

Certifying Induced Subgraphs in Large Graphs

Ulrich Meyer¹ Hung Tran^{1,2} Konstantinos Tsakalidis³

¹Goethe University Frankfurt, Germany

²Frankfurt Institute for Advanced Studies, Germany

³University of Liverpool, United Kingdom

Submitted:	Reviewed:	Revised:
Accepted:	Final:	Published:
Article type:		Communicated by:

Abstract.

We introduce I/O-efficient certifying algorithms for the recognition of bipartite, split, threshold, bipartite chain, and trivially perfect graphs. When the input graph is a member of the respective class, the certifying algorithm returns a certificate that characterizes this class. Otherwise, it returns a forbidden induced subgraph as a certificate for non-membership. On a graph with n vertices and m edges, our algorithms take $\mathcal{O}(\text{sort}(n + m))$ I/Os in the worst case for split, threshold and trivially perfect graphs. In the same complexity bipartite and bipartite chain graphs can be certified with high probability. We provide implementations and an experimental evaluation for split and threshold graphs.

1 Introduction

Certifying algorithms [19] ensure the correctness of an algorithm's output without having to trust the algorithm itself. The user of a certifying algorithm inputs x and receives the output y with a *certificate* or *witness* w that proves that y is a correct output for input x . In a subsequent step, the certificate can be inspected using an authentication algorithm that considers the input, output and certificate and returns whether the output is indeed correct. Certifying the bipartiteness of a graph is a textbook example where the returned witness w is a bipartition of the vertices (YES-certificate) or an *odd-length cycle* subgraph, i.e. a cycle of vertices with an odd number of edges (NO-certificate).

Konstantinos Tsakalidis was supported by the International Exchanges Grant IES\R3\203041 of the Royal Society. Ulrich Meyer and Hung Tran were supported by the Deutsche Forschungsgemeinschaft (DFG) under grant ME 2088/5-1 (FOR 2975 — Algorithms, Dynamics, and Information Flow in Networks). A preliminary version of this paper appeared in WALCOM 2023 [21].

E-mail addresses: umeyer@ae.cs.uni-frankfurt.de (Ulrich Meyer) htran@ae.cs.uni-frankfurt.de (Hung Tran) K.Tsakalidis@liverpool.ac.uk (Konstantinos Tsakalidis)



This work is licensed under the terms of the [CC-BY](https://creativecommons.org/licenses/by/4.0/) license.

28 Emerging big data applications need to process large graphs efficiently. Standard models of com-
 29 putation in internal memory (RAM, pointer machine) do not capture the algorithmic complexity
 30 of processing graphs with size that exceed the main memory. The *I/O-model* by Aggarwal and Vit-
 31 ter [1] is suitable for studying large graphs stored in an external memory hierarchy, e.g. comprised
 32 of cache, RAM and hard disk memories. The input data elements are stored in *external memory*
 33 (EM) packed in *blocks* of at most B elements and computation is free in *main memory* for at most
 34 M elements. The *I/O-complexity* is measured in *I/O-operations* (*I/Os*) that transfer a block from
 35 external to main memory and vice versa. Common tasks of many algorithms include reading or
 36 writing n contiguous items (which is referred to as scanning) requiring $\text{scan}(n) := \Theta(n/B)$ I/Os
 37 and sorting n consecutive elements requiring $\text{sort}(n) := \Theta((n/B) \log_{M/B}(n/B))$ I/Os.

38 1.1 Previous Work

39 Certifying bipartiteness in internal memory takes linear time in the number of edges by any traver-
 40 sal of the graph. In external memory, however, breadth-first search [20, 2] and depth-first search [5]
 41 algorithms take suboptimal $\omega(\text{sort}(n + m))$ I/Os for an input graph with n vertices and m edges.

42 Heggenes and Kratsch [14] present optimal internal memory algorithms for certifying whether
 43 a graph belongs to the classes of split, threshold, bipartite chain, and trivially perfect graphs. They
 44 return in linear time a YES-certificate characterizing the corresponding class or a forbidden induced
 45 subgraph of the class (NO-certificate). The YES- and NO-certificates are authenticated in linear and
 46 constant time, respectively. A straightforward application to the I/O-model leads to suboptimal
 47 certifying algorithms since graph traversal algorithms in external memory are much more involved
 48 and no worst-case efficient algorithms are known.

49 1.2 Our Results

50 We present I/O-efficient certifying algorithms for *split*, *threshold*, *bipartite chain*, and *trivially*
 51 *perfect* graphs. All algorithms return in the membership case, a YES-certificate w characterizing the
 52 graph class, or a $\mathcal{O}(1)$ -size NO-certificate in the non-membership case. The YES- and NO-certificates
 53 can be authenticated using $\mathcal{O}(\text{sort}(n + m))$ and $\mathcal{O}(1)$ I/Os, respectively. As a subroutine for the
 54 certification of bipartite chain graphs we develop a certifying algorithm to recognize bipartite
 55 graphs using $\mathcal{O}(\text{sort}(n + m))$ I/Os with high probability. Additionally, we perform experiments
 56 for split and threshold graphs showing scaling well beyond the size of main memory.

57 2 Preliminaries and Notation

58 For a graph $G = (V, E)$, let $n = |V|$ and $m = |E|$ denote the number of vertices V and edges
 59 E , respectively. We assume that the vertices $V = \{v_1, \dots, v_n\}$ are ordered by their indices. For
 60 a vertex $v \in V$ we denote by $N(v)$ the *neighborhood* of v and by $N[v] = N(v) \cup \{v\}$ the *closed*
 61 *neighborhood* of v . The *degree* $\text{deg}(v)$ of a vertex v is given by $\text{deg}(v) = |N(v)|$. A vertex v is
 62 called *simplicial* if $N(v)$ is a clique and *universal* if $N[v] = V$.

63 **Subgraphs and Orderings** The subgraph of G that is induced by a subset $A \subseteq V$ of vertices
 64 is denoted by $G[A]$. The *substructure* (subgraph) of a cycle on k vertices is denoted by C_k and
 65 of a path on k vertices is denoted by P_k . The $2K_2$ is a graph that is isomorphic to the following
 66 constant size graph: $(\{a, b, c, d\}, \{ab, cd\})$.

67 Henceforth we refer to different types of orderings of vertices: an ordering (u_1, \dots, u_n) is a (i)
 68 *perfect elimination ordering (peo)* if u_i is simplicial in $G[\{u_i, u_{i+1}, \dots, u_n\}]$ for all $i \in \{1, \dots, n\}$,
 69 and a (ii) *universal-in-a-component ordering (uco)* if u_i is universal in its connected component in
 70 $G[\{u_i, u_{i+1}, \dots, u_n\}]$ for all $i \in \{1, \dots, n\}$. For a subset $X = \{u_1, \dots, u_k\} \subseteq V$, we call (u_1, \dots, u_k)
 71 a *nested neighborhood ordering (nno)* if $(N(u_1) \setminus X) \subseteq (N(u_2) \setminus X) \subseteq \dots \subseteq (N(u_k) \setminus X)$.

72 Finally for any given ordering, we partition the set of neighbors $N(v)$ into $L(v) = \{x \in N(v) :$
 73 v is ranked higher than $x\}$ and $H(v) = \{x \in N(v) : v$ is ranked lower than $x\}$ where $L(v)$ and
 74 $H(v)$ denote the lower and higher ranked neighbors, respectively.

75 **Graph Relabeling** A *relabeling* of a graph $G = (V, E)$ is defined by a bijection $f : V \rightarrow V$
 76 where each edge $\{u, v\} \in E$ is reflected by an edge $\{f(u), f(v)\}$ of relabeled endpoints. For an
 77 ordering $\alpha = (u_1, \dots, u_n)$, a relabeling of G by α corresponds to the mapping where each v_i is
 78 mapped to its rank in α , e.g. $f(v_i) = v_r$ where r is the rank of v_i in α .

79 Employing this subroutine can lead to a more suitable representation of the graph in memory
 80 and often allows for more efficient data processing. The relabeling can be done I/O-efficiently
 81 in a constant number of scanning and sorting steps incurring $\mathcal{O}(\text{sort}(n+m))$ I/Os [4]. As all
 82 our algorithms perform an initial relabeling according to some ordering, we use the vertex labels
 83 obtained by this initial relabeling.

84 **Graph Representation** We assume an *adjacency array representation* [23] where the graph
 85 $G = (V, E)$ is represented by two arrays $P = [P_i]_{i=1}^n$ and $E = [u_i]_{i=1}^m$. The neighbors of a vertex
 86 v_i are then given in sorted order by the vertices at position P_i to $P_{i+1}-1$ in E . This representation
 87 allows for efficient straight-forward processing of G : (i) scanning k consecutive adjacency lists
 88 consisting of m' edges requires $\mathcal{O}(\text{scan}(m'))$ I/Os and (ii) computing and scanning the degrees of
 89 k consecutive vertices requires $\mathcal{O}(\text{scan}(k))$ I/Os.

90 **Time-Forward Processing** *Time-forward processing (TFP)* is a generic technique to manage
 91 data dependencies of external memory algorithms [18]. These dependencies are typically modeled
 92 by a directed acyclic graph $G = (V, E)$ where every vertex $v_i \in V$ models the computation of z_i
 93 and an edge $(v_i, v_j) \in E$ indicates that z_i is required for the computation of z_j .

94 Computing a solution then requires the algorithm to traverse G according to some topological
 95 order \prec_T of the vertices V . The TFP technique achieves this in the following way: after z_i has been
 96 calculated, the algorithm inserts a message $\langle v_j, z_i \rangle$ into a minimum priority-queue data structure
 97 for every successor $(v_i, v_j) \in E$ where the items are sorted by the recipients according to \prec_T . By
 98 construction, v_j receives all required values z_i of its predecessors $v_i \prec_T v_j$ as messages in the data
 99 structure. Since these predecessors already removed their messages from the data structure, items
 100 addressed to v_j are currently the smallest elements in the data structures and thus can be dequeued
 101 with a delete-minimum operation. By using suitable external memory priority-queues [3], TFP
 102 incurs $\mathcal{O}(\text{sort}(k))$ I/Os, where k is the number of messages sent.

3 Certifying Graph Classes in External Memory

3.1 Split Graphs

A split graph is a graph that can be partitioned into two sets of vertices (K, I) where K and I induce a clique and an independent set, respectively. The partition (K, I) is called the *split partition*. They are additionally characterized by the forbidden induced subgraphs $2K_2, C_4$ and C_5 , meaning that any vertex subset of a split graph cannot induce these substructures [13]. Since split graphs are a subclass of chordal graphs, there exists a perfect elimination ordering of the vertices for every split graph [11]. In fact, any non-decreasing degree ordering of a split graph is a perfect elimination ordering [14].

Our algorithm adapts the internal memory certifying algorithm of Heggernes and Kratsch [14] to external memory by adopting time-forward processing. As output it either returns the split partition (K, I) as a YES-certificate or one of the forbidden substructures C_4, C_5 or $2K_2$ as a NO-certificate. We present the certifying algorithm and its corresponding authentication algorithm and provide details in Proposition 1 and Proposition 2 and conclude with Theorem 1 at the end of the subsection.

Algorithm Description First, we compute a non-decreasing degree ordering $\alpha = (v_1, \dots, v_n)$ and relabel the graph according to α . Thereafter we check whether α is a perfect elimination ordering in $\mathcal{O}(\text{sort}(n+m))$ I/Os by Proposition 1. In the case that α is not a perfect elimination ordering, the algorithm returns three vertices v_j, v_k, v_i where $\{v_i, v_j\}, \{v_i, v_k\} \in E$ but $\{v_j, v_k\} \notin E$ and $i < j < k$, violating that v_i is simplicial in $G[\{v_i, \dots, v_n\}]$. In order to return a forbidden substructure we find additional vertices that complete the induced subgraphs. Note that (v_k, v_i, v_j) already forms a P_3 and may extend to a C_4 if $N(v_k) \cap N(v_j)$ contains a vertex $z \neq v_i$ that is not adjacent to v_i .

Computing $(N(v_k) \cap N(v_j)) \setminus N(v_i)$ requires scanning the adjacencies of three vertices totaling to $\mathcal{O}(\text{scan}(n))$ I/Os. If $(N(v_k) \cap N(v_j)) \setminus N(v_i)$ is empty we try to extend the P_3 to a C_5 or output a $2K_2$ otherwise. To do so, we find vertices $x \neq v_i$ and $y \neq v_i$ for which $\{x, v_j\}, \{y, v_k\} \in E$ but $\{x, v_k\}, \{y, v_j\} \notin E$ that are also not adjacent to v_i , i.e. $\{x, v_i\}, \{y, v_i\} \notin E$. Both x and y exist due to the ordering α [14] and are found using $\mathcal{O}(1)$ scanning steps requiring $\mathcal{O}(\text{scan}(n))$ I/Os. If $\{x, y\} \in E$ then (v_j, v_i, v_k, y, x) is a C_5 , otherwise $G[\{v_j, x, v_k, y\}]$ constitutes a $2K_2$. Determining whether $\{x, y\} \in E$ requires scanning $N(x)$ and $N(y)$ using $\mathcal{O}(\text{scan}(n))$ I/Os.

In the membership case, α is a perfect elimination ordering and the algorithm proceeds to verify first the clique K and then the independent set I of the split partition (K, I) . Note that for a split graph the maximum clique of size k must consist of the k -highest ranked vertices in α [14] where k can be computed using $\mathcal{O}(\text{sort}(n+m))$ I/Os by Proposition 2. Therefore, it suffices to verify for each of the k candidates v_i whether it is connected to $\{v_{i+1}, \dots, v_n\}$ since the graph is undirected. For a sorted sequence of edges relabeled by α , we check this property using $\mathcal{O}(\text{scan}(m))$ I/Os. If we find a vertex $v_i \in \{v_{n-k+1}, \dots, v_n\}$ where $\{v_i, v_j\} \notin E$ with $i < j$ then $G[\{v_i, \dots, v_n\}]$ already does not constitute a clique and we have to return a NO-certificate. Since the maximum clique has size k , there are k vertices with degree at least $k-1$. By these degree constraints there must exist an edge $\{v_i, x\} \in E$ where $x \in \{v_1, \dots, v_{i-1}\}$ [14]. Additionally, it holds that $\{x, v_j\} \notin E$ and there exists an edge $\{z, v_j\} \in E$ where $z \in \{v_1, \dots, v_{i-1}\}$ that cannot be connected to x , i.e. $\{x, z\} \notin E$ [14]. Thus, we first scan the adjacency lists of v_i and v_j to find x and z in $\mathcal{O}(\text{scan}(n))$ I/Os and return $G[\{v_i, v_j, x, z\}]$ as the $2K_2$ NO-certificate. Otherwise let $K = \{v_{n-k+1}, \dots, v_n\}$.

Lastly, the algorithm verifies whether the remaining vertices form an independent set. We verify

Algorithm 1: Recognizing Perfect Elimination in External Memory

Data: edges E of graph G , non-decreasing degree ordering $\alpha = (v_1, \dots, v_n)$
Output: bool whether α is a peo, three invalidating vertices $\{v_i, v_j, v_k\}$ if not a peo

```

1 Relabel  $G$  according to  $\alpha$ 
2 for  $i = 1, \dots, n$  do
3   Retrieve  $H(v_i)$  from  $E$ 
4   if  $H(v_i) \neq \emptyset$  then
5     Let  $u$  be the smallest successor of  $v_i$  in  $H(v_i)$ 
6     for  $x \in H(v_i) \setminus \{u\}$  do
7       PQ.push( $\langle u, x, v_i \rangle$ ) // inform  $u$  of  $x$  coming from  $v_i$ 
8     while  $\langle v, v_k, v_j \rangle \leftarrow$  PQ.top() where  $v = v_i$  do // for each message to  $v_i$ 
9       if  $v_k \notin H(v_i)$  then //  $v_i$  does not fulfill peo property
10        return FALSE,  $\{v_i, v_j, v_k\}$ 
11      PQ.pop()
12 return TRUE
  
```

147 that each candidate v_i is not connected to $\{v_{i+1}, \dots, v_{n-k}\}$, since the graph is undirected. For
 148 this, it suffices to scan over $n - k$ consecutive adjacency lists in $\mathcal{O}(\text{scan}(m))$ I/Os. More precisely,
 149 we scan the adjacency lists from v_{n-k} to v_1 and in case an edge $\{v_i, v_j\}$ where $i < j \leq n - k$ is
 150 found we find two more vertices to again complete a $2K_2$. For the first occurrence of such a vertex
 151 v_i , we remark that $\{v_{i+1}, \dots, v_{n-k}\}$ and $\{v_{n-k+1}, \dots, v_n\}$ form an independent set and a clique,
 152 respectively. Therefore there exists a vertex $y \in K$ that is adjacent to x but not to v_i [14]. We
 153 find y by scanning $N(x)$ and $N(v_i)$ in $\mathcal{O}(\text{scan}(n))$ I/Os. To complete the $2K_2$ we similarly find
 154 $z \in N(y) \setminus (N(x) \cup N(v_i))$ in $\mathcal{O}(\text{scan}(n))$ I/Os which is guaranteed to exist [14].

155 **Authentication** Given G and a split partition (K, I) we can verify in $\mathcal{O}(\text{sort}(n + m))$ I/Os that
 156 G is indeed a split partition. After relabeling G by a non-decreasing degree ordering α , we verify
 157 that the relabeled vertices of K correspond to the k -highest ranked vertices in α . By a subsequent
 158 scan over the relabeled edges we check whether any edge runs between vertices of I and that the
 159 last k vertices form a clique.

160 For a graph G and any of the forbidden substructures $2K_2, C_4$ or C_5 we not only return the
 161 corresponding vertex subsets but also the edge positions in the adjacency array representation
 162 for both edges and non-edges. To do so, we revert the relabeling in $\mathcal{O}(\text{sort}(n + m))$ I/Os and
 163 access all corresponding adjacency lists in $\mathcal{O}(\text{scan}(n))$ I/Os and return appropriate pointers to the
 164 adjacency array representation. For edges that are present in the substructure we directly point to
 165 the corresponding entry. Conversely, as non-edges are not present, we instead return pointers to
 166 the position the edge would have occupied if it existed using the fact that the individual adjacency
 167 lists are sorted. Since all NO-certificates are of constant size, authentication therefore only requires
 168 $\mathcal{O}(1)$ I/Os by direct accesses to memory.

169 **Proposition 1** Verifying that a non-decreasing degree ordering $\alpha = (v_1, \dots, v_n)$ of a graph G is
 170 a perfect elimination ordering takes $\mathcal{O}(\text{sort}(n + m))$ I/Os.

171 **Proof:** We follow the approach of [12, Theorem 4.5] and adapt it to the external memory using
 172 time-forward processing, see Algorithm 1.

Algorithm 2: Maximum Clique Size for Chordal Graphs in External Memory

Data: edges E of input graph G , perfect elimination ordering $\alpha = (v_1, \dots, v_n)$ **Output:** maximum clique size χ

```

1 Relabel  $G$  according to  $\alpha$ 
2  $\chi \leftarrow 0$ 
3 for  $i = 1, \dots, n$  do
4   Retrieve  $H(v_i)$  from  $E$  // scan  $E$ 
5   if  $H(v_i) \neq \emptyset$  then
6     Let  $u$  be the smallest successor of  $v_i$  in  $H(v_i)$ 
7     PQ.push( $\langle u, |H(v_i)| - 1 \rangle$ ) //  $v_i$  simplicial  $\Rightarrow G[N(v_i)]$  is clique
8      $S(v_i) \leftarrow -\infty$ 
9     while  $\langle v, S \rangle \leftarrow$  PQ.top() where  $v = v_i$  do
10     $S(v_i) \leftarrow \max\{S(v_i), S\}$  // compute maximum over all
11    PQ.pop()
12   $\chi \leftarrow \max\{\chi, 1 + S(v_i)\}$ 
13 return  $\chi$ 

```

173 After relabeling and sorting the edges by α , we iterate over the vertices in the order given by α .
174 For a vertex v_i the set of neighbors $N(v_i)$ needs to be a clique in order for v_i to be simplicial. In
175 order to verify this for all vertices, we iterate over α and at vertex v_i retrieve $H(v_i)$ by a continuous
176 scan over E . Then let $u \in H(v_i)$ be the smallest higher ranked neighbor. As $u \in H(v_i) \subseteq N(v_i)$ is
177 adjacent to v_i , it has to be verified that it also is adjacent to the remaining neighbors. We verify
178 this property partially for higher ranked neighbors in time-forward fashion. To do so, we insert a
179 message $\langle u, w \rangle$ into a priority-queue where $w \in H(v_i) \setminus \{u\}$ to inform u of every vertex it should be
180 adjacent to. For any given v_i it is therefore verified that $N(v_i)$ is a clique after the processing of all
181 neighbors has finished. Conversely, after sending all required adjacency information, we retrieve
182 for v_i all messages $\langle v_i, - \rangle$ directed to v_i and check that all received vertices are indeed neighbors
183 of v_i by comparison to the existing adjacencies as seen by the scan over E .

184 Relabeling and sorting the edges takes $\mathcal{O}(\text{sort}(m))$ I/Os. Every vertex v_i inserts at most all
185 its higher ranked neighbors into the priority-queue totaling up to $\mathcal{O}(m)$ messages which takes
186 $\mathcal{O}(\text{sort}(m))$ I/Os. Checking that all received vertices are indeed neighbors only requires a concur-
187 rent scan over all edges since vertices are handled in ascending order by α . \square

188 **Proposition 2** *Computing the size of a maximum clique in a split graph takes $\mathcal{O}(\text{sort}(n + m))$ I/Os.*

189 **Proof:** Note that split graphs are both chordal and co-chordal [13]. For chordal graphs, computing
190 the size of a maximum clique in internal memory takes linear time [12, Theorem 4.17] and can be
191 adapted straight-forwardly to an external memory algorithm using $\mathcal{O}(\text{sort}(m))$ I/Os.

192 To do so, we simulate the data accesses of the internal memory variant using priority-queues to
193 employ time-forward processing, see Algorithm 2. The algorithm proceeds similar to Algorithm 1
194 but relays different information forward in time. For a vertex v_i we instead inform the smallest
195 successor $u \in H(v_i)$ of the fact it is in a clique of size $|H(v_i)|$, namely the higher ranked neighbors
196 of v_i . Conversely, at each vertex v_i , we collect all sent messages and compute the size of the
197 maximum clique that v_i is a part of and update the global maximum accordingly. \square

198 By the above description and Proposition 1 and Proposition 2 it follows that split graphs can
 199 be recognized using $\mathcal{O}(\text{sort}(n + m))$ I/Os which we summarize in Theorem 1.

200 **Theorem 1** *A graph G can be recognized whether it is a split graph or not in $\mathcal{O}(\text{sort}(n + m))$
 201 I/Os. In the membership case the algorithm returns the split partition (K, I) as the YES-certificate,
 202 and otherwise it returns an $\mathcal{O}(1)$ -size NO-certificate.*

203 3.2 Threshold Graphs

204 Threshold graphs [8, 12, 17] are split graphs with the additional property that the independent
 205 set I of the split partition (K, I) has a nested neighborhood ordering. Its corresponding forbidden
 206 substructures are $2K_2, P_4$ and C_4 . Alternatively, threshold graphs can be characterized by a
 207 graph generation process: repeatedly add universal or isolated vertices to an initially empty graph.
 208 Conversely, by repeatedly removing universal and isolated vertices from a threshold graph the
 209 resulting graph must be the empty graph. In comparison to certifying split graphs, threshold
 210 graphs thus require additional steps.

211 Our algorithm adapts the internal memory certifying algorithm of Heggenes and Kratsch [14]
 212 to external memory. As output it either returns a nested neighborhood ordering β of I as a YES-
 213 certificate or one of the forbidden induced subgraphs C_4, P_4 or $2K_2$ as a NO-certificate. We again
 214 present the certifying algorithm and its corresponding authentication algorithm and provide details
 215 in Proposition 3 and conclude with Theorem 2 at the end of the subsection.

216 **Algorithm Description** First, the algorithm certifies whether the input is a split graph. In
 217 the non-membership case, if the returned NO-certificate is a C_5 we extract a P_4 otherwise we
 218 return the substructure immediately. For the membership case, we recognize whether the input
 219 is a threshold graph by repeatedly removing universal and isolated vertices using the previously
 220 computed perfect elimination ordering α in $\mathcal{O}(\text{sort}(n + m))$ I/Os by Proposition 3 (see below).
 221 If the remaining graph is empty, we return the independent set I with its non-decreasing degree
 222 ordering. Note that after removing a universal vertex v_i , vertices with degree one become isolated.
 223 Since low-degree vertices are at the front of α , an I/O-efficient algorithm cannot determine and
 224 remove them on-the-fly after removing a high-degree vertex. Therefore preprocessing is required.

225 For every vertex v_i we compute the number of vertices $S(v_i)$ that become isolated after the
 226 removal of $\{v_i, \dots, v_n\}$. To do so, we iterate over α in ascending order and consider vertices
 227 v_i where $L(v_i) = \emptyset$. Since v_i has no lower ranked neighbors, it would become isolated after
 228 removing all vertices in $H(v_i)$, in particular this happens when the last successor with smallest
 229 index $v_j \in H(v_i)$ is removed. To capture this information, we save v_j in a vector \mathcal{S} and sort \mathcal{S}
 230 in non-ascending order incurring $\mathcal{O}(\text{sort}(m))$ I/Os. The number of consecutive occurrences of any
 231 vertex v_j in \mathcal{S} correspond to the number of isolated vertices that are created by the removal of
 232 the vertices $\{v_j, \dots, v_n\}$. Thus, the aforementioned values $S(v_n), \dots, S(v_1)$ are now accessible by
 233 a scan over \mathcal{S} after counting the occurrences of each v_j in $\mathcal{O}(\text{scan}(m))$ I/Os.

234 The algorithm now proceeds to check whether removing universal and isolated vertices leads
 235 to an empty graph. By iterating in reverse order of α , vertices are considered in non-increasing
 236 degree order and verified to be universal using the values that are computed in the preprocessing
 237 stage without the need to actually remove them. This incurs a total of $\mathcal{O}(\text{scan}(n))$ I/Os. In
 238 the membership case, the resulting graph would be empty and we return a non-decreasing degree
 239 ordering β on the vertices of the independent set I . In the non-membership case, there must exist
 240 a P_4 since the input is a split graph and can therefore not contain a C_4 or a $2K_2$.

Algorithm 3: Recognizing Threshold Graphs for Split Graphs in External Memory

Data: edges E of split graph G , peo $\alpha = (v_1, \dots, v_n)$
Output: bool whether G is threshold

```

1 Relabel  $G$  according to  $\alpha$ 
2 Vector  $\mathcal{S}$ 
3 for  $i = 1, \dots, n$  do
4   if  $L(v_i) = \emptyset$  then
5     Let  $v_j$  be the smallest successor of  $v_i$  in  $H(v_i)$ 
6      $\mathcal{S}.push(v_j)$  //  $v_i$  would be isolated after deleting  $\{v_j, \dots, v_n\}$ 
7 Sort  $\mathcal{S}$  in non-ascending order
8  $n_{del} \leftarrow 0$  // number of deleted universal/isolated vertices
9 for  $i = n, \dots, 1$  do
10  if  $L(v_i) \neq \emptyset$  then //  $v_i$  not isolated in  $G[\{v_1, \dots, v_n\}]$ 
11    if  $|L(v_i)| < (n - 1) - n_{del}$  then //  $v_i$  not universal
12      return FALSE
13     $n_{del} \leftarrow n_{del} + 1 + \text{occurrences of } v_i$  //  $v_i$  removed, scan  $\mathcal{S}$ 
14 return TRUE

```

241 To find a P_4 , we can disregard further vertices from the remaining graph that cannot be part of
242 a P_4 . For this, let $K' \subset K$ and $I' \subset I$ be the remaining vertices when the non-universal vertex is
243 discovered. Any $v \in K$ where $N(v) \cap I' = \emptyset$ and any $v \in I$ where $N(v) \cap K' = K'$ cannot be part of
244 a P_4 and can therefore be disregarded [14]. We proceed by considering and removing vertices of K
245 by non-descending degree and vertices of I by non-ascending degree. After this process, we retrieve
246 the highest-degree vertex v in I for which there exists $\{v, y\} \notin E$ and $\{y, z\} \in E$ where $y \in K$ and
247 $z \in I$ [14]. Additionally, there is a neighbor $w \in K$ of v for which $\{w, z\} \notin E$ [14] and we return
248 the P_4 given by $G[\{v, w, y, z\}]$. Finding the P_4 therefore only requires $\mathcal{O}(\text{scan}(n + m))$ I/Os.

249 **Authentication** Given G and a nested neighborhood ordering β , we authenticate that the im-
250 plicitly given split partition (K, I) certifies that G is a split graph using $\mathcal{O}(\text{sort}(n + m))$ I/Os, as
251 detailed in subsection 3.1. It remains to verify that $\beta = (v_1, \dots, v_{|I|})$ is indeed a nested neighbor-
252 hood ordering of I . To do so, we verify for increasing i that $N(v_i) \subseteq N(v_{i+1})$ by a concurrent scan
253 over both neighborhoods requiring a total of $\mathcal{O}(\text{scan}(m))$ I/Os for all i .

254 Since the NO-certificates are again of constant size, authenticating in the non-membership case
255 takes $\mathcal{O}(1)$ I/Os, as detailed in subsection 3.1.

256 **Proposition 3** Verifying that G emits an empty graph after repeatedly removing universal and
257 isolated vertices requires $\mathcal{O}(\text{sort}(n + m))$ I/Os.

258 **Proof:** The described algorithm can be seen in Algorithm 3. Relabeling of G by any non-decreasing
259 degree ordering takes $\mathcal{O}(\text{sort}(n + m))$ I/Os. Generating the values $S(v_n), \dots, S(v_1)$ requires a scan
260 over all adjacency lists in ascending order and sorting \mathcal{S} which takes $\mathcal{O}(\text{scan}(m) + \text{sort}(n))$ I/Os.
261 After preprocessing, the algorithm only requires a reverse scan over the vertices v_n, \dots, v_1 . While
262 iterating over α we check for each v_i whether $L(v_i) = \emptyset$. If v_i is not isolated it must be universal.
263 Therefore we compare its current degree $\text{deg}(v_i)$ with the value $(n - 1) - n_{del}$ where $n_{del} =$
264 $\sum_{j=i+1}^n S(v_j)$. All operations take $\mathcal{O}(\text{scan}(m))$ I/Os in total. \square

265 By the above description and Proposition 3 it follows that there exists a certifying algorithm for
 266 the recognition of threshold graphs using $\mathcal{O}(\text{sort}(n+m))$ I/Os which is summarized in Theorem 2.

267 **Theorem 2** *A graph G can be recognized whether it is a threshold graph or not in $\mathcal{O}(\text{sort}(n+m))$ I/Os.*
 268 *In the membership case the algorithm returns a nested neighborhood ordering β as the YES-certificate,*
 269 *and otherwise it returns an $\mathcal{O}(1)$ -size NO-certificate.*

270 3.3 Trivially Perfect Graphs

271 Trivially perfect graphs have no vertex subset that induces a P_4 or a C_4 [12]. In contrast to split
 272 graphs, any non-increasing degree ordering of a trivially perfect graph is a universal-in-a-component
 273 ordering [14]. In fact, this is a one-to-one correspondence: a non-increasing sorted degree sequence
 274 of a graph is a universal-in-a-component ordering if and only if the graph is trivially perfect [14].

275 In external memory this can be verified using time-forward processing by adapting the algorithm
 276 in [14]. As output it either returns a universal-in-a-component ordering γ as a YES-certificate or
 277 one of the two forbidden subgraphs C_4, P_4 as a NO-certificate. We again present the certifying
 278 algorithm and its corresponding authentication algorithm and provide details in Proposition 4 and
 279 conclude with Theorem 3 at the end of the subsection.

280 **Algorithm Description** After computing a non-increasing degree ordering γ the algorithm
 281 relabels the edges of the graph according to γ and sorts them. Now we iterate over the vertices in
 282 ascending order of γ , process for each vertex v_i its received messages and relay further messages
 283 forward in time. Initially all vertices are labeled with 0. Then, at step i vertex v_i checks that all
 284 adjacent vertices $N(v_i)$ have the same label as v_i . After this, v_i relabels each vertex $u \in N(v_i)$
 285 with its own index i and is then removed from the graph.

286 In the external memory setting we cannot access labels of vertices and relabel them on-the-fly
 287 but rather postpone the comparison of the labels to the adjacent vertices instead. To do so, v_i
 288 forwards its own label $\ell(v_i)$ to $u \in H(v_i)$ by sending two messages $\langle v_i, v_i, \ell(v_i) \rangle$ and $\langle u, v_i, i \rangle$ to u ,
 289 signaling that u should compare its own label to v_i 's label $\ell(v_i)$ and then update it to i . Since the
 290 label of any adjacent vertex is changed after processing a vertex, when arriving at vertex v_j an
 291 odd number of messages will be targeted to v_j , where the last one corresponds to its actual label
 292 at step j . Then, after collecting all received labels, we compare disjoint consecutive pairs of labels
 293 and check whether they match. In the membership case, we do not find any mismatch and return
 294 γ as the YES-certificate. Otherwise, we have to return a P_4 or C_4 .

295 In the description of [14] the authors stop at the first anomaly where v_i detects a mismatch
 296 in its own label and one of its neighbors. We simulate the same behavior by writing out every
 297 anomaly we find, e.g. that v_j does not have the expected label of v_i via an entry $\langle v_i, v_j, k \rangle$ where
 298 k denotes the label of v_j . After sorting the entries, we find the earliest anomaly $\langle v_i, v_j, k \rangle$ with the
 299 largest label k of v_i 's neighbors in $\mathcal{O}(\text{sort}(m))$ I/Os. Since v_j received the label k from v_k , but v_i
 300 did not, it is clear that v_k is not universal in its connected component in $G[\{v_k, v_{k+1}, \dots, v_n\}]$ and
 301 we thus find a P_4 or C_4 . Note that (v_k, v_j, v_i) already constitutes a P_3 where $\deg(v_k) \geq \deg(v_j)$,
 302 since v_j received the label k . Since v_j is adjacent to both v_k and v_i and $\deg(v_k) \geq \deg(v_j)$, there
 303 must exist a vertex $x \in N(v_k)$ where $\{v_j, x\} \notin E$. Thus, $G[\{v_k, v_j, v_i, x\}]$ is a P_4 if $\{v_i, x\} \notin E$
 304 and a C_4 otherwise. Finding x and determining whether the forbidden subgraph is a P_4 or a C_4
 305 requires scanning $\mathcal{O}(1)$ adjacency lists using $\mathcal{O}(\text{scan}(n))$ I/Os.

Algorithm 4: Recognizing Universal-in-a-Component Orderings in External Memory

Data: edges E of graph G , non-increasing degree ordering $\gamma = (v_1, \dots, v_n)$
Output: bool whether γ is a uco

```

1 Relabel  $G$  according to  $\gamma$ 
2 for  $i = 1, \dots, n$  do
3   Vector  $\mathcal{L} = [0]$  // initialize with 0
4   while  $\langle v, v_j, \ell \rangle \leftarrow \text{PQ.top}()$  where  $v = v_i$  do //  $v_i$ 's received labels
5      $\mathcal{L}.push(\ell)$ 
6      $\text{PQ.pop}()$ 
7   for  $i = 1, \dots, \mathcal{L}.size/2$  do //  $\mathcal{L}.size$  is even
8     if  $\mathcal{L}[2i] \neq \mathcal{L}[2i+1]$  and  $\mathcal{L}.size > 1$  then // mismatch / anomaly
9       return FALSE
10   $\ell(v_i) \leftarrow \mathcal{L}[\mathcal{L}.size]$  // assign label of  $v_i$ 
11  Retrieve  $H(v_i)$  from  $E$  // scan  $E$ 
12  for  $u \in H(v_i)$  do
13     $\text{PQ.push}(\langle u, v_i, \ell(v_i) \rangle)$ 
14     $\text{PQ.push}(\langle u, v_i, i \rangle)$ 
15 return TRUE

```

306 **Authentication** Given G and a universal-in-a-component ordering we run [Algorithm 4](#) using
307 $\mathcal{O}(\text{sort}(n + m))$ I/Os by [Proposition 4](#). In the case of non-membership, we find the given substructure
308 in $\mathcal{O}(1)$ I/Os, as detailed in [subsection 3.1](#).

309 **Proposition 4** *Verifying that a non-increasing degree ordering $\gamma = (v_1, \dots, v_n)$ of a graph G with*
310 *n vertices and m edges is a universal-in-a-component ordering requires $\mathcal{O}(\text{sort}(m))$ I/Os.*

311 **Proof:** Every vertex v_i receives exactly two messages per neighbor in $L(v_i)$ and verifies that all
312 consecutive pairs of labels match. Then, either the label i is sent to each higher ranked neighbor of
313 $H(v_i)$ via time-forward processing or it is verified that γ is not a universal-in-a-component ordering.
314 Since at most $\mathcal{O}(m)$ messages are forwarded, the resulting overall complexity is $\mathcal{O}(\text{sort}(m))$ I/Os.
315 Correctness follows from [\[14\]](#) since the adapted algorithm performs the same operations but only
316 delays the label comparisons. \square

317 By the above description and [Proposition 4](#) it follows that there exists a certifying algorithm
318 for the recognition of trivially perfect graphs using $\mathcal{O}(\text{sort}(n + m))$ I/Os which we summarize in
319 [Theorem 3](#).

320 **Theorem 3** *A graph G can be recognized whether it is a trivially perfect graph or not in $\mathcal{O}(\text{sort}(n + m))$*
321 *I/Os. In the membership case the algorithm returns the universal-in-a-component ordering γ as*
322 *the YES-certificate, and otherwise it returns an $\mathcal{O}(1)$ -size NO-certificate.*

323 3.4 Bipartite Chain Graphs

324 Bipartite chain graphs are bipartite graphs where one part of the bipartition has a nested neigh-
325 borhood ordering [\[24\]](#) similar to threshold graphs. Interestingly, for chain graphs one side of the

326 bipartition exhibits this property if and only if both partitions do [24]. Its forbidden induced sub-
 327 structures are $2K_2$, C_3 and C_5 . By definition, bipartite chain graphs are bipartite graphs which
 328 therefore requires I/O-efficient bipartiteness testing.

329 Our algorithm adapts the internal memory certifying algorithm of Heggenes and Kratsch [14]
 330 to external memory. As a byproduct, we develop a certifying algorithm to recognize whether an
 331 input graph is bipartite or not and use it as a subroutine, see Lemma 1. The algorithm either
 332 returns a bipartition $(U, V \setminus U)$ with two nested neighborhood orderings on U and $V \setminus U$ as a YES-
 333 certificate or one of the forbidden induced subgraphs C_3 , C_5 or $2K_2$ as a NO-certificate. We present
 334 the full certifying algorithm first and provide details in Lemma 1, Corollary 1 and conclude with
 335 Theorem 4 at the end of the subsection.

336 **Algorithm Description** We follow the linear time internal memory approach of [14] with slight
 337 adjustments to accommodate the external memory setting. First, we check whether the input is
 338 indeed a bipartite graph. Instead of using breadth-first search which is very costly in external
 339 memory, even for constrained settings [2], we can use a more efficient approach with spanning
 340 trees which is presented in Lemma 1. In case the input is not connected, we simply return two
 341 edges of two different components as the $2K_2$. If the graph is connected, we proceed to verify that
 342 the graph is bipartite and return a NO-certificate in the form of a C_3 , C_5 or $2K_2$ in case it is not.
 343 In order to find a C_3 , C_5 or $2K_2$ some modifications to Lemma 1 are necessary. Essentially, the
 344 algorithm instead returns a minimum odd cycle that is built from T and a single non-tree edge.
 345 Due to minimality we can then find a C_3 , C_5 or a $2K_2$. The result is summarized in Corollary 1.

346 Then, it remains to show that each side of the bipartition has a nested neighborhood ordering.
 347 Let U be the larger side of the partition. By [17] it suffices to show that the input is a bipartite
 348 chain graph if and only if the graph obtained by adding all possible edges with both endpoints in
 349 U is a threshold graph. Instead of materializing the threshold graph, we implicitly represent the
 350 new adjacencies of vertices in U to retain the same I/O-complexity and apply Theorem 2 using
 351 $\mathcal{O}(\text{sort}(n+m))$ I/Os. Note that, in this threshold graph vertices of U have higher degrees than
 352 vertices in $V \setminus U$ since U is the larger side of the bipartition. If the input is bipartite but not bipartite
 353 chain, we repeatedly delete vertices that are connected to all other vertices of the other side and
 354 the resulting isolated vertices, similar to subsection 3.3 and [14]. After this, the vertex v with
 355 highest degree has a non-neighbor y in the other partition. By similar arguments to subsection 3.2
 356 y is adjacent to another vertex z that is adjacent to a vertex x where $\{v, x\} \notin E$ [14]. As such,
 357 $G[\{v, y, z, x\}]$ is a $2K_2$ and can be found in $\mathcal{O}(\text{scan}(n))$ I/Os and returned as the NO-certificate.

358 **Authentication** Given G and a bipartition $(U, V \setminus U)$ with two nested neighborhood orderings on
 359 U and $V \setminus U$ we first confirm that U and $V \setminus U$ are indeed independent sets using $\mathcal{O}(\text{sort}(n+m))$ I/Os
 360 similar to subsection 3.1. After this, we confirm that the provided orderings are nested neighbor-
 361 hood orderings as detailed in subsection 3.2.

362 As the NO-certificates are of constant size, authentication again only takes $\mathcal{O}(1)$ I/Os in the
 363 non-membership case similar to subsection 3.1.

364 **Lemma 1** *A graph G can be recognized whether it is a bipartite graph or not in $\mathcal{O}(\text{sort}(n+m))$
 365 I/Os, given a spanning forest of the input graph. In the membership case the algorithm returns
 366 a bipartition $(U, V \setminus U)$ as the YES-certificate, and otherwise it returns an odd cycle as the NO-
 367 certificate.*

368 **Proof:** In case there are multiple connected components, we operate on each individually and
 369 thus assume that the input is connected. Let T be the edges of the spanning tree and $E \setminus T$ the

370 non-tree edges. Any edge $e \in E \setminus T$ may produce an odd cycle by its addition to T . In fact,
 371 the input is bipartite if and only if $T \cup \{e\}$ is bipartite for all $e \in E \setminus T$ ¹. We check whether
 372 an edge $e = \{u, v\}$ closes an odd cycle in T by computing the distance $d_T(u, v)$ of its endpoints
 373 in T . Since this is required for every non-tree edge $E \setminus T$, we resort to batch-processing. Note
 374 that T is a tree and hence after choosing a designated root $r \in V$ it holds that $d_T(u, v) =$
 375 $d_T(u, \text{LCA}_T(u, v)) + d_T(v, \text{LCA}_T(u, v))$ where $\text{LCA}_T(u, v)$ is the lowest common ancestor of u
 376 and v in T . Therefore for every edge $E \setminus T$ we compute its lowest common ancestor in T using
 377 $\mathcal{O}((1 + m/n) \cdot \text{sort}(n)) = \mathcal{O}(\text{sort}(m))$ I/Os [7].

378 Additionally, for each vertex $v \in V$ we compute its depth in T in $\mathcal{O}(\text{sort}(m))$ I/Os using Euler
 379 Tours [7] and inform each incident edge of this value by a few scanning and sorting steps. Similarly,
 380 each edge $e = \{u, v\}$ is provided of the depth of $\text{LCA}_T(u, v)$. Then, after a single scan over $E \setminus T$
 381 we compute $d_T(u, v)$ and check if it is even. If any value is even, we return the odd cycle as a
 382 NO-certificate or a bipartition in T as the YES-certificate. Both can be computed using Euler Tours
 383 in $\mathcal{O}(\text{sort}(m))$ I/Os. □

384 **Corollary 1** *If a connected graph G contains a C_3, C_5 or $2K_2$ then any of these subgraphs can be*
 385 *found in $\mathcal{O}(\text{sort}(n + m))$ I/Os given a spanning tree of G .*

386 **Proof:** We extend the algorithm presented in Lemma 1 to either return the induced cycles C_3
 387 and C_5 or a $2K_2$. While iterating over the edges to find an odd cycle we save the smallest one by
 388 keeping a copy of the edge $e \in E \setminus T$ and the length of the minimum odd cycle. In case we find a C_3
 389 or a C_5 we are done and return the NO-certificate immediately otherwise for an odd (non-induced)
 390 cycle of length k with $k = 2\ell + 1 > 5$ we return a $2K_2$ by finding a matching edge to the non-tree
 391 edge $e \in E \setminus T$ in the cycle.

392 Let $C = (u_1, \dots, u_k, u_1)$ be the returned cycle where $\{u_k, u_1\}$ is the non-tree edge. In this
 393 case we return for the $2K_2$ the graph $(\{u_\ell, u_{\ell+1}, u_1, u_k\}, \{\{u_1, u_k\}, \{u_\ell, u_{\ell+1}\}\})$. If ℓ is odd,
 394 the non-edges of the $2K_2$ cannot exist since otherwise any of the following smaller odd cycles
 395 $(u_1, u_2, \dots, u_{\ell+1}, u_k, u_1)$, $(u_1, u_2, \dots, u_\ell, u_1)$, $(u_\ell, u_{\ell+1}, \dots, u_k, u_\ell)$ and $(u_1, u_{\ell+1}, u_{\ell+2}, \dots, u_k, u_1)$
 396 would be present, contradicting the minimality of C . For the other case where ℓ is even, a similar
 397 argument can be found. The I/O-complexity therefore remains the same. □

398 We summarize our findings for bipartite chain graphs in Theorem 4.

399 **Theorem 4** *A graph G can be recognized whether it is a bipartite chain graph or not in $\mathcal{O}(\text{sort}(n + m))$*
 400 *I/Os with high probability. In the membership case the algorithm returns a bipartition $(U, V \setminus U)$*
 401 *and nested neighborhood orderings of both partitions as the YES-certificate, and otherwise it returns*
 402 *a $\mathcal{O}(1)$ -size NO-certificate.*

403 **Proof:** Computing a spanning tree T requires $\mathcal{O}(\text{sort}(n + m))$ I/Os with high probability by an
 404 external memory variant of the Karger, Klein and Tarjan minimum spanning tree algorithm [7].
 405 By Corollary 1 we find a C_3, C_5 or $2K_2$ if the input is not bipartite or not connected. We proceed
 406 by checking the nested neighborhood orderings of both partitions in $\mathcal{O}(\text{sort}(n + m))$ I/Os using
 407 Theorem 2. □

¹Since T is bipartite, one can think of T as a representation of a 2-coloring on T .

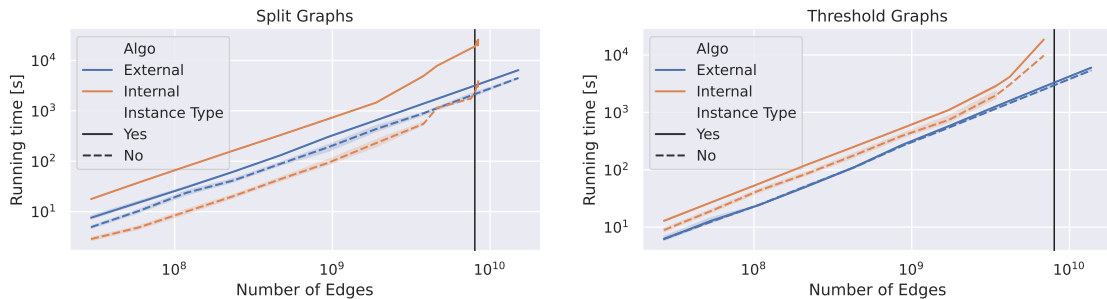


Figure 1: Running times of the certifying algorithms for split (left) and threshold graphs (right) for different random graph instances. The black vertical lines depict the number of elements that can concurrently be held in internal memory.

4 Experimental Evaluation

We implemented our external memory certifying algorithms for split and threshold graphs in C++ using the STXXL library [9]. To provide a comparison of our algorithms, we also implemented the internal memory state-of-the-art algorithms by Heggernes and Kratsch [14]. STXXL offers external memory versions of fundamental algorithmic building blocks like scanning, sorting and several data structures. Our benchmarks are built with GNU g++-10.3 and executed on a machine equipped with an AMD EPYC 7302P processor and 64 GB RAM running Ubuntu 20.04 using six 500 GB solid-state disks.

In order to validate the predicted scaling behavior we generate our instances parameterized by n . For YES-instances of split graphs we generate a split partition (K, I) with $|K| = n/10$ and add each possible edge $\{u, v\}$ with probability $1/4$ for $u \in I$ and $v \in K$. Analogously, YES-instances of threshold graphs are generated by repeatedly adding either isolated or universal vertices with probability $9/10$ and $1/10$, respectively. We additionally attempt to generate NO-instances by adding $\mathcal{O}(1)$ many random edges to the YES-instances. In a last step, we randomize the vertex indices to remove any biases emerging from the generation process.

In Figure 1 we present the running times of all algorithms on multiple YES- and NO-instances. It is clear that the performance of both external memory algorithms is not impacted by the main memory barrier while the running time of their internal memory counterparts already increases when at least half the main memory is used. This effect is amplified immensely after exceeding the size of main memory for split graphs.

Certifying the produced NO-instances of split graphs seems to require less time than their corresponding unmodified YES-instances as the algorithm typically stops early. Furthermore, due to the low data locality of the internal memory variant it is apparent that the external memory algorithm is superior for the YES-instances. The performance on both YES- and NO-instances is very similar in external memory. This is in part due to the fact that the common relabeling step is already relatively costly. For threshold graphs, however, the external memory variant outperforms the internal memory variant due to improved data locality.

5 Conclusions

We have presented the first I/O-efficient certifying recognition algorithms for split, threshold, trivially perfect, bipartite and bipartite chain graphs. Our algorithms require $\mathcal{O}(\text{sort}(n + m))$ I/Os matching common lower bounds for many algorithms in external memory. In our experiments we show that the algorithms perform well even for graphs exceeding the size of main memory.

Further, it would be interesting to extend the scope of certifying recognition algorithms to more graph classes for the external memory regime. In internal memory, a plethora of graph classes are efficiently certifiable which currently have no efficient external memory pendant, e.g. circular-arc graphs [10], HHD-free graphs [22], interval graphs [16], normal helly circular-arc graphs [6], permutation graphs [16], proper interval graphs [15], proper interval bigraphs [15] and many more. Due to limited data locality, straight-forward applications of these algorithms are highly inefficient for use in external memory. In turn, new algorithmic techniques are necessary to bridge the gap to larger processing scales.

References

- [1] A. Aggarwal and J. S. Vitter. The input/output complexity of sorting and related problems. *Communications of the ACM*, 31(9):1116–1127, 1988. doi:10.1145/48529.48535.
- [2] D. Ajwani and U. Meyer. Design and engineering of external memory traversal algorithms for general graphs. In J. Lerner, D. Wagner, and K. A. Zweig, editors, *Algorithmics of Large and Complex Networks - Design, Analysis, and Simulation [DFG priority program 1126]*, volume 5515 of *Lecture Notes in Computer Science*, pages 1–33. Springer, 2009. doi:10.1007/978-3-642-02094-0_1.
- [3] L. Arge. The buffer tree: A technique for designing batched external data structures. *Algorithmica*, 37(1):1–24, 2003. doi:10.1007/s00453-003-1021-x.
- [4] G. S. Brodal, R. Fagerberg, D. Hammer, U. Meyer, M. Penshuck, and H. Tran. An experimental study of external memory algorithms for connected components. In D. Coudert and E. Natale, editors, *19th International Symposium on Experimental Algorithms, SEA 2021, June 7-9, 2021, Nice, France*, volume 190 of *LIPICs*, pages 23:1–23:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.SEA.2021.23.
- [5] A. L. Buchsbaum, M. H. Goldwasser, S. Venkatasubramanian, and J. R. Westbrook. On external memory graph traversal. In D. B. Shmoys, editor, *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA*, pages 859–860. ACM/SIAM, 2000.
- [6] Y. Cao. Direct and certifying recognition of normal helly circular-arc graphs in linear time. In J. Chen, J. E. Hopcroft, and J. Wang, editors, *Frontiers in Algorithmics - 8th International Workshop, FAW 2014, Zhangjiajie, China, June 28-30, 2014. Proceedings*, volume 8497 of *Lecture Notes in Computer Science*, pages 13–24. Springer, 2014. doi:10.1007/978-3-319-08016-1_2.
- [7] Y. Chiang, M. T. Goodrich, E. F. Grove, R. Tamassia, D. E. Vengroff, and J. S. Vitter. External-memory graph algorithms. In K. L. Clarkson, editor, *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, 22-24 January 1995. San Francisco*,

- 475 *California, USA*, pages 139–149. ACM/SIAM, 1995. URL: <http://dl.acm.org/citation.cfm?id=313651.313681>.
- 477 [8] V. Chvátal and P. L. Hammer. Set-packing and threshold graphs. *Research Report, Computer*
478 *Science Department, Univeristy Waterloo, 1973*, 1973.
- 479 [9] R. Dementiev, L. Kettner, and P. Sanders. STXXL: standard template library for XXL data
480 sets. *Software - Practice and Experience*, 38(6):589–637, 2008. doi:10.1002/spe.844.
- 481 [10] M. C. Francis, P. Hell, and J. Stacho. Forbidden structure characterization of circular-arc
482 graphs and a certifying recognition algorithm. In P. Indyk, editor, *Proceedings of the Twenty-*
483 *Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA,*
484 *USA, January 4-6, 2015*, pages 1708–1727. SIAM, 2015. doi:10.1137/1.9781611973730.
485 114.
- 486 [11] D. Fulkerson and O. Gross. Incidence matrices and interval graphs. *Pacific Journal of Math-*
487 *ematics*, 15(3):835–855, 1965.
- 488 [12] M. C. Golumbic. *Algorithmic graph theory and perfect graphs*. Elsevier, 2004.
- 489 [13] P. L. Hammer and S. Földes. Split graphs. *Congressus Numerantium*, 19:311–315, 1977.
- 490 [14] P. Heggernes and D. Kratsch. Linear-time certifying recognition algorithms and forbidden
491 induced subgraphs. *Nordic Journal of Computing*, 14(1-2):87–108, 2007.
- 492 [15] P. Hell and J. Huang. Certifying lexbfs recognition algorithms for proper interval graphs
493 and proper interval bigraphs. *SIAM Journal on Discrete Mathematics*, 18(3):554–570, 2004.
494 doi:10.1137/S0895480103430259.
- 495 [16] D. Kratsch, R. M. McConnell, K. Mehlhorn, and J. P. Spinrad. Certifying algorithms for
496 recognizing interval graphs and permutation graphs. *SIAM Journal on Computing*, 36(2):326–
497 353, 2006. doi:10.1137/S0097539703437855.
- 498 [17] N. V. Mahadev and U. N. Peled. Threshold graphs and related topics. *Annals of Discrete*
499 *Mathematics*, 56, 1995.
- 500 [18] A. Maheshwari and N. Zeh. A survey of techniques for designing i/o-efficient algorithms. In
501 U. Meyer, P. Sanders, and J. F. Sibeyn, editors, *Algorithms for Memory Hierarchies, Advanced*
502 *Lectures [Dagstuhl Research Seminar, March 10-14, 2002]*, volume 2625 of *Lecture Notes in*
503 *Computer Science*, pages 36–61. Springer, 2002. doi:10.1007/3-540-36574-5_3.
- 504 [19] R. M. McConnell, K. Mehlhorn, S. Näher, and P. Schweitzer. Certifying algorithms. *Computer*
505 *Science Review*, 5(2):119–161, 2011. doi:10.1016/j.cosrev.2010.09.009.
- 506 [20] K. Mehlhorn and U. Meyer. External-memory breadth-first search with sublinear I/O. In R. H.
507 Möhring and R. Raman, editors, *Algorithms - ESA 2002, 10th Annual European Symposium,*
508 *Rome, Italy, September 17-21, 2002, Proceedings*, volume 2461 of *Lecture Notes in Computer*
509 *Science*, pages 723–735. Springer, 2002. doi:10.1007/3-540-45749-6_63.
- 510 [21] U. Meyer, H. Tran, and K. Tsakalidis. Certifying induced subgraphs in large graphs. In
511 C. Lin, B. M. T. Lin, and G. Liotta, editors, *WALCOM: Algorithms and Computation -*
512 *17th International Conference and Workshops, WALCOM 2023, Hsinchu, Taiwan, March*
513 *22-24, 2023, Proceedings*, volume 13973 of *Lecture Notes in Computer Science*, pages 229–
514 241. Springer, 2023. doi:10.1007/978-3-031-27051-2_20.

- 515 [22] S. D. Nikolopoulos and L. Palios. An $o(nm)$ -time certifying algorithm for recognizing HHD-
516 free graphs. *Theoretical Computer Science*, 452:117–131, 2012. doi:10.1016/j.tcs.2012.
517 06.006.
- 518 [23] P. Sanders, K. Mehlhorn, M. Dietzfelbinger, and R. Dementiev. *Sequential and Paral-*
519 *lel Algorithms and Data Structures - The Basic Toolbox*. Springer, 2019. doi:10.1007/
520 978-3-030-25209-0.
- 521 [24] M. Yannakakis. Node-deletion problems on bipartite graphs. *SIAM Journal on Computing*,
522 10(2):310–327, 1981. doi:10.1137/0210022.