

For the converse, suppose that I-block $(N(C), C)$ is feasible. Let T be a canonical tree-decomposition of $G(C)$ of width k or smaller. Orient the edges of T as in the proof of Lemma 5: orient the edge from X to Y if and only if Y intersects the outbound full component associated with $X \cap Y$. We need to stress here that the notion of outbound components used in this orientation is with respect to the entire graph G and not with respect to $G(C)$, the graph of which T is a tree-decomposition. As $N(C)$ is a clique in $G(C)$, T contains a bag that contains $N(C)$. In the subtree of T induced by those bags containing $N(C)$, let X_0 be a sink with respect to the above orientation. As T is canonical, X_0 is a potential maximal clique of $G(C)$ and hence of G by Lemma 6. We show below that X_0 is feasible.

Let A be the outbound full component associated with $N(C)$. As $N(C) \subseteq X_0$ and $A \cap N[C] = \emptyset$, A is a component associated with X_0 . We claim that $N(C) = \text{outlet}(X_0)$. Suppose otherwise that there is some outbound component A' associated with X_0 such that $N(C)$ is a proper subset of $N(A')$. Then, as A' is not confined to $N(C)$, $C = \text{crib}(N(C), X_0)$ contains A' . Therefore, there is some bag X adjacent to X_0 in T such that $X \cap A' \neq \emptyset$. Since $N(C)$ is a minimal separator that separates A from A' , X must contain $N(C)$. But, since A' is an outbound component associated with X_0 , the edge between X_0 and X is oriented from X_0 to X . This contradicts the choice of X_0 and we conclude that $N(C) = \text{outlet}(X_0)$.

It remains to verify that each $D \in \text{support}(X_0)$ is feasible. This is true since the tree of bags obtained from T by intersecting each bag with $N[D]$ is a tree-decomposition of $G(D)$ required for the feasibility of D . \square

Lemma 8 *Let K be a cliquish vertex set, \mathcal{C} a non-empty subset of $\text{support}(K)$, and $S = \bigcup_{C \in \mathcal{C}} N(C)$. If S is a proper subset of K then $\text{crib}(S, K)$ is outbound.*

Proof Let K , \mathcal{C} and S be as in the lemma. Since K is cliquish, $\text{crib}(S, K)$ is a full component associated with S that contains $K \setminus S$, by Lemma 3. To show that it is outbound, it suffices to show that no other full component associated with S is outbound. Let A be an arbitrary full component associated with S that is distinct from $\text{crib}(S, K)$. As A does not intersect K , it is a component associated with K . Let C be an arbitrary member of \mathcal{C} . Then, C is confined to S by the definition of S . On the other hand C is not confined to $\text{outlet}(K)$ since $C \in \text{support}(K)$. Therefore, S is not a subset of $\text{outlet}(K)$. A cannot be outbound, since it would imply that $S = N(A) \subseteq \text{outlet}(K)$. Therefore, A is inbound and, since this holds for every full component associated with S other than $\text{crib}(S, K)$, $\text{crib}(S, K)$ is outbound. \square

The following lemma is crucial for our PID result: the algorithm described in the next section generates all buildable potential maximal cliques and we need to guarantee all feasible maximal cliques to be among them.

Lemma 9 *Let Ω be a feasible potential maximal clique. Then, Ω is buildable.*

Proof Let $S = \bigcup_{C \in \text{support}(\Omega)} N(C)$.

Suppose first that $S \cup \text{outlet}(\Omega) \neq \Omega$ and let v be an arbitrary member of $\Omega \setminus (S \cup \text{outlet}(\Omega))$. Since Ω is cliquish and v is not in $N(C)$ for any component C associated with Ω , v is adjacent to every other vertex in Ω . Therefore, $\Omega \subseteq N[v]$. Let C be an arbitrary component associated with Ω . If C is confined to $\text{outlet}(\Omega)$ then $v \notin N(C)$ since $v \notin \text{outlet}(\Omega)$. Otherwise, $C \in \text{support}(\Omega)$ and hence $v \notin N(C)$ as $v \notin S$. Therefore, $N(v) \setminus \Omega$ is empty and hence we have $\Omega = N[v]$. Thus, Ω is buildable, the first case of buildability.

Suppose next that $S \cup \text{outlet}(\Omega) = \Omega$. We have two cases to consider: $S = \Omega$ and $S \neq \Omega$.

Consider the case where $S = \Omega$. Let \mathcal{C}_0 be an arbitrary minimal subset of $\text{support}(\Omega)$ such that $\bigcup_{C \in \mathcal{C}_0} N(C) = \Omega$. Since Ω does not have a full component associated with it, \mathcal{C}_0 has at least two members. Let C_0 be an arbitrary member of \mathcal{C}_0 and let $\mathcal{C}_1 = \mathcal{C}_0 \setminus \{C_0\}$. From the minimality of \mathcal{C}_0 , $S_1 = \bigcup_{C \in \mathcal{C}_1} N(C)$ is a proper subset of Ω . By Lemmas 3 and 8, $A_1 = \text{crib}(S_1, \Omega)$ is a full component associated with S_1 and is outbound. Therefore, (S_1, A_1) is an O-block and is feasible since every member of $\mathcal{C}_1 \subseteq \text{support}(\Omega)$ is feasible as potential maximal clique Ω is feasible. Thus, the second case in the definition of feasible potential maximal cliques applies.

Finally, suppose that $S \neq \Omega$. Let $A = \text{crib}(S, \Omega)$. Then, A is a full component associated with S and is outbound, by Lemmas 3 and 8. Since $S = \bigcup_{C \in \text{support}(\Omega)} N(C)$ and Ω is feasible, the O-block (S, A) is feasible. Let x be an arbitrary vertex in $\Omega \setminus S$. Since we are assuming that $S \cup \text{outlet}(\Omega) = \Omega$ we have $x \in \text{outlet}(\Omega) \setminus S$. Let v be an arbitrary vertex in $\Omega \setminus \text{outlet}(\Omega)$. Observe that there is no component C associated with Ω such that $N(C)$ contains both x and v : $x \notin N(C)$ for every $C \in \text{support}(\Omega)$ and $v \notin N(C)$ for every C that is confined to $\text{outlet}(\Omega)$. Since Ω is cliquish, it follows that x and v are adjacent to each other. Therefore, we have $\Omega \setminus S \subseteq N(v)$. Moreover, A contains $\Omega \setminus S$ by Lemma 3. Finally, $A \setminus \Omega$ is disjoint from $N(v)$, since every component D associated with Ω such that $v \in N(D)$ is not confined to $\text{outlet}(\Omega)$ and hence contained in C . Therefore, we have $\Omega = S \cup (N(v) \cap A)$, and the third case in the definition of buildable potential maximal cliques applies. \square

4 Algorithm

Given graph G and positive integer k , our algorithm generates all I-blocks, O-blocks, and potential maximal cliques that are feasible. In the algorithm description below, the following variables, with suffixes, are used: \mathcal{I} for listing feasible I-blocks, \mathcal{O} for feasible O-blocks, \mathcal{P} for buildable potential maximal cliques, and \mathcal{S} for feasible potential maximal cliques. We note that each member of \mathcal{I} and \mathcal{O} is actually the component part of an I- or O-block.

Algorithm PID-BT

Input Graph G and positive integer k

Output “YES” if $\text{tw}(G) \leq k$; “NO” otherwise

Procedure

1. Let $\mathcal{I}_0 = \emptyset$ and $\mathcal{O}_0 = \emptyset$.
2. Initialize \mathcal{P}_0 and \mathcal{S}_0 to \emptyset .

3. Set $j = 0$.
4. For each $v \in V(G)$, if $N[v]$ is a potential maximal clique with $|N[v]| \leq k + 1$ then add $N[v]$ to \mathcal{P}_0 and if, moreover, $\text{support}(N[v]) = \emptyset$ then do the following.
 - (a) Add $N[v]$ to \mathcal{S}_0 .
 - (b) If $\text{outlet}(N[v]) \neq \emptyset$ then let $C = \text{crib}(\text{outlet}(N[v]), N[v])$ and, provided that $C \neq C_h$ for $1 \leq h \leq j$, increment j and let $C_j = C$.
5. Set $i = 0$.
6. Repeat the following and stop repetition when j is not incremented during the iteration step.
 - (a) While $i < j$, do the following.
 - i. Increment i and let \mathcal{I}_i be $\mathcal{I}_{i-1} \cup \{C_i\}$.
 - ii. Initialize \mathcal{O}_i to \mathcal{O}_{i-1} , \mathcal{P}_i to \mathcal{P}_{i-1} , and \mathcal{S}_i to \mathcal{S}_{i-1} .
 - iii. For each $B \in \mathcal{O}_{i-1}$ such that $C_i \subseteq B$ and $|N(C_i) \cup N(B)| \leq k + 1$, let $K = N(C_i) \cup N(B)$ and do the following.
 - A. If K is a potential maximal clique, then add K to \mathcal{P}_i .
 - B. If $|K| \leq k$ and there is a full component A associated with K (which is unique), then add A to \mathcal{O}_i .
 - iv. Let A be the full component associated with $N(C_i)$ and add A to \mathcal{O}_i .
 - v. For each $A \in \mathcal{O}_i \setminus \mathcal{O}_{i-1}$ and $v \in N(A)$, let $K = N(A) \cup (n(v) \cap A)$ and if $|K| \leq k + 1$ and K is a potential maximal clique then add K to \mathcal{P}_i .
 - vi. For each $K \in \mathcal{P}_i \setminus \mathcal{S}_{i-1}$, if $\text{support}(K) \subseteq \mathcal{I}_i$ then add K to \mathcal{S}_i and do the following: if $\text{outlet}(K) \neq \emptyset$ then let $C = \text{crib}(\text{outlet}(K), K)$ and, provided that $C \neq C_h$ for $1 \leq h \leq j$, increment j and let $C_j = C$.
7. If there is some $K \in \mathcal{S}_j$ such that $\text{outlet}(K) = \emptyset$, then answer “YES”; otherwise, answer “NO”.

Theorem 1 Algorithm PID-BT, given G and k , answers “YES” if and only if $\text{tw}(G) \leq k$.

Proof We show that \mathcal{S}_J computed by the algorithm, where J denotes the final value of j , is exactly the set of feasible potential maximal cliques for the given G and k . The theorem then follows by Lemma 5.

In the following proof, \mathcal{O}_i , \mathcal{P}_i , and \mathcal{S}_i for each i stand for the final values of these program variables.

We first show by induction on i that the following conditions are satisfied.

1. For every $1 \leq h \leq i$, $(N(C_j), C_j)$ is a feasible I-block.
2. $\mathcal{S}_i = \{C_h \mid 1 \leq h \leq i\}$.
3. For every $A \in \mathcal{O}_i$, $(N(A), A)$ is a feasible O-block.
4. Every $K \in \mathcal{P}_i$ is a buildable potential maximal clique.
5. Every $K \in \mathcal{S}_i$ is a feasible potential maximal clique.

Consider the base case $i = 0$. Condition 1 vacantly holds. Conditions 2 and 3 also hold since $\mathcal{S}_0 = \mathcal{O}_0 = \emptyset$. Condition 4 holds: $N[v]$ is confirmed to be a potential maximal clique before it is added to \mathcal{P}_0 and is buildable by the definition of buildability

(case 1). Condition 5 holds since $\text{support}(N[v]) = \emptyset$ implies that the potential maximal clique $N[v]$ is feasible.

Suppose $i > 0$ and that the above conditions are satisfied for smaller values of i .

1. When C_i is defined, there is some $i' < i$ and $K \in \mathcal{S}_{i'}$ such that $\text{outlet}(K) \neq \emptyset$ and $C_i = \text{crib}(\text{outlet}(K), K)$. By the induction hypothesis, K is a feasible potential maximal clique and hence, by Lemma 7, $(N(C_i), C_i)$ is a feasible I-block.
2. As $\mathcal{S}_{i-1} = \{C_h \mid 1 \leq h \leq i - 1\}$ and $\mathcal{S}_i = \mathcal{S}_{i-1} \cup \{C_i\}$, $\mathcal{S}_i = \{C_h \mid 1 \leq h \leq i\}$ holds.
3. Let $A \in \mathcal{O}_i \setminus \mathcal{O}_{i-1}$. Then there is some $B \in \mathcal{O}_{i-1}$ such that A is outbound, $|N(A)| \leq k$, and $N(A) = N(C_i) \cup N(B)$. From the first two conditions, $(N(A), A)$ is an O-block. By the induction hypothesis, $(N(B), B)$ is a feasible O-block and hence $N(B) = \bigcup_{D \in \mathcal{C}} N(D)$ for some set \mathcal{C} of feasible inbound components. As C_i is feasible by 1 above and $N(A) = \bigcup_{D \in \mathcal{C} \cup \{C_i\}} N(D)$, O-block $(N(A), A)$ is feasible.
4. Let $K \in \mathcal{P}_i \setminus \mathcal{P}_{i-1}$. Then, K is added to \mathcal{P}_i either at step 6-(a)-iii-A or at step 6-(a)-v. Consider the first case, Then, $K = N(B) \cup N(C_i)$ where $(N(B), B)$ is a feasible O-block and hence $N(B) = \bigcup_{D \in \mathcal{C}} N(D)$ for some set \mathcal{C} of feasible inbound components. As C_i is feasible, K satisfies all the conditions in the second case of the definition of buildable potential maximal cliques. Consider next the second case, K is obtained at step 6-(a)-v. Then, $K = N(A) \cup (n(v) \cap A)$, where $(N(A), A)$ is a feasible O-block, and the third case in the definition of buildable potential maximal cliques applies.
5. Let $K \in \mathcal{S}_i \setminus \mathcal{S}_{i-1}$. Then, $K \in \mathcal{P}_i$ and is a buildable potential maximal clique by 4 above. The confirmed condition $\text{support}(K) \subseteq \mathcal{S}_i$ ensures that K is feasible, since every member of \mathcal{S}_i is feasible by 1 and 2 above.

We conclude that every member of \mathcal{S}_J is a feasible potential maximal clique.

In showing the converse, the following observation is crucial. Let $(N(A), A)$ be a feasible O-block such that $N(A) = \bigcup_{C \in \mathcal{C}} N(C)$ for some set \mathcal{C} of feasible components and suppose $\mathcal{C} \subseteq \mathcal{S}_i$. Then, $A \in \mathcal{O}_i$. The proof is a straightforward induction on i .

The proof of the converse consists in showing the following by induction on m .

1. For each feasible I-block $(N(C), C)$, with $|C| = m$, there is some i such that $C = C_i$.
2. For each feasible O-block $(N(A), A)$ with $|A| = |V(G)| - m$, there is some i such that $A \in \mathcal{O}_i$.
3. For each buildable potential maximal clique Ω such that $|\bigcup_{C \in \text{support}(\Omega)} C| = m$, there is some i such that $\Omega \in \mathcal{P}_i$.
4. For each feasible potential maximal clique Ω such that $|\bigcup_{C \in \text{support}(\Omega)} C| = m$, there is some i such that $\Omega \in \mathcal{S}_i$.

The base case $m = 0$ is vacantly true. Suppose $m > 0$ and the statements hold for smaller values of m .

1. Let $(N(C), C)$ be a feasible I-block with $|C| = m$. Then, by Lemma 7, there is some feasible potential maximal clique Ω such that $N(C) = \text{outlet}(\Omega)$ and $C = \text{crib}(N(C), \Omega)$. We have $|\bigcup_{C \in \text{support}(\Omega)} C| < m$, since this union is a subset

- of $C \setminus (\Omega \setminus N(C))$. Therefore, by the induction hypothesis, there is some i such that $\Omega \in \mathcal{S}_i$. Therefore, C is constructed as C_j either at step 4-(b) or at step 6-(a)-vi.
2. Let $(N(A), A)$ be a feasible O-block with $|A| = |V(G)| - m$. Let \mathcal{C} be a set of feasible components such that $N(A) = \bigcup_{C \in \mathcal{C}} N(C)$ and let C be an arbitrary member of \mathcal{C} . As C, A , and $N(C)$ are pairwise disjoint, we have $|C| < m$. Therefore, there is some i_C such that $C_{i_C} = C$. Set $i = \max\{i_C \mid C \in \mathcal{C}\}$. Then, $\mathcal{C} \subseteq \mathcal{S}_i$ and hence $A \in \mathcal{O}_i$, by the observation above.
 3. Let Ω be a buildable potential maximal clique with $|\bigcup_{C \in \text{support}(\Omega)} C| = m$. In the first case of the definition of buildability, Ω is added to \mathcal{P}_0 at step 4. In the second case, we have $\Omega = \bigcup_{C \in \mathcal{C}} N(C)$ for some $\mathcal{C} \subseteq \text{support}(\Omega)$ such that every member of \mathcal{C} is feasible. Choose \mathcal{C} to be minimal subject to these conditions. Let C be an arbitrary member of \mathcal{C} . As $|C| \leq m$, by the induction hypothesis and 1 above, there is some i_C such that $\mathcal{C} \subseteq \mathcal{S}_{i_C}$. Choose $C \in \mathcal{C}$ so that i_C is the largest and let the chosen be D . Let $\mathcal{C}' = \mathcal{C} \setminus \{D\}$ and let $S = \bigcup_{C \in \mathcal{C}'} N(C)$. By the minimality of \mathcal{C} , S is a proper subset of Ω . Therefore, $\text{crib}(S, \Omega)$ is a full component associated with S and there is an outbound full component A associated with S . As all members of \mathcal{C}' is feasible and $|S| \leq k$, (S, A) is a feasible O-block. By the choice of D , we have $\mathcal{C}' \subseteq \mathcal{S}_{i_D-1}$ and hence $A \in \mathcal{O}_{i_D-1}$ by the observation above. At step 6-(a)-iii-A in the iteration for $i = i_D$, Ω is put into \mathcal{P}_{i_D} .
 4. Let Ω be a feasible potential maximal clique with $|\bigcup_{C \in \text{support}(\Omega)} C| = m$. Then, by 3 above, there is some i_1 such that $\Omega \in \mathcal{P}_{i_1}$. Furthermore, as every member C of $\text{support}(\Omega)$ is feasible and $|C| \leq m$, there is some i_2 such that $\text{support}(\Omega) \subseteq \mathcal{S}_{i_2}$, by 1 above. At step 7 in the iteration for $i = \max\{i_1, i_2\}$, Ω is put into \mathcal{S}_i .

We conclude that every feasible potential maximal clique is in \mathcal{S}_J . This completes the proof. □

5 Running time analysis

The running time of our algorithm is stated in terms of the the number of positive subproblem instances. Given G and $k > 0$, let \mathcal{I}_G^k denote the set of feasible I-blocks and \mathcal{O}_G^k the set of feasible O-blocks.

Observation 1 *Given G and $k > 0$, algorithm PID-BT runs in $O^*(|\mathcal{I}_G^k| \cdot |\mathcal{O}_G^k|)$ time.*

Proof The number of iteration in step 6, where i is incremented each time, is $|\mathcal{I}_G^k|$. In each iteration step, every computation step may be charged to each element of \mathcal{O}_{i-1} and the total number of steps charged to a single element of \mathcal{O}_{i-1} is $n^{O(1)}$. Since $|\mathcal{O}_{i-1}| \leq |\mathcal{O}_G^k|$, we have the claimed time bound. □

The bound in this observation is incomparable to the previous bounds on non-PID versions of the BT algorithm, which run in $O^*(|\Pi_G|)$ time when Π_G , the set of potential maximal cliques in G , is given. In Fomin and Villanger (2012), in addition to a combinatorial bound of $|\Pi_G| = O(1.7549^n)$, it was shown that Π_G can be computed in $O^*(\Pi_G)$ time.

It should be emphasized, however, that it is not known whether the decision problem version of the treewidth problem with given k can be solved in $O^*(|\Pi_G^{k+1}|)$ time, where Π_G^k is the set of potential maximal cliques of cardinality at most k in G . The bottleneck here is the time to list all members of Π_G^{k+1} . Although a nontrivial upper bound on $|\Pi_G^{k+1}|$ in terms of n and k , together with a running time bound based on it, is given in Fomin and Villanger (2012), a huge gap between the actual value $|\Pi_G^{k+1}|$ and the upper bound is observed in practice, as shown in the next section. This is the gap that makes the bound in Observation 1 interesting.

6 Experimental analysis

To study the strength of the running time bound of Observation 1 from a practical view point, we have performed some experiments, in which we count the number of combinatorial objects involved in the treewidth computation. We first compare the actual number of relevant potential maximal cliques (that is, of cardinality at most $k + 1$ where k is the treewidth) with the theoretical upper bounds on that number: the naive bound of $\binom{n}{k+1}$ and an asymptotically stronger bound of $n \left(\binom{\lceil (2n+k+7)/3 \rceil}{k+2} + \binom{\lceil (2n+k+4)/2 \rceil}{k+1} \right)$ given in Fomin and Villanger (2012). Table 1 shows the results on some random instances, where the number of vertices n is 20, 30, 40, or 50, the number of edges m is $2n$, $3n$, $4n$ or $5n$, and the graph for each pair (n, m) is chosen uniformly at random from the set of all graphs with n vertices and m edges. Huge gaps between the actual number and the upper bounds are apparent.

Since the running time bound in Observation 1 involves the quantity $|\mathcal{O}_G^k|$ which is not theoretically upper-bounded by a function of $|\Pi_G^{k+1}|$, the gaps observed in Table 1 alone may not be sufficient to support the importance of this running time bound. To address this issue, we have counted more combinatorial objects involved in our PID computation on the same graph instances: in addition to relevant potential maximal cliques counted above, all potential maximal cliques, relevant minimal separators, all minimal separators, feasible I-blocks, feasible O-blocks and feasible potential maximal cliques. Here, the input k to the decision problem is set to the treewidth of the graph.

Table 2 shows the result. We see that the number of feasible O-blocks is smaller than the number of relevant potential maximal cliques, as far as these instances are concerned. This, together with what we have observed in Table 1, provides an evidence that the running time bound of Observation 1 is more relevant from a practical point of view than the running time bounds of known theoretical algorithms.

We also see that the number of all potential maximal cliques grows much faster than the number of relevant potential maximal cliques. This shows the advantage of our algorithm which avoids generating all potential maximal cliques.

To summarize, our PID algorithm has advantages over the standard BT algorithms because the running time upper bounds of those algorithms are either in terms of a combinatorial *upper bound* on the number of relevant potential maximal cliques or in terms of the actual number of *all* potential maximal cliques: our experiments reveal huge gaps between the actual number of relevant potential maximal cliques and both

Table 1 The numbers of relevant potential maximal cliques and their upper bounds

$n = V $	$ E $	$k = tw$	PMCs ($\leq k + 1$)	$\binom{n}{k+1}$	$n \left(\frac{\binom{2n+k+7}{31} + \binom{k+2}{\binom{(n+k+4)/2}{k+1}} \right)$
20	40	6	115	77,520	1,003,860
20	60	8	96	167,960	2,076,360
20	80	11	121	125,970	1,921,680
20	100	11	37	125,970	1,921,680
30	60	7	559	5,852,925	67,393,950
30	90	11	682	86,493,225	352,580,340
30	120	14	1137	155,117,520	430,361,970
30	150	16	768	119,759,850	426,140,550
40	80	8	5341	273,438,880	2,705,471,600
40	120	14	10,372	40,225,345,056	91,260,807,600
40	160	18	17,360	131,282,408,400	135,562,547,400
40	200	20	6820	131,282,408,400	157,012,867,200
50	100	10	6029	37,353,738,800	201,991,095,800
50	150	16	48,068	9,847,379,391,150	10,332,510,412,500
50	200	20	36,388	67,327,446,062,800	53,246,262,826,500
50	250	24	47,729	126,410,606,437,752	52,230,760,068,000

of these quantities. Note that, if there is an efficient method of generating relevant potential maximal cliques, a non-PID version of the BT algorithm might outperform our PID version.

7 Implementation

In this section, we sketch two important ingredients of our implementation. Although both are crucial in obtaining the result reported in Sect. 8, our work on this part is preliminary and improvements are the subject of future research.

7.1 Data structures

The crucial elementary operation in our algorithm is the following. We have a set \mathcal{O} of feasible O-blocks obtained so far and, given a new feasible I-block $(N(C), C)$, need to find all members $(N(A), A)$ of \mathcal{O} such that $C \subseteq A$ and $|N(C) \cup N(A)| \leq k + 1$. As the experimental analysis in the previous section shows, there is only a few such A on average for the tested instances even though \mathcal{O} is usually huge. To support an efficient query processing, we introduce an abstract data structure we call a block sieve.

Let G be a graph and k a positive integer. A *block sieve* for graph G and width k is a data structure storing vertex sets of $V(G)$ which supports the following operations.

- store(U) : store vertex set U in in the block sieve.
- supersets(U) : return the list of entries W stored in the block sieve such that $U \subseteq W$ and $|N(U) \cup N(W)| \leq k + 1$.

Data structures for superset query have been studied (Savnik 2013). The second condition above on the retrieved sets, however, appears to make this data structure new. For each $U \subseteq V(G)$, we define the *margin* of U to be $k + 1 - |N(U)|$. Our implementation of block sieves described below exploits an upper bound on the margins of vertex sets stored in the sieve.

We first describe how such block sieves with upper bounds on margins are used in our algorithm. Let \mathcal{O} be the current set of O-blocks. We use t block sieves $\mathcal{B}_1, \dots, \mathcal{B}_t$, each \mathcal{B}_i having a predetermined upper bound m_i on the margins of the sets stored. We have $0 < m_1 < m_2 < \dots < m_t = k$. We set $m_0 = 0$ for notational ease below. In our implementation, we choose roughly $t = \log_2 k$ and $m_i = 2^i$ for $0 < i < t$. For each $(N(A), A)$ in \mathcal{O} , A is stored in \mathcal{B}_i such that the margin $k + 1 - |N(A)|$ is m_i or smaller but larger than m_{i-1} . When we are given an I-block $(N(C), C)$ and are to list relevant blocks in \mathcal{O} , we query all of the t blocks with the operations supersets(C). These queries as a whole return the list of all vertex sets A such that $(N(A), A) \in \mathcal{O}$, $C \subseteq A$, and $|N(A) \cup N(C)| \leq k + 1$.

We implement a block sieve by a trie \mathcal{T} . The upper bound m on margin is not used in the construction of the sieve; it is used in the query time. In the following, we assume $V(G) = \{1, \dots, n\}$ and, by an interval $[i, j]$, $1 \leq i \leq j \leq n$, we mean the set $\{v : i \leq v \leq j\}$ of vertices. Each non-leaf node p of \mathcal{T} is labelled with a non-empty interval $[s_p, f_p]$, such that $s_r = 0$ for the root r , $s_p = f_q + 1$ if p is a child of q , and $f_p = n$ if p is a parent of a leaf. Each edge (p, q) which connects node p and a child

Table 2 The numbers of principal objects in treewidth computation

V	E	tw	Minimal separators		PMCs		Feasible objects		
			All	$\leq tw$	All	$\leq tw + 1$	I-blocks	O-blocks	PMCs
20	40	6	98	51	376	115	19	26	37
20	60	8	191	48	796	96	46	108	93
20	80	11	185	122	698	376	121	158	370
20	100	11	107	25	354	37	24	32	36
30	60	7	535	185	3122	559	114	170	334
30	90	11	2983	247	20,154	682	228	708	618
30	120	14	2713	376	16,736	1137	352	804	1055
30	150	16	1913	281	10,535	768	240	498	647
40	80	8	14,842	1070	178,661	5341	840	2965	4154
40	120	14	164,773	2356	1,740,644	10,372	2080	8637	8577
40	160	18	134,485	3952	1,251,656	17,360	3289	10,023	13,646
40	200	20	52,182	1790	423,691	6820	1502	4749	5347
50	100	10	96,499	1361	1,123,621	6029	779	2171	2914
50	150	16	1,792,713	9152	>2,000,000	48,068	8099	36,881	39,803
50	200	20	2,130,811	7878	>2,000,000	36,388	6956	28,247	29,842
50	250	24	1,452,449	10,571	>2,000,000	47,729	8949	30,834	37,115

q of p , is labelled with a subset $S_{(p,q)}$ of the interval $[s_p, f_p]$. Thus, for each node p , the union of the labels of the edges along the path from the root to p is a subset of the interval $[1, s_p - 1]$, or $[1, n]$ when p is a leaf, which we denote by S_p . The choice of interval $[s_p, f_p]$ for each node p is heuristic. It is chosen so that the number of descendants of p is not too large or too small. In our implementation, the interval size is adaptively chosen from 8, 16, 32, and 64.

Each leaf q of trie \mathcal{T} represents a single set stored at this leaf, namely S_q as defined above. We denote by $S(\mathcal{T})$ the set of all sets stored in \mathcal{T} . Then, for each node p of \mathcal{T} , the set of sets stored under p is $\{U \mid U \cap [1, p] = S_p\}$.

We now describe how a query is processed against this data structure. Suppose query U is given. The goal is to visit all leaves q such that $U \subseteq S_q$ and $|N(U) \cup N(S_q)| \leq k + 1$. This is done by a depth-first traversal of the trie \mathcal{T} . When we visit node p , we have the invariant that $U \cap [1, f_p] \subseteq S_p$, since otherwise no leaf in the subtree rooted at p stores a superset of U . Therefore, we descend from p to a child p' of p only if this invariant is maintained. Moreover, we keep track of the quantity $i(p, U) = |N(U) \cap S_p|$ in order to make further pruning of search possible. For each leaf q below p such that $U \subseteq S_q$, we have $i(q, U) \geq i(p, U)$. Combining this with equality $|N(U) \setminus N(S_q)| = |N(U) \cap S_q| = i(q, U)$, we have $|N(U) \cup N(S_q)| \geq |N(S_q)| + i(p, U)$. Since we know an upper bound m on the margin $k + 1 - |N(S_q)|$ of S_q , or lower bound $k + 1 - m$ on $|N(S_q)|$, we may prune the search under node p if $i(p, U) > m$, since this inequality implies $|N(U) \cup N(S_q)| > k + 1$ for every leaf q under p . When we reach a leaf q , we test if $|N(U) \cup N(S_q)| \leq k + 1$ indeed holds.

7.2 Safe separators

The notion of safe separators for treewidth was introduced by Bodlaender and Koster (2006): a separator S of G is *safe* if completing S into a clique does not change the treewidth of G . If we find a safe separator S then the problem of deciding tree width of G reduces to that of deciding the treewidth of $G(C)$ for each component C associated with S . Preprocessing G into such independent subproblems is highly desirable whenever possible.

The above authors observed that a powerful sufficient condition for safeness can be formulated based on graph minors. A *labelled minor* of G is a graph obtained from G by zero or more applications of the following operations. (1) Edge contraction: choose an edge $\{u, v\}$, replace u and v by a single new vertex and let all neighbors of u and v be adjacent to this new vertex; name the new vertex as either u or v . (2) Vertex deletion: delete a vertex together with all incident edges. (3) Edge deletion.

Lemma 10 (Bodlaender and Koster 2006) *A separator S of G is safe if, for every component C associated with S , $G[V(G) \setminus C]$ contains clique S as a labelled minor.*

Call a separator *minor-safe* if it satisfies the sufficient condition for safeness stated in this lemma. Bodlaender and Koster (2006) showed that if S is a minimal separator and is an almost clique (deleting some single vertex makes it a clique) then S is minor-safe and moreover that the set of all almost clique minimal separators can be found in $O(n^2m)$ time, where n is the number of vertices and m is the number of edges.

We aim at capturing as many minor-safe separators as possible, at the expense of theoretical running time bounds on the algorithm for finding them. Thus, in our approach, both the algorithm for generating candidate separators and the algorithm for deciding minor-safeness are heuristic. For candidate generation, we use greedy heuristic for treewidth such as min-fill and min-degree: the separators in the resulting tree-decomposition are all candidates for safe separators.

When we apply our heuristic decision algorithm for minor-safeness to candidate separator S , one of the following occurs.

1. The algorithm answers “YES”. In this case, a required labelled clique minor has been found for every component associated S and hence S is minor-safe.
2. The algorithm answers “DON’T KNOW”. In this case, the algorithm has failed to find a labelled clique minor for at least one component, and hence it is not known if S is minor-safe or not.
3. The algorithm aborts, after reaching the prescribed number of execution steps.

Our heuristic decision algorithm works in two phases. Let S be a separator, C a component associated with S , and $R = V(G) \setminus (S \cup C)$. In the first phase, we contract edges in R and obtain a graph B on vertex set $S \cup R'$, where each vertex of R' is a contraction of some vertex set of R and B has no edge between vertices in R' . For each pair u, v of distinct vertices in S , let $N(u, v)$ denote the common neighbors of u and v in graph B . The contractions are performed with the goal of making $|N(u, v) \cap R'|$ large for each missing edge $\{u, v\}$ in S . In the second phase, for each missing edge $\{u, v\}$, we choose a common neighbor $w \in N(u, v) \cap R'$ and contract either $\{u, w\}$ or $\{v, w\}$. The choice of the next missing edge to be processed and the choice of the common neighbor are done as follows. Suppose the contractions in the second phase are done for some missing edges in S . For each missing edge $\{u, v\}$ not yet “processed”, let $N'(u, v)$ be the set of common neighbors of u and v that are not yet contracted with any vertex in S . We choose $\{u, v\}$ with the smallest $|N'(u, v) \cap R'|$ to be processed next. Tie-breaking when necessary and the choice of the common neighbor w in $N'(u, v) \cap R'$ to be contracted with u or v is done in such a way that the minimum of $|(N'(x, y) \cap R') \setminus \{w\}|$ is maximized over all remaining missing edges $\{x, y\}$ in S .

The performance of these heuristics strongly depends on the instances. For PACE 2017 public instances, they work quite well. Table 3 shows the preprocessing result on the last 10 of those instances. See Sect. 8 for the description of those instances and the computational environment for the experiment. For each instance, the number of safe separators found and the maximum subproblem size in terms of the number of vertices, after the graph is decomposed by the safe separators found, are listed. The results show that these instances, which are deemed the hardest among all the 100 public instances, are quickly decomposed into manageable subproblems by our preprocessing.

On the other hand, these heuristics turned out useless for most of the DIMACS graph coloring instances: no safe separators are found for those instances. We suspect that this is not the limitation of the heuristics but is simply because those instances lack minor-safe separators. We need, however, further study to get a firm conclusion.

Table 3 Safe separator preprocessing on PACE 2017 instances

Name	$ V $	$ E $	$tw(G)$	Safe separators found	Max subproblem	Time (s)
ex181	109	732	18	18	89	0.078
ex183	265	471	11	173	76	0.031
ex185	237	793	14	142	52	0.046
ex187	240	453	10	138	81	0.031
ex189	178	4517	70	6	161	0.062
ex191	492	1608	15	184	132	0.171
ex193	1391	3012	10	791	119	3.17
ex195	216	382	10	114	84	0.015
ex197	303	1158	15	176	56	0.062
ex199	310	537	9	157	131	0.046

Table 4 Results on the DIMACS graph coloring instances

Name	$ V $	$ E $	tw	Time(s)	Prev UB	Prev LB
Anna	138	493	12	0.078	12	12
David	87	406	13	0.031	13	13
DSJC125.5	125	3891	108	459	108	56
DSJC125.9	125	6961	119	0.062	119	119
DSJC250.9	250	27,897	243	0.44	243	212
DSJC500.9	500	112,437	492	14	492	433
DSJR500.5	500	58,862	246	546	–	–
DSJR500.1c	500	121,275	485	2.12	485	485
fpsol2.i.1	496	11,654	66	3.30	66	66
fpsol2.i.2	451	8691	31	5.66	31	31
fpsol2.i.3	425	8688	31	5.68	31	31
games120^a	120	638	32	94,738	32	24
homer^a	561	1628	30	2765	31	26
huck	74	301	10	0.012	10	10
inithx.i.1	864	18,707	56	8.10	56	56
inithx.i.2	645	13,979	31	8.14	31	31
inithx.i.3	621	13,969	31	10	31	31
jean	80	254	9	0.031	9	9
miles250	128	387	9	0.000	9	9
miles500	128	1170	22	0.11	22	22
miles750	128	2113	36	0.23	36	35
miles1000	128	3216	49	0.33	49	49
miles1500	128	5198	77	0.45	77	77
mulsol.i.1	197	3925	50	1.41	50	50
mulsol.i.2	188	3885	32	1.77	32	32
mulsol.i.3	184	3916	32	1.80	32	32
mulsol.i.4	185	3946	32	1.78	32	32
mulsol.i.5	186	3973	31	1.80	31	31
myciel2	5	5	2	0.000	2	2
myciel3	11	20	5	0.000	5	5
myciel4	23	71	10	0.015	10	10
myciel5	47	236	19	0.33	19	19
myciel6	95	755	35	419	35	29
queen5_5	25	160	18	0.000	18	18
queen6_6	36	290	25	0.031	25	25
queen7_7	49	476	35	0.19	35	35
queen8_8	64	728	45	4.16	45	25
queen9_9	81	1056	58	274	58	35
queen8_12	96	1368	65	649	–	39
queen10_10	100	1470	72	20,934	72	39

Table 4 continued

Name	$ V $	$ E $	tw	Time(s)	Prev UB	Prev LB
zeroin.i.1	211	4100	50	1.09	50	50
zeroin.i.2	211	3541	32	1.64	32	32
zeroin.i.3	206	3540	32	1.55	32	31

Previous upper bounds from Gogate and Dechter (2004) and Musliu (2008); previous lower bounds from Gogate and Dechter (2004) and Bodlaender et al. (2006).

^a24GB heap space is used for these instances

8 Performance results

We have tested our implementation on two sets of instances. The first set comes from the DIMACS graph coloring challenge (Johnson 1996) and has served as a standard benchmark suite for treewidth in the literature (Gogate and Dechter 2004; Bodlaender and Koster 2006; Musliu 2008; Samer and Veith 2009; Bodlaender et al. 2012; Berg and Järvisalo 2014). The other is the set of public instances posed by the exact treewidth track (PACE 2017).

The computing environment for the experiment is as follows. CPU: Intel Core i7-7700K, 4.20GHz; RAM: 32GB; Operating system: Windows 10, 64bit; Programming language: Java 1.8; JVM: jre1.8.0_121. The maximum heap space size is 6GB by default and is 24GB where it is stated so. The implementation is single threaded, except that multiple threads may be invoked for garbage collection by JVM. The time measured is the CPU time, which includes the garbage collection time.

To determine the treewidth of a given instance we use our decision procedure with k being incremented one by one, starting from the obvious lower bound, namely the minimum degree of the graph. Binary search is not used because the cost of overshooting the exact treewidth can be huge. We do not feel the need of using stronger lower bounds either, since the cost of executing the decision procedure for k below such lower bounds is usually quite small.

Table 4 shows the results on DIMACS graph coloring instances. Each row shows the name of the instance, the number of vertices, the number of edges, the exact treewidth computed by our algorithm, CPU time in seconds, and the previously best known upper and lower bounds on the treewidth. Rows in bold face show the newly solved instances. For all but three of them, the previous best upper bound has turned out optimal: only the lower bound was weaker. In this experiment, however, no knowledge of previous bounds are used and our algorithm independently determines the exact treewidth.

The results on “queen” instances illustrate how far our algorithm has extended the practical limit of exact treewidth computation. Queen7_7 with 49 vertices is the largest instance previously solved, while queen10_10 with 100 vertices is now solved. Also note that all previously solved instances are fairly easy for our algorithm: all of them are solved within 10 seconds per instance and many of them within a second.

Table 5 shows the lower bounds obtained by our algorithm on unsolved DIMACS graph coloring instances. Lower bound entries in bold face are improvements over the previously known lower bounds. Computation time of the previously best lower bounds

Table 5 New lower bounds on the treewidth of unsolved DIMACS graph coloring instances

Name	V	E	Lower bounds computed			Previous bounds	
			1 sec	1 min	30 min	Lower	Upper
DSJC125.1	125	736	25	30	36	20	60
DSJC250.1	250	3218	45	57	66	43	167
DSJC250.5	250	15,668	180	197	211	114	229
DSJC500.1	500	12,458	–	94	115	87	409
DSJC500.5	500	62,624	–	360	388	231	479
DSJC1000.1	1000	49,629	–	172	189	183	896
DSJC1000.5	1000	249,826	–	724	742	469	977
DSJC1000.9	1000	449,449	–	983	987	872	991
le450_5a	450	5714	29	50	59	79	243
le450_5b	450	5734	–	49	57	–	246
le450_5c	450	9803	–	84	100	106	265
le450_5d	450	9757	–	94	99	–	265
le450_15a	450	8168	24	40	49 ^a	94	262
le450_15b	450	8169	23	32	47 ^a	–	258
le450_15c	450	16,680	–	114	132	139	350
le450_15d	450	16,750	–	112	131	–	353
le450_25a	450	8260	11	23	25 ^a	96	216
le450_25b	450	8263	16	26	30 ^a	–	219
le450_25c	450	17,343	43	89	109	144	320
le450_25d	450	17,425	–	93	112	–	327
myciel7	191	2360	22	31	35	52	66

Table 5 continued

Name	V	E	Lower bounds computed			Previous bounds	
			1 sec	1 min	30 min	Lower	Upper
queen11_11	121	1980	61	70	77	40	87
queen12_12	144	2596	71	76	84	55	103
queen13_13	169	3328	70	82	91	51	121
queen14_14	196	4186	74	87	98	55	140
queen15_15	225	5180	78	93	104	73	162
queen16_16	256	6320	83	99	110	79	186
school1	385	19,095	73	112	125	149	178
school1_nsh	352	14,612	78	105	118	132	152

Previous upper bounds from Musliu (2008); previous lower bounds from Bodlaender et al. (2006).

^aOut of memory before time out

Table 6 Results on the PACE 2017 public instances

Name	V	E	tw	Time (s)	Name	V	E	tw	Time (s)
ex001	262	648	10	1.48	ex101	1038	29,1034	540	12
ex003	92	2113	44	8.92	ex103	237	419	10	3.01
ex005	377	597	7	14	ex105	1038	29,1037	540	12
ex007	137	451	12	0.046	ex107	166	396	12	1.44
ex009	466	662	7	13	ex109	1212	1794	7	43
ex011	465	1004	9	0.50	ex111	395	668	9	4.33
ex013	56	280	29	15	ex113	93	488	14	0.046
ex015	177	669	15	0.046	ex115	963	419,877	908	18
ex017	330	571	9	1.11	ex117	77	181	13	18
ex019	291	752	11	40	ex119	84	479	23	16
ex021	318	572	9	2.80	ex121	204	1164	34	76
ex023	690	1355	8	0.91	ex123	122	635	35	14
ex025	92	472	20	1.61	ex125	320	8862	70	8.19
ex027	274	715	11	51	ex127	228	527	10	0.20
ex029	238	411	9	1.33	ex129	737	2826	14	0.97
ex031	219	382	8	12	ex131	292	1386	18	0.17
ex033	363	541	7	50	ex133	522	1296	11	3.94
ex035	247	804	14	3.60	ex135	2822	129,474	87	49
ex037	272	615	10	3.43	ex137	196	1098	19	0.34
ex039	56	280	32	58	ex139	334	568	9	8.34
ex041	205	341	9	0.63	ex141	226	1168	34	117
ex043	279	513	9	3.34	ex143	130	660	35	52
ex045	600	865	7	7.80	ex145	48	96	12	18

Table 6 continued

Name	V	E	tw	Time (s)	Name	V	E	tw	Time (s)
ex047	1854	21,118	21	140	ex147	101	606	16	0.093
ex049	117	332	13	0.078	ex149	698	2604	12	0.75
ex051	136	254	10	0.62	ex151	279	733	12	210
ex053	218	383	9	1.98	ex153	772	11654	47	57
ex055	197	813	18	0.078	ex155	758	11580	47	103
ex057	281	9075	117	0.093	ex157	260	467	9	6.42
ex059	298	780	10	0.47	ex159	582	2772	18	2.37
ex061	158	1058	22	9.59	ex161	1046	3906	12	2.84
ex063	103	582	34	4.76	ex163	244	445	10	4.69
ex065	50	175	25	79	ex165	222	742	14	0.23
ex067	235	424	10	2.70	ex167	509	969	10	7.96
ex069	235	441	9	1.43	ex169	3706	42,236	22	530
ex071	253	434	9	2.42	ex171	647	2175	14	0.77
ex073	712	1085	7	15	ex173	536	1011	10	5.05
ex075	111	360	8	0.28	ex175	227	1000	17	113
ex077	237	423	10	2.70	ex177	227	759	14	0.23
ex079	314	4943	42	1.64	ex179	187	346	10	14
ex081	188	638	6	0.55	ex181	109	732	18	0.20
ex083	213	380	10	3.05	ex183	265	471	11	8.61
ex085	229	370	8	11	ex185	237	793	14	0.33
ex087	380	5790	47	46	ex187	240	453	10	2.80
ex089	318	576	9	11	ex189	178	4517	70	3.59

Table 6 continued

Name	$ V $	$ E $	tw	Time (s)	Name	$ V $	$ E $	tw	Time (s)
ex091	193	336	9	31	ex191	492	1608	15	21
ex093	454	664	7	27	ex193	1391	3012	10	3.80
ex095	220	555	11	0.59	ex195	216	382	10	6.11
ex097	286	4079	48	2.01	ex197	303	1158	15	0.36
ex099	616	923	7	88	ex199	310	537	9	23

ranges from a few minutes to a week (Bodlaender et al. 2006). Detailed comparison of lower bound methods, which requires the normalization of machine speeds, is not intended here. Rather, the table is meant to show the potential of our algorithm as a lower bound procedure.

For many of the instances the improvements are significant. It can also be seen from this table that our algorithm performs rather poorly on relatively sparse graphs with a large number of vertices.

Table 6 shows the results on PACE 2017 instances. The prefix “ex” in the instance names means that they are for the exact treewidth track. Odd numbers mean that they are public instances disclosed prior to the competition for testing and experimenting. Even numbered instances, not in the list, are secret and to be used in evaluating submissions. The time allowed to be spent for each instance is 30 minutes. As can be seen from the table, our algorithm solves all of the public instances with a large margin in time.

Acknowledgements The author thanks Hiromu Ohtsuka for his help in implementing the block sieve data structure. He also thanks Yasuaki Kobayashi for helpful discussions and especially for drawing the author’s attention to the notion of safe separators. This work would have been non-existent if not motivated by the timely challenges of PACE 2016 and 2017. The author is deeply indebted to their organizers, especially Holger Dell, for their dedication and excellent work. Funding was provided by JSPS KAKENHI (Grant No. JP26330021).

References

- Amborg S, Corneil DG, Proskurowski A (1987) Complexity of finding embeddings in a k -tree. *SIAM J Algebr Discrete Methods* 8:277–284
- Berg J, Järvisalo M (2014) SAT-based approaches to treewidth computation: an evaluation. In: Proceedings of the IEEE 26th international conference on tools with artificial intelligence, pp 328–335
- Bodlaender HL (1996) A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J Comput* 25(6):1305–1317
- Bodlaender HL, Koster AMCA (2006) Safe separators for treewidth. *Discrete Math* 306(3):337–350
- Bodlaender HL, Koster AMCA (2008) Combinatorial optimization on graphs of bounded treewidth. *Comput J* 51(3):255–269
- Bodlaender HL, Fomin FV, Koster AM, Kratsch D, Thilikos DM (2012) On exact algorithms for treewidth. *ACM Trans Algorithms* 9(1):12
- Bodlaender HL, Wollé T, Koster AMCA (2006) Contraction and treewidth lower bounds. *J Graph Algorithms Appl* 10(1):5–49
- Bouchitté V, Todinca I (2001) Treewidth and minimum fill-in: grouping the minimal separators. *SIAM J Comput* 31(1):212–232
- Bouchitté V, Todinca I (2002) Listing all potential maximal cliques of a graph. *Theor Comput Sci* 276:17–32
- Dell H, Husfeldt T, Jansen B M, Kaski P, Komusiewicz C, Rosamond F A (2017) The first parameterized algorithms and computational experiments challenge. In: *LIPICs-Leibniz international proceedings in informatics* 63,
- Fomin FV, Kratsch D, Todinca I, Villanger Y (2008) Exact algorithms for treewidth and minimum fill-in. *SIAM J Comput* 38(3):1058–1079
- Fomin F, Villanger Y (2012) Treewidth computation and extremal combinatorics. *Combinatorica* 32(3):289–308
- Github repository: <https://github.com/TCS-Meiji/PACE2017-TrackA>
- Gogate V, Dechter R (2004) A complete anytime algorithm for treewidth. In: Proceedings of the 20th conference on Uncertainty in artificial intelligence, AUAI Press,

- Johnson DS, Trick MA (eds) (1996) Cliques, coloring, and satisfiability: second DIMACS implementation challenge, vol 26. Series in discrete mathematics and theoretical computer science. American Mathematical Society, Providence
- Musliu N (2008) An iterative heuristic algorithm for tree decomposition. In: Cotta C, van Hemert J (eds) Recent Advances in Evolutionary Computation for Combinatorial Optimization. Studies in Computational Intelligence, vol 153. Springer, Berlin, Heidelberg
- PACE (2017). website: <https://pacechallenge.wordpress.com/>
- Robertson N, Seymour PD (1986) Graph minors. II. Algorithmic aspects of tree-width. *J Algorithms* 7:309–322
- Robertson N, Seymour PD (2004) Graph minors. XX. Wagner’s conjecture. *J Comb Theory Ser B* 92(2):325–357
- Samer M, Veith H (2009) Encoding treewidth into SAT. In: Proceedings of international conference on theory and applications of satisfiability testing, pp 45–50
- Savnik I (2013) Index data structure for fast subset and superset queries. In: Proceedings of international conference on availability, reliability, and security, pp. 134–148,