Efficient Enumeration of Markov Equivalent DAGs*

Marcel Wienöbst,[1] **Malte Luttermann,**[2] **Max Bannach,**[1] **Maciej Liśkiewicz**[1]

[1] Institute for Theoretical Computer Science, University of Lübeck, Germany
[2] Institute of Information Systems, University of Lübeck, Germany
{wienoebst@tcs, luttermann@ifis, bannach@tcs, liskiewi@tcs}.uni-luebeck.de

Abstract

Enumerating the directed acyclic graphs (DAGs) of a Markov equivalence class (MEC) is an important primitive in causal analysis. The central resource from the perspective of computational complexity is the delay, that is, the time an algorithm that lists all members of the class requires between two consecutive outputs. Commonly used algorithms for this task utilize the rules proposed by Meek (1995) or the transformational characterization by Chickering (1995), both resulting in super-linear delay. In this paper, we present the first linear-time delay algorithm. On the theoretical side, we show that our algorithm can be generalized to enumerate DAGs represented by models that incorporate background knowledge, such as MPDAGs; on the practical side, we provide an efficient implementation and evaluate it in a series of experiments. Complementary to the linear-time delay algorithm, we also provide intriguing insights into Markov equivalence itself: All members of an MEC can be enumerated such that two successive DAGs have structural Hamming distance at most three.

1 Introduction

Graphical causal models endow researchers with an intuitive and mathematically sound language to infer causal relations between random variables from observational and interventional data. Directed acyclic graphs (DAGs), whose edges encode direct causal influences between the variables, belong to the most popular models and are used in many areas of empirical research (Spirtes, Glymour, and Scheines 2000; Rothman et al. 2008; Pearl 2009; Koller and Friedman 2009; Elwert 2013). However, there is usually not a unique DAG that can be learned from observational or limited experimental data as multiple models can encode the same statistical properties. These DAGs form a *Markov equivalence class* (MEC) and each of them explains the data equally well (Andersson, Madigan, and Perlman 1997; Pearl 2009).

Exploring the structural and quantitative properties of MECs are challenging tasks in graphical causal analysis and, despite extensive research efforts, several basic issues involving MECs fundamental to causal discovery remain open as e. g., calculating the number of MECs on n variables (Gillispie and Lemieux 2001; Steinsky 2003) or enumerating them efficiently (Chen, Choi, and Darwiche 2016).

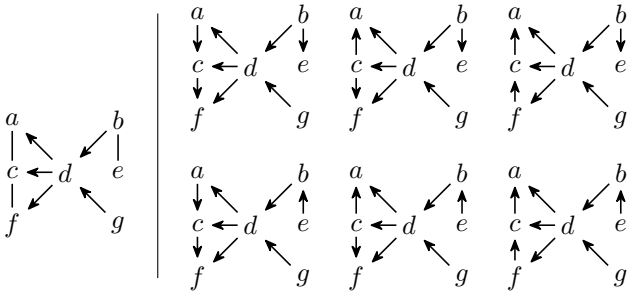

Figure 1: A Markov equivalence class (MEC) on the right, which consists of six DAGs. This class is represented by the left CPDAG, which uniquely represents the MEC by including undirected edges if two DAGs differ in their direction.

In this work, we study the properties of a *single* MEC, encoded as a *completed partially directed acyclic graph* (CPDAG) (Andersson, Madigan, and Perlman 1997), see Fig. 1 for an example. The CPDAG representation is often learned directly by causal discovery algorithms (Spirtes, Glymour, and Scheines 2000; Chickering 2002) and, thus, it is of high practical value to offer efficient implementations for their analysis. Recently, Wienöbst, Bannach, and Liśkiewicz (2021b) have shown that computing the size of an MEC as well as uniformly sampling from it can be done in polynomial time. We deal with the closely related problem of *enumerating the DAGs in an MEC* given its CPDAG, that is listing each member of the MEC exactly once. This task is an important primitive in causal analysis, used as a subroutine to solve more complex problems in software packages such as `pcalg` (Kalisch et al. 2012) and `causaldag` (Squires 2018). Enumeration of an MEC's members can be applied to solve many important downstream tasks in causal inference. For example, one can estimate the causal effect of the exposure variable on the outcome for each DAG in the equivalence class, which is learned from the observed data (Maathuis, Kalisch, and Bühlmann 2009). One could also check for every DAG whether it conforms to additional domain information or background knowledge in order to find the most plausible DAG (Meek 1995), or select intervention targets to distinguish between certain DAGs in the class (He and Geng 2008; Hauser and Bühlmann 2012). While there are

*Extended version of paper accepted to the Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-23).

custom algorithms, which avoid the use of explicit enumeration (sometimes by settling for approximate solutions), for many of these cases, it remains a flexible and very general tool that can be utilized even when other methods fail. The main drawback is, of course, its high computational cost, which we aim to address in this work.

Any method for enumerating the DAGs in an MEC requires exponential time in the worst-case, due to the basic fact that there are classes with exponential size. A crucial feature from a computational perspective is the *delay:* the algorithm's run-time between two consecutive output DAGs. Another desirable property would be that the subsequent DAGs smoothly change their structure, i. e., share most of their edge orientations, which constitutes a more plausible enumeration from the causal point of view. In the present work, we take both these aspects into account.

To the best of our knowledge, no study has been published that performs a systematic analysis of the enumeration problem, including its algorithmic aspects. One commonly used folklore algorithm utilizes the rules proposed by Meek (1995) to transform a causal graph (e. g., a CDPAG or PDAG) into its maximal orientation. Applying these rules has the property that any remaining undirected edge $a - b$ is oriented $a \to b$ in at least one and $a \leftarrow b$ in another DAG represented by the graph. Consequently, the DAGs can be enumerated by successively trying both possible orientations. This yields a polynomial delay algorithm, but the degree of the corresponding polynomial is rather large since the Meek rules have to be applied at *every* step. Another folklore approach is based upon the transformational characterization of MECs given by Chickering (1995), which states that two DAGs in the same MEC can be transformed into each other by successive single-edge reversals. Hence, the MEC can be explored through such edge reversals starting from an arbitrary DAG in the class. The issue with this approach is that already output DAGs need to be stored and every time a new DAG is explored it has to be checked that it has not been output before. This leads to a relatively large delay and memory demand. As both algorithms (which we call MEEK-ENUM and CHICKERING-ENUM) have not been explicitly stated in a publication, we give a formal description of both in Appendix A as well as a rigorous analysis of their delay.

The main contribution of this paper is the first $O(n + m)$-delay algorithm that, for a given CPDAG representing an MEC, lists all members of the class.[1] We also show that the algorithm can be generalized to enumerate DAGs represented by a PDAG or MPDAG – causal models incorporating background knowledge. To achieve these results, we utilize the Maximum Cardinality Search (MCS) (Tarjan and Yannakakis 1984) originating from the chordal graph literature. In addition to the theoretical results, we give an efficient practical implementation, which is significantly faster than implementations of MEEK-ENUM and CHICKERING-ENUM.

We also propose a complementary method with the property that during enumeration subsequent DAGs gradually change their structure. This method utilizes the results by (Chickering 1995), but performs the traversal of the MEC in a more refined way. Using such an approach, it is possible to output all Markov equivalent DAGs in sequence with the property that two successive DAGs have *structural Hamming distance* (SHD) at most three. This result is tight in the sense that there are MECs whose members cannot be enumerated in a sequence with maximal distance at most two. We also show that our ideas can be used in the more general setting of enumerating maximal ancestral graphs (MAGs) which encode conditional independence relations in DAG models with latent variables (Richardson and Spirtes 2002).

[1] We denote the number of vertices by n and the number of edges by m.

2 Preliminaries

A graph $G = (V, E)$ consists of a set of vertices V and a set of edges $E \subseteq V \times V$. An edge $u - v$ is undirected if $(u, v), (v, u) \in E$ and directed $u \to v$ if $(u, v) \in E$ and $(v, u) \notin E$. Vertices linked by an edge (of any type) are *adjacent* and *neighbors* of each other. We say that u is a *parent* of v if $u \to v$. We denote by $Pa(v)$ and $Ne(v)$ the set of parents and neighbors of v. The *degree* $\delta(v)$ of vertex v is the number of its neighbors $|Ne(v)|$. The structural Hamming distance (SHD) of $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ is denoted by $\mathrm{shd}(G_1, G_2)$ and defined as number of pairs $(a, b) \in V^2$ with differing edge relations, i. e., $E_1 \cap \{(a, b), (b, a)\} \neq E_2 \cap \{(a, b), (b, a)\}$. Given a graph $G = (V, E)$ and a vertex set S, the *induced subgraph* $G[S]$ contains the edges $E \cap (S \times S)$ from G that are incident only to vertices in S. The union of a set of graphs $\{G_1 = (V, E_1), \ldots, G_k = (V, E_k)\}$ is the graph $G = (V, \bigcup_{i=1}^{k} E_k)$. A path π between two vertices v_1 and v_p is a sequence of distinct vertices $\pi = \langle v_1, \ldots, v_p \rangle$ with $p \geq 2$ such that each vertex v_i is adjacent to v_{i+1} for $i = 1, \ldots, p - 1$. An *undirected connected component* is a maximal induced subgraph in which every pair of vertices is connected by a path of undirected edges. A path of the form $v_1 \to v_2 \to \ldots \to v_p$ is directed or *causal*. A graph is *acyclic* if there is no directed path from a vertex u to v with $v \to u$. An acyclic graph with only directed edges is called a DAG. An undirected graph is called *chordal* if no subset of four or more vertices induces an undirected cycle.

The *skeleton* of G, denoted by $\mathrm{skel}(G)$, is a graph with the same vertex set in which every edge is replaced by an undirected edge. A *v-structure* is an ordered triple of vertices (u, c, v) that induces the subgraph $u \to c \leftarrow v$. A Markov equivalence class (MEC) consists of DAGs encoding the same set of conditional independence relations among the variables. (Verma and Pearl 1990; Frydenberg 1990) showed that two DAGs are Markov equivalent iff they have the same skeleton and the same v-structures. An MEC can be represented by a CPDAG (*completed partially directed acyclic graph*), which is the union graph of the DAGs in the equivalence class it represents. The set $[G]$ denotes all DAGs in the MEC represented by CPDAG G.

Subclasses of MECs can be represented by *partially directed acyclic graphs* (PDAGs), which are restricted only in that they may not contain a directed cycle, and *maximally oriented PDAGs* (MPDAGs), which consist of PDAGs closed under the four *Meek rules*. Explicitly, we only use the first

Meek rule in this work, which states that an induced subgraph $a \to b - c$ is oriented into $a \to b \to c$. A formal description of all four Meek rules can be found in Appendix A.

The input of the enumeration algorithms is the representation of an MEC in form of its CPDAG G. Hence, before approaching the enumeration task, it is important, first and foremost, to understand how one can derive a DAG in $[G]$ from the CPDAG representation (this is also called the *extension task* and such a DAG is called a *consistent extension* of G). To extend G into a consistent DAG, it is necessary to find an orientation of its undirected edges. This orientation needs to be acyclic and contain the same v-structures as G.

Fact 1 (Andersson, Madigan, and Perlman (1997); He, Jia, and Yu (2015))**.** *The undirected components of a CPDAGs are* undirected and connected chordal graphs *(UCCGs). Orienting each UCCG independently, without introducing a v-structure, gives a DAG in* $[G]$.

The acyclic orientations without a v-structure of a chordal graph are called *AMOs* (acyclic moral orientations). *Every* DAG in $[G]$ may be computed by finding appropriate AMOs for the UCCGs. In Fig. 1, to obtain the DAG at the top left, the UCCG $a - c - f$ is oriented as AMO $a \to c \to f$ and $b - e$ is oriented as $b \to e$. For the former there are three AMOs, for the latter two, corresponding to the six DAGs in the MEC (as the UCCGs can be oriented independently). Thus, when tackling the enumeration task for CPDAGs, it suffices to enumerate the AMOs of a chordal graph.

Fact 2 (Implicit in Wienöbst, Bannach, and Liśkiewicz (2021b))**.** *Every AMO of a UCCG G can be obtained by orienting the edges according to an MCS ordering ($a \to b$ if a comes before b in the ordering).*

An MCS ordering is a linear ordering of the vertices produced by running the graph traversal algorithm *Maximum Cardinality Search* (MCS) (Tarjan and Yannakakis 1984), which was originally proposed for testing chordality of a graph in linear time. Notably, the reverse direction holds as well, that is, an MCS ordering will always produce an AMO of G. A brief introduction to chordal graphs, AMOs and the MCS algorithm is given in Appendix 12.

The output of an MCS depends on the choices of the "next" vertex to visit in each step of the graph traversal. This vertex is taken from the set of vertices with the largest number of already visited neighbors (called the vertices with *highest label*) and, from Fact 2, we can conclude that there are such choices, which may produce any AMO of a given UCCG.

3 Enumerating AMOs with Linear Delay

The observations from the previous section yield a new approach: Instead of producing a single AMO by choosing an arbitrary vertex in each step of the MCS, we perform multiple choices and recur for each of them – eventually listing all AMOs. There is one pitfall however: We *cannot* simply choose *every* vertex from the highest-label set one after the other as some graphs would be output multiple times. Fig. 2 illustrates this issue: If vertex a has been visited, the next vertex could be b, c, d, e, or f (all have one visited neighbor, namely a). But choosing, say, b or e may lead to the same AMO as the order of b and e after choosing a is irrelevant.

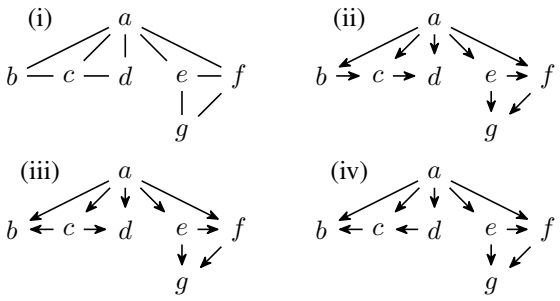

Figure 2: An example showing for the chordal graph (i) that, at a given step of the algorithm, not all vertices in the highest-label set can be chosen. If the MCS starts with vertex a, all neighbors b, c, d, e, f have the same label, namely 1. While an MCS may choose any one of them, we *cannot* choose *all* one-after-the-other in our enumeration. Choosing b or e as the second vertex may yield the same AMO as a, b, c, d, e, f, g and a, e, f, g, b, c, d are both topological orderings of (ii). However, choosing highest-label vertices from one connected component in $G[\{b, c, d, e, f, g\}]$ such as $\{b, c, d\}$ one-after-another will yield distinct AMOs (in (iii) and (iv) AMOs with c, d chosen as second vertex are given).

This issue can be addressed as follows: While the choice of the first vertex v from the highest-label set is arbitrary, every other vertex x has to be connected to v in the remaining graph (the induced subgraph over the unvisited vertices). If they are connected, then the order of v and x matters. Otherwise, choosing x instead of v would lead to duplicate AMOs being output. In our example, this means that if b is the first considered vertex with highest label after a has been visited, the other choices we would consider are c and d as these are the vertices reachable from b in the induced subgraph over the unvisited vertices.

Lemma 1. *Given a chordal graph $G = (V, E)$ and the sequence of previously visited vertices τ with $|\tau| = k < n$ produced by an MCS with current highest-label set S.*

1. *If $x, y \in S$ are connected in $G[V \setminus \tau]$, the set of AMOs produced by choosing x next is disjoint from the set produced by choosing y next.*
2. *If $x, y \in S$ are unconnected in $G[V \setminus \tau]$, any AMO produced by choosing y as the next vertex can be produced by choosing a vertex in S connected to x in $G[V \setminus \tau]$ next.*

Proof. We show item 1 first and let p be the shortest path between x and y in $G[V \setminus \tau]$. Since AMOs do not contain v-structures, any AMO choosing x before y must orient p as $x \to \cdots \to y$ while any other AMO yields $x \leftarrow \cdots \leftarrow y$.

For item 2 let $C_1, \ldots, C_k$ be the connected components of $G[V \setminus \tau]$ with $x \in C_1$. Any topological ordering with prefix τ can be rewritten as $\tau, C_{\pi(1)}, \ldots, C_{\pi(k)}$ for an arbitrary permutation π – in particular for $\pi = \mathrm{id}$. □

The following Lemma provides a simplified way of testing whether two vertices of highest label are connected.

input : A UCCG $G = (V, E)$.
output : All AMOs of G.

```
1  A := array of n initially empty sets
2  τ := empty list
3  A[0] := V
4  enumerate(G, A, τ)
5  function enumerate(G, A, τ)
6      if |τ| = n then
7          Output AMO of G according to ordering τ
8      end
9      i := highest index of non-empty set in A
10     v := any vertex from A[i]
11     x := v
12     do
13         delete x from A[i]
14         append x to τ
15         foreach w ∈ (Ne(x) \ τ) do
16             j := index of set in A that w is in
17             delete w from A[j]
18             insert w in A[j + 1]
19         end
20         enumerate(G, A, τ)
21         foreach w ∈ (Ne(x) \ τ) do
22             j := index of set in A that w is in
23             delete w from A[j]
24             insert w in A[j − 1]
25         end
26         insert x in A[i]
27         pop x from τ
28         if x = v then
29             R := {a | a reachable from v in G[A[i]]}
30         end
31     while R is non-empty, x := pop(R)
32 end
```

Algorithm 1: Linear-time delay enumeration algorithm MCS-ENUM for listing all AMOs of a chordal graph.

Lemma 2. *Given a connected chordal graph $G = (V, E)$ and a sequence of visited vertices τ produced by an MCS with the current highest-label set S. Vertices $x, y \in S$ are connected in $G[S]$ iff they are connected in $G[V \setminus \tau]$.*

Algorithm MCS-ENUM utilizes these results to enumerate the AMOs of a chordal graph. Lines 1 to 4 in Algorithm 1 initialize the necessary data structures for an MCS (in particular an array A, which includes the set of vertices with i visited neighbors at index $A[i]$). Afterward, the recursive function `enumerate` is called. It first chooses any vertex v from the vertices with most visited neighbors (just as a normal MCS). This vertex is then removed from $A[i]$ in line 13, it is appended to τ in line 14, which stores the traversal sequence (later used to impose edge directions based on its ordering) and the neighbors are moved from $A[j]$ to $A[j+1]$ in lines 15 to 19. The recursive calls to `enumerate` are repeated until, at some point, the first AMO is output in line 7 (up to this point there is no difference to an MCS).

After the recursive call, however, the changes are reversed in lines 21 to 27 and then, in line 29, the vertices reachable from v are computed. These are then iterated in the do-while loop, meaning we also recursively go through the AMOs produced by choosing those vertices instead of v, as discussed above. Note that reachability in line 29 is only performed once in a call of `enumerate` for the initial vertex v.

Theorem 1. *Given a chordal graph G,* MCS-ENUM *enumerates all AMOs of G.*

Proof. Every DAG output in line 7 is an AMO, as it is generated by a linear ordering produced by an MCS. This holds as any chosen vertex is from the highest-index non-empty set in A. To see that the algorithm outputs *all* AMOs of G, recall that every AMO can be represented by an MCS ordering by Fact 2. In principle, Algorithm 1 considers all possible courses an MCS could take, except the pruning of vertices unreachable from v. By Lemma 2, it suffices to inspect only connected vertices in $G[A[i]]$ and by item 2 of Lemma 1 those unreachable vertices would not lead to any new AMO.

Finally, we argue that no AMO is output twice. Every output is obtained by constructing a directed graph based on the ordering given by the graph traversal. Assume for the sake of contradiction that we have two such sequences τ_1 and τ_2 representing the same AMO. Let x and y be the vertices in τ_1 and τ_2 at the first differing position, respectively. Note that x and y are connected and, hence, by item 1 of Lemma 1 it follows that τ_1 and τ_2 yield different AMOs. □

Theorem 2. MCS-ENUM *has worst-case delay $O(n + m)$.*

Proof. Let us partition the steps between two outputs in three phases: (i) the recursion goes "upwards" from an output; (ii) it reaches its "top" in the recursion tree; and (iii) the recursion goes "downwards" towards the next output.

We show that each phase runs in time $O(n + m)$. In phase (i), lines 21 to 31 are executed and the do-while loop stops (otherwise we would be in phase (ii)). The for-loop in lines 21 to 25 has time complexity $O(\delta(x))$. Moreover, the reachability query executed in line 29 does not yield any vertices (otherwise the do-while loop would continue), meaning it takes time $O(\delta(x))$ (all neighbors of v are checked once and the search stops). The run-time is therefore $O(m)$ as every edge is considered at most twice.

The main costs of phase (iii) are produced by the for-loop in lines 15 to 19, which requires time $O(\delta(x))$ leading to an overall time of $O(m)$. Both for-loops are executed in phase (ii), resulting in overall time $O(\delta(x))$. The reachability query from v in line 29 costs time $O(m)$. As this is done only once, we obtain a worst-case delay of $O(n + m)$. □

The result immediately generalizes to CPDAGs.

Theorem 3. *The Markov equivalence class $[G]$ of a CPDAG G can be enumerated with worst-case delay $O(n + m)$.*

Proof. Algorithm 1 also works for unconnected chordal graphs without any modifications. Hence, it can be called for the graph obtained by removing all directed edges of G. After computing an AMO of this graph, the directed edges can be re-added and the output is a member of $[G]$ produced with delay $O(n + m)$. The correctness follows from Fact 1. □

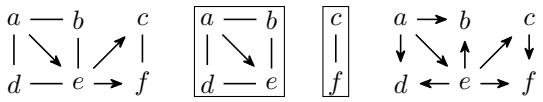

Figure 3: The MPDAG on the left, contains two buckets with chordal skeleton (middle), which lead to a consistent extension when each oriented into an AMO (right).

4 PDAGs and MPDAGs

One can naturally generalize the enumeration task to only consider members of the MEC, which conform to certain background knowledge given in the form of additional directed edges. Such a subclass of an MEC is commonly represented by an MPDAG (or PDAG). As before, a DAG is a consistent extension of G if it has the same skeleton and v-structures.It is easy to see that MEEK-ENUM works in this generalized setting as well, as does CHICKERING-ENUM (see also Corollary 2 in Section 5). In this section, we show how MCS-ENUM can be adapted, first to handle MPDAGs and, building on this, PDAGs.

Definition 1. *We term the graph induced by the vertices of a undirected component in an MPDAG a* bucket.

Buckets are similar to the undirected chordal components that we considered in the previous section, in that orienting each bucket in an MPDAG without cycles or v-structures will yield a consistent extension. However, a bucket may already contain directed edges. An illustration is given in Fig. 3.

Fact 3 (Wienöbst, Bannach, and Liśkiewicz (2021a))**.** *An MPDAG is extendable (i. e., has a consistent extension) iff the skeleton of every bucket is chordal. An extension can be computed in time $O(n + m)$.*

An MPDAG can be extended by running a modified MCS for each bucket B in the following way: The graph traversal is performed on the skeleton of B with the restriction that only vertices in the highest-label set S that have no unvisited parent in B are considered. These are the vertices x with $Pa(x) \setminus \tau = \emptyset$, given that τ contains the visited vertices, and we denote the set of such vertices by $S^+ \subseteq S$. This way, the MCS conforms to the background edges.

With these insights, an analogue modification of Algorithm 1 suggests itself to enumerate *all* AMOs of a bucket: Perform the algorithm on the skeleton of the bucket and only consider vertices in S^+.

Lemma 3. *Let B be a bucket and τ be a sequence of visited vertices with $|\tau| = k < n$ produced by the modified MCS using S^+. Then it holds that:*

1. *If $x, y \in S^+$ are connected in skel$(B[V \setminus \tau])$, the set of AMOs produced by choosing x next is disjoint from the set produced by choosing y next.*
2. *If $x, y \in S^+$ are unconnected in skel$(B[V \setminus \tau])$, any AMO produced by choosing y as the next vertex can also be produced by choosing a vertex in S^+ connected to x in skel$(B[V \setminus \tau])$ next.*

Proof. For item 1 consider the shortest path between x and y and assume that it contains directed edges (if not the argument of Lemma 1 applies). Then x or y have an incoming edge due to the non-applicability of the first Meek rule in the original bucket B. This violates the assumption that $x, y \in S^+$, meaning that τ_1 and τ_2 imply different AMOs. For item 2 the same argument as in Lemma 1 holds. □

Reachability can again be tested in a simplified way:

Lemma 4. *Given a bucket B and a sequence of visited vertices τ produced by the modified MCS using S^+. Vertices $x, y \in S^+$ are connected in $B[V \setminus \tau]$ iff they are connected in $B[S^+]$.*

Using these results, we can show that:

Theorem 4. *There is an algorithm that enumerates all AMOs of a given bucket B with worst-case delay $O(n + m)$.*

Proof sketch. Consider the just sketched algorithm, i. e., which proceeds as Algorithm 1 for the skeleton of B with the modification of choosing vertices and performing reachability with regard to S^+ (the algorithm is given explicitly in Appendix 2). By using S^+, the resulting AMOs conform with the directed edges in the bucket and due to Lemma 3 and 4 and by similar arguments as for Theorem 1 every such AMO is output exactly once. The linear-time delay follows as before, notably, S^+ can be efficiently maintained by storing the in-degree of each vertex. □

Using Fact 3, the result for buckets immediately generalizes to MPDAGs.

Corollary 1. *There is an algorithm that enumerates all consistent extensions of a given MPDAG with linear-time delay.*

The matter for PDAGs is similar as they can be maximally oriented into an equivalent MPDAG by Meek's rules.

Theorem 5. *There is an algorithm that enumerates all consistent extensions of a given PDAG with linear-time delay after an initialization step of time $O(n^3)$.*

Proof. The graph is initially transformed into its MPDAG. This is possible in time $O(n^3)$ as shown in (Wienöbst, Bannach, and Liśkiewicz 2021a). Afterward, apply the algorithm from Corollary 1. □

Wienöbst, Bannach, and Liśkiewicz (2021a) showed that the initialization step of maximally orienting a PDAG is likely not possible in linear time. However, using a finer complexity analysis, it can be performed in time $O(dm)$, where d is the *degeneracy* of the input's skeleton, which implies linear-time on many natural graph classes such as planar graphs, bounded-degree and bounded-treewidth graphs.

5 Another Approach for Enumerating Markov Equivalent DAGs

The results of the previous sections settle the worst-case complexity of enumerating the members of an MEC (at least if every DAG is output separately, a run-time of $o(n + m)$ is not achievable as this would be less than the size of the graph). In this section, we complement these results with an

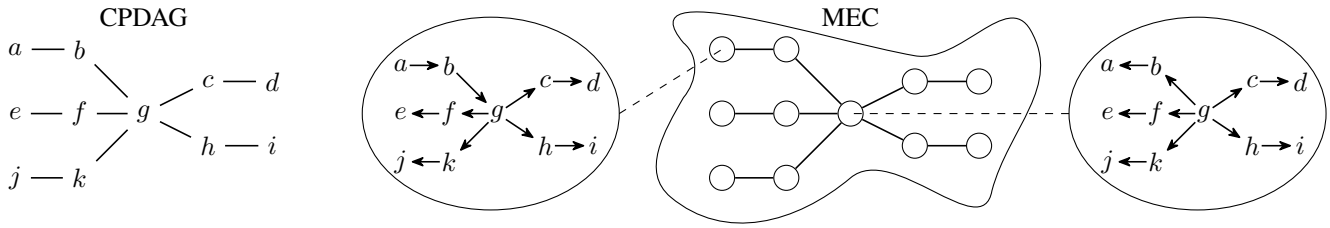

Figure 4: An example that has no sequence of SHD two that enumerates all Markov equivalent DAGs. Two DAGs in the MEC are connected by an edge if they can be transformed into each other by a single edge reversal. For trees, the resulting topology coincides with the one of the CPDAG, each DAG in the MEC can be represented by its unique source vertex. During the enumeration, the DAG in the center can be used only once, which makes it impossible to list all "leaf DAGs" when allowing only for distance at most two.

enumeration sequence of small changes between consecutive DAGs. While having a worse delay, such sequences are more natural from the causal perspective, with only a few edge orientations changing at a time, and provide structural insights into Markov equivalence itself. In more detail, we show that all graphs in an MEC can be enumerated in a sequence such that every two consecutive DAGs have structural Hamming distance (SHD) at most three. Our results are based on the following characterization of Markov equivalence, which is also the basis for CHICKERING-ENUM (see also Appendix 15):

Fact 4 (Chickering (1995))**.** *For any two Markov equivalent DAGs D and D' there is a sequence of Markov equivalent DAGs $\langle D = D_1, \dots, D_k = D' \rangle$ such that D_i and D_{i+1} have SHD one.*

The statement also holds for two consistent extensions of a *PDAG* in the sense that all intermediate DAGs are also consistent extensions of this PDAG.

Corollary 2. *For any two consistent extensions C and C' of PDAG G there is a sequence of consistent extensions of G $\langle C = D_1, \dots, D_k = C' \rangle$ such that D_i and D_{i+1} have SHD one.*

Proof. Only differing edges between C and C' are reversed in the constructive proof of Chickering (1995). Hence, all background edges stay fixed during the transformation. □

This means that it is possible to go from one DAG to another with only single edge reversals for CPDAGs as well as for PDAGs and MPDAGs. The task we are trying to solve, however, is to enumerate *all* members of an MEC, meaning the goal is to find a sequence in which *every* DAG occurs exactly once. It can be shown that such a sequence with SHD at most one does indeed *not* exist. Fig. 4 provides an example that does not even allow a sequence of SHD *two*.

However, if we permit *three* edge reversals between consecutive DAGs we can always find such a sequence:

Theorem 6. *Every MEC can be represented as sequence $\langle D_1, D_2, \dots \rangle$ of Markov equivalent DAGs such that D_i and D_{i+1} have SHD at most three.*

Proof. For a constructive, proof consider the graph that contains all DAGs in the MEC as nodes.[2] In that graph connect two nodes with an edge if the DAGs can be transformed into each other by a single edge reversal (hence, these have SHD one). By Fact 4 the graph is connected.

Every connected graph has a sequence $\langle p_1, p_2, \dots \rangle$ that contains every node exactly once such that the distance between consecutive nodes is at most three.[3] This sequence can be constructed by performing a depth-first-search (DFS) starting at an arbitrary node r and appending nodes with an even distance from r in the DFS tree when they are discovered and nodes with an odd distance from r when they are fully processed (essentially mixing pre- and post-order depending on the layer of the DFS tree). The SHD between two output nodes is never larger than three: When going down the DFS tree, every second node is output, when going up (after last outputting in odd layer i) the node in layer $i-2$ is output after it is finished. Hence, if it has no unvisited neighbors, the SHD is two. If it does, one of these gets explored and, hence, immediately output as it is in even layer $i-1$. In this case, the SHD to the last output is three. □

Due to Corollary 2, this result generalizes to PDAGs:

Corollary 3. *The consistent extensions of PDAG G can be represented as sequence $\langle D_1, D_2, \dots \rangle$ of consistent extensions of G with SHD at most three.*

We note that MEEK-ENUM, CHICKERING-ENUM and MCS-ENUM do not have this property.

Lemma 5. *Sequences of DAGs produced by* MEEK-ENUM *may contain consecutive DAGs with SHD larger than three.*

Proof. Consider the CPDAG shown in Fig. 4. Since MEEK-ENUM has no preferences on the edge it orients first, it may start with the edge $a \to b$. All other edge directions would then follow from the first Meek rule yielding the output DAG shown in Fig. 4. The orientation $a \leftarrow b$ is tried afterward,

[2]We use the term node instead of vertex here to avoid confusion with the vertices of the DAGs.

[3]The authors became aware of this graph property due to a problem posed by Jorke de Vlas in the annual programming contest BAPC (Problem H at BAPC 2021: https://2021.bapc.eu/).

which would result in no further directed edges. Then, assume the next undirected edge picked by the algorithm is the ones between h and i. It may be oriented as $h \leftarrow i$ yielding a DAG with SHD 4 to the previously output DAG. □

Similar arguments hold for CHICKERING-ENUM and MCS-ENUM, the former could end up in a state where the only DAGs left are the one with edge $a \rightarrow b$ and the one with $h \leftarrow i$ and the latter could start with vertex a and afterward choose i as the first vertex – yielding again the same DAGs.

Corollary 4. *Sequences of DAGs produced by* CHICKERING-ENUM *and* MCS-ENUM *may contain consecutive DAGs with structural Hamming distance larger than three.*

Computationally, our results do not imply a better bound on the delay in producing the sequence from Theorem 6 and we leave this as an open problem. The constructive algorithm (which we call SHD3-ENUM) given in the proof of Theorem 6 behaves similar to CHICKERING-ENUM and has delay $O(m^2)$ as every DAG may have m neighbors and we have to check for each of them whether they were already visited (between two outputs a constant number of recursive calls are handled; see Appendix 7). It seems unlikely that this can be improved without further structural insights.

Lastly, we remark that the same idea can also be used in the more general setting of enumerating maximal ancestral graphs (MAGs) without selection bias, which are causal models allowing for latent confounders and for which a similar transformational characterization exists (Zhang and Spirtes 2005). A brief introduction to MAGs and a more detailed analysis are given in Appendix C.

Corollary 5. *Every MEC of MAGs without selection bias can be represented as sequence $\langle M_1, M_2, \dots \rangle$ of Markov equivalent MAGs with SHD at most three.*

In contrast, approaches such as MEEK-ENUM do not exist for MAGs as analogue rules proposed by Zhang (2008) only complete the graph in case the edge marks are inferred by observational data. If edge marks are chosen for the sake of enumeration, these rules are not known to be complete. Generally, it is an interesting direction for future work to investigate the computational aspects of the enumeration of MECs of MAGs.

6 Experiments

In addition to the theoretical results, we also show that MCS-ENUM and its generalizations are practically implementable and significantly faster than previously used algorithms.

In Fig. 5, we compare the average delay of the four approaches (MEEK-ENUM, CHICKERING-ENUM, MCS-ENUM, SHD3-ENUM) implemented in Julia (Bezanson et al. 2017). For each instance, the programs were terminated after two minutes if the enumeration was not completed. As the enumeration problem reduces to listing the AMOs of a chordal graph (as shown in Fact 1), we consider as instances undirected graphs generated by randomly inserting edges, which do not violate chordality, until a graph with $3 \cdot n$ edges is reached. Note that these instances are all CPDAGs, just fully undirected ones, and thus all approaches can be applied to this setting. In Appendix D, we also compare the results for

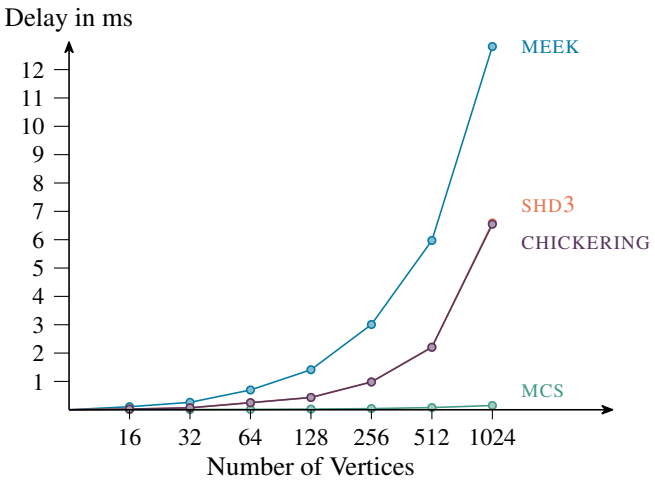

Figure 5: Average delay in milliseconds for enumerating the AMOs of random chordal graphs with $m = 3 \cdot n$ edges. We compare the algorithms MEEK-ENUM, CHICKERING-ENUM, MCS-ENUM and SHD3-ENUM.

CPDAGs *with* directed edges as well as for PDAGs, which both lead to very similar results. Moreover, we discuss the distribution of the delay for the various algorithms.[4]

The results clearly show that MCS-ENUM is by far the fastest among the algorithms. This is mainly due to the fact that the other algorithms always incur a cost of at least $\Omega(n + m)$, whenever a single edge is (re-)oriented. MEEK-ENUM needs to apply the four completion rules, whereas CHICKERING-ENUM and SHD3-ENUM require checking whether the resulting DAG was already output (which might often be the case). Still, the latter algorithms are significantly faster than MEEK-ENUM (at the cost of higher memory demand), and notably have both very similar delay (showing that the enumeration with SHD at most three gives mainly structural insights into Markov equivalence and has in itself no computational advantage).

7 Conclusion

We have given the first formal and exhaustive treatment of the fundamental problem of enumerating Markov equivalent DAGs. Our main results are twofold: (i) we significantly improve the run-time of enumeration by giving the first linear-time delay algorithm, which is also practically effective and (ii) we give structural insights into Markov equivalence by constructing an enumeration sequence with minimal distance between successive graphs. The concepts for (ii) are so general that they directly apply to MAGs without selection bias as well.

As an open problem, it remains to find more efficient enumeration algorithms for MAGs, where, currently, approaches in the spirit of both MEEK-ENUM and MCS-ENUM cannot be applied, because similar structure does not exist for Markov equivalence of MAGs or, at least, is not known.

[4]The implementations of the algorithms are available at https://github.com/mwien/mec-enum.

Acknowledgments

The research of Malte Luttermann was partly supported by the Medical Cause and Effects Analysis (MCEA) project. The work of the last author was partly supported by the Deutsche Forschungsgemeinschaft (DFG) grant 471183316.

References

Albert, R.; and Barabási, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1): 47–97.

Andersson, S. A.; Madigan, D.; and Perlman, M. D. 1997. A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25(2): 505–541.

Bezanson, J.; Edelman, A.; Karpinski, S.; and Shah, V. B. 2017. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1): 65–98.

Chen, E. Y.-J.; Choi, A. C.; and Darwiche, A. 2016. Enumerating equivalence classes of Bayesian networks using EC graphs. In *Artificial Intelligence and Statistics*, 591–599. PMLR.

Chickering, D. M. 1995. A transformational characterization of equivalent Bayesian network structures. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence, UAI '95*, 87–98.

Chickering, D. M. 2002. Optimal Structure Identification With Greedy Search. *Journal of Machine Learning Research*, 3: 507–554.

Elwert, F. 2013. Graphical Causal Models. In *Handbook of Causal Analysis for Social Research*, Handbooks of Sociology and Social Research, 245–273. Springer.

Frydenberg, M. 1990. The chain graph Markov property. *Scandinavian Journal of Statistics*, 333–353.

Gillispie, S. B.; and Lemieux, C. 2001. Enumerating Markov Equivalence Classes of Acyclic Digraph Models. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, UAI '01*, 171–177.

Hauser, A.; and Bühlmann, P. 2012. Characterization and Greedy Learning of Interventional Markov Equivalence Classes of Directed Acyclic Graphs. *Journal of Machine Learning Research*, 13: 2409–2464.

He, Y.; Jia, J.; and Yu, B. 2015. Counting and Exploring Sizes of Markov Equivalence Classes of Directed Acyclic Graphs. *Journal of Machine Learning Research*, 16(79): 2589–2609.

He, Y.-B.; and Geng, Z. 2008. Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research*, 9(Nov): 2523–2547.

Kalisch, M.; Mächler, M.; Colombo, D.; Maathuis, M. H.; and Bühlmann, P. 2012. Causal inference using graphical models with the R package pcalg. *Journal of statistical software*, 47: 1–26.

Koller, D.; and Friedman, N. 2009. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press. ISBN 978-0-262-01319-2.

Maathuis, M. H.; Kalisch, M.; and Bühlmann, P. 2009. Estimating High-Dimensional Intervention Effects from Observational Data. *The Annals of Statistics*, 37(6A): 3133–3164.

Meek, C. 1995. Causal Inference and Causal Explanation with Background Knowledge. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence, UAI '95*, 403–410.

Pearl, J. 2009. *Causality*. Cambridge University Press. ISBN 978-0521895606.

Richardson, T.; and Spirtes, P. 2002. Ancestral graph Markov models. *The Annals of Statistics*, 30(4): 962–1030.

Rothman, K. J.; Greenland, S.; Lash, T. L.; et al. 2008. *Modern epidemiology*, volume 3. Wolters Kluwer Health/Lippincott Williams & Wilkins Philadelphia.

Spirtes, P.; Glymour, C.; and Scheines, R. 2000. *Causation, Prediction, and Search, Second Edition*. MIT Press. ISBN 978-0-262-19440-2.

Squires, C. 2018. `causaldag`*: creation, manipulation, and learning of causal models*.

Steinsky, B. 2003. Enumeration of labelled chain graphs and labelled essential directed acyclic graphs. *Discrete mathematics*, 270(1-3): 267–278.

Tarjan, R. E.; and Yannakakis, M. 1984. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on computing*, 13(3): 566–579.

Verma, T.; and Pearl, J. 1990. Equivalence and Synthesis of Causal Models. In *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence, UAI'90*, 255–270.

Wienöbst, M.; Bannach, M.; and Liśkiewicz, M. 2021a. Extendability of Causal Graphical Models: Algorithms and Computational Complexity. In *Proceedings of the 37th Conference in Uncertainty in Artificial Intelligence, UAI '21*. AUAI Press.

Wienöbst, M.; Bannach, M.; and Liśkiewicz, M. 2021b. Polynomial-Time Algorithms for Counting and Sampling Markov Equivalent DAGs. In *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI 2021*, volume 35, 12198–12206. AAAI Press.

Zhang, J. 2008. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17): 1873–1896.

Zhang, J.; and Spirtes, P. 2005. A transformational characterization of Markov equivalence for directed acyclic graphs with latent variables. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence, UAI '05*, 667–674.

Appendix

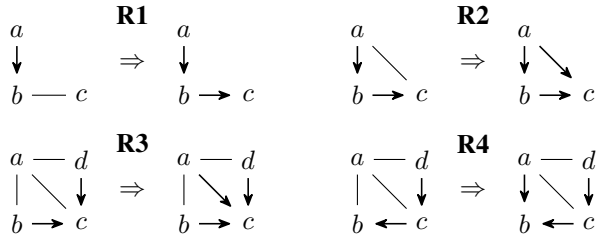

Figure 6: The four Meek rules that are used to characterize MPDAGs (Meek 1995).

A Formal Description of Related Algorithms

Enumerating Markov Equivalent DAGs Based on Meek's Rules

Meek (1995) gave a complete set of four rules (R1 - R4), which, when applied repeatedly, transform a PDAG into its maximal orientation, i. e., orient all undirected edges, which are fixed in the DAGs represented by the PDAG. Graphs completed under the Meek rules are also called *Maximally oriented PDAGs* (MPDAGs).

See R1 in Figure 6 as an example for the application of the rules. The edge $b - c$ necessarily has a fixed orientation because $a \to b \leftarrow c$ and $a \to b \to c$ are in different MECs (the first graph contains a v-structure, the latter does not) and hence cannot be represented by the same PDAG. The single DAG represented by the PDAG in this case is $a \to b \to c$, the graph which does not introduce a new v-structure, and it is immediately obtained after applying R1. In most cases some undirected edges remain even after application of the Meek rules.

In the converse, a graph completed under the Meek rules has the property that every undirected edge $x - y$ is *not* fixed in the DAGs it represents, i. e., there is a DAG in the class which contains $x \to y$ and one which contains $x \leftarrow y$. Thus, the members can be enumerated by first orienting the edge as $x \to y$ (and recursively orienting the remaining edges one-by-one, always after completing the graph under the Meek rules to ensure that every orientation yields a valid DAG) and afterward orienting the edge as $x \leftarrow y$. This approach we call MEEK-ENUM is depicted in Algorithm 2.

MEEK-ENUM has polynomial-delay, particularly as every orientation leads to a valid DAG (there are no "dead-ends" in the recursive search). However, it requires the application of Meek's rules in every step, which amounts to significant effort to perform between two consecutive outputs. Naively estimated, checking whether one of the four Meek rules applies takes time $O(n^4)$ and in a worst-case scenario this check has to be performed multiple times as more and more edges get oriented successively. This would lead to a polynomial time requirement of large degree, even for applying the Meek rules once. However, there have been recent algorithmic improvements in this regard.

input : A CPDAG $G = (V, E)$.
output : $[G]$.

function `enumerate`(G)
  $H :=$ maximal orientation of G
  undir $:= \{\{u, v\} \mid (u, v) \in E_H \wedge (v, u) \in E_H\}$
  if undir *is empty* **then**
    Output H
  else
    $\{u, v\} :=$ any element from undir
    $E_H := E_H \setminus \{(u, v)\}$ `// Orient` $u \leftarrow v$
    `enumerate`(H)
    $E_H := E_H \cup \{(u, v)\}$ `// Undo`
    $E_H := E_H \setminus \{(v, u)\}$ `// Orient` $u \to v$
    `enumerate`(H)
    $E_H := E_H \cup \{(v, u)\}$ `// Undo`
  end
end

Algorithm 2: Enumeration algorithm of Markov equivalent DAGs based on Meek's rules (MEEK-ENUM).

Fact 5 (Wienöbst, Bannach, and Liśkiewicz (2021a))**.** *Given a PDAG G, Meek's rules can be applied exhaustively on G in time $O(n^3)$.*

This gives a slightly better bound on the achievable delay, formally stated in the following Theorem.

Theorem 7. MEEK-ENUM *can be implemented such that the MEC is enumerated with worst-case delay $O(m \cdot n^3)$.*

Proof. Clearly, every DAG in the MEC is output as every undirected edge is oriented in both directions (by soundness of the Meek rules every directed edge is actually fixed in the DAGs in the class). No DAG is output twice because in each step, an undirected edge is oriented and fixed, i.e., every recursive call receives a different graph as input. Nor is an invalid DAG output, because every undirected edge can be oriented in both directions leading to a valid DAG by completeness of the Meek rules. In the worst case, the algorithm needs m recursive calls until a DAG is output, and a single call costs time $O(n^3)$ due to the application of Meek's rules by Fact 5. Hence, the delay between two outputs is bounded by $O(m \cdot n^3)$. □

The same result holds for PDAGs and MPDAGs without modifications. We note that in the experiments, we decided to implement the Meek rules in the "naive" way as this is typically used in implementations, which makes for a more meaningful comparison.

Enumerating Markov Equivalent DAGs Based on Chickering's Transformational Characterization

Another approach for enumerating Markov equivalent DAGs is build on the characterization by Chickering (1995), also stated briefly in Section 5 above. It is based on the notion of *covered edges*:

Definition 2 (Chickering (1995))**.** *An edge $x \to y$ is called* covered *if* $\mathrm{Pa}(x) \cup \{x\} = \mathrm{Pa}(y)$.

This allows the following characterization of Markov equivalence:

Theorem 8 (Chickering (1995))**.** *Let D and D' be two Markov equivalent DAGs. There exists a sequence of $\mathrm{shd}(D, D')$ edge reversals in D with the following properties:*

1. *Each edge reversed in D is a covered edge.*
2. *After each reversal, D is a DAG Markov equivalent to D'.*
3. *After all reversals, $D = D'$.*

Hence, starting from a DAG D we reach all DAGs in the same Markov equivalence class by reversals of covered edges. This permits the following enumeration approach, which we call CHICKERING-ENUM.

input : A CPDAG $G = (V, E)$.
output : $[G]$.

```
D := any DAG in [G]
vis := {D}
enumerate(D)
function enumerate(D, vis)
    Output D
    foreach DAG D' obtained by reversing a covered edge in D do
        if D' ∉ vis then
            vis := vis ∪ {D'}
            enumerate(D', vis)
        end
    end
end
```

Algorithm 3: Enumeration algorithm of Markov equivalent DAGs based on Chickerings' characterization (Chickering 1995) (CHICKERING-ENUM).

Starting from D, similar to a depth-first-search (DFS), all neighbors of D (graphs with a single reversed covered edge) are explored and this is continued recursively. Eventually all Markov equivalent DAGs are reached by Theorem 8 above. In order to not visit any DAG twice, it is necessary to store a set of all visited DAGs.

Theorem 9. CHICKERING-ENUM *enumerates a Markov equivalence class and can be implemented with worst-case delay $O(m^3)$.*

Proof. Clearly, any DAG is output exactly once by Theorem 8 and storing all visited DAGs ensures not outputting any DAG multiple times.

For the delay, it is first useful to analyze when the most steps between two outputs are necessary. This happens when the algorithm traverses back up the recursion tree (from a leaf), potentially up to the root, without outputting any new DAG and then down into another subtree. In principle, the recursion depth might be exponential and this would, in turn, lead to an exponential-time delay.

However, it is possible to bound the recursion depth by m (by Theorem 8 every edge is reversed at most once) and one will still find all DAGs in the MEC. It remains to estimate the cost at each recursion step. Going up the recursion tree, at each DAG, $O(m)$ neighbors might be checked whether they have been visited before, e. g., if they are in vis, and all of them might indeed are. The check takes time $\Omega(m)$ per neighboring DAG D' as the whole DAG has to be "read" at least once for lookup. If one uses a hash table for storing the DAGs, expected time $O(m)$ can be reached.

So, in total, we have $O(m)$ recursion steps between two outputs, with at most $O(m)$ neighboring DAGs being considered, each of which consuming $O(m)$ time for lookup in vis. This leads to a worst-case delay of $O(m^3)$. □

Computing a Consistent Extension of a CPDAG based on the MCS Algorithm

In this subsection, we revisit the closely related problem of finding an arbitrary DAG in an MEC represented by a CPDAG G instead of listing all of them. This is known as the *extension* task and it is well-known that it is possible to perform it in linear-time (Hauser and Bühlmann 2012) for CPDAGs using the so-called Lexicographic Breadth-First Search (LBFS) algorithm, which is a graph traversal algorithm originally proposed for testing chordality of a graph. For MPDAGs and PDAGs, this problem is discussed in depth in (Wienöbst, Bannach, and Liśkiewicz 2021a).

As we base our results on the Maximum Cardinality Search (MCS) algorithm, which has been used less frequently in the causality community compared to LBFS, we give a brief introduction. The MCS algorithm has been designed, as LBFS, for testing chordality of a graph. It is a graph traversal algorithm, which visits the vertices of an undirected graph in an order representing[5] an AMO (*acyclic moral orientation*, that is, an orientation without directed cycles and v-structures) iff the graph is chordal.[6]

The connection between AMOs and CPDAGs stems from the fact that the undirected components of a CPDAG are chordal and replacing each by an AMO leads to a DAG in the Markov equivalence class. The reason for this is that a CPDAG is extended into a DAG by orienting its undirected edges without creating a directed cycle or a new v-structure. This is why AMOs are needed, and the undirected components are chordal, because only chordal graphs have AMOs.

[5] A linear ordering of the vertices represents the orientation of an undirected graph, in which edges are directed $x \to y$ if x comes before y in the ordering.

[6] In the chordal graph literature the term *perfect elimination ordering* is more commonly used in this setting. A linear-ordering of the vertices describes an AMO iff its reverse is a perfect elimination ordering. We will stick to AMOs in this paper to avoid confusion.

input : A CPDAG $G = (V, E)$.
output : DAG $D \in [G]$.

1 $D :=$ copy of G
2 $U := (V, \{(u, v) \mid (u, v)$ and $(v, u) \in E\})$, i. e., the graph with only the undirected edges of G
3 **foreach** *connected component* $C = (V_C, E_C)$ *of* U **do**
4 $\quad \tau_C = \texttt{mcs}(C)$
5 $\quad$ Orient edges in $D[V_C]$ according to τ_C, i. e., $x - y$ as $x \to y$ if x comes before y in τ_C
6 **end**
7 **function** `mcs` (C)
8 $\quad A :=$ array of n initially empty sets
9 $\quad \tau :=$ empty list
10 $\quad A[0] := V$
11 $\quad$ **while** $|\tau| < |V_C|$ **do**
12 $\quad\quad i :=$ highest index of non-empty set in A
13 $\quad\quad v :=$ any vertex from $A[i]$
14 $\quad\quad$ delete v from $A[i]$
15 $\quad\quad$ append v to τ
16 $\quad\quad$ **foreach** $w \in (Ne(v) \setminus \tau)$ **do**
17 $\quad\quad\quad j :=$ index of set in A that w is in
18 $\quad\quad\quad$ delete w from $A[j]$
19 $\quad\quad\quad$ insert w in $A[j + 1]$
20 $\quad\quad$ **end**
21 $\quad$ **end**
22 $\quad$ **return** τ
23 **end**

Algorithm 4: Computing a DAG in the MEC represented by CPDAG G in linear-time $O(n + m)$.

We can now look at Algorithm 4, which implements this strategy, in more detail. We discussed the first part (lines 2 to 6) in detail. The UCCGs (i.e., the connected components when all directed edges are removed from G) are oriented one-by-one. For this, an ordering τ is computed according to which the edges are oriented. This ordering τ is the traversal sequence of an MCS. This algorithm is given in lines 7 to 23. It differs from standard graph traversals in the way the next vertex to visit is chosen (line 13). The algorithm chooses one of the vertices with most previously visited neighbors. This is implemented efficiently by having a set of vertices with 0 previously visited neighbors ($A[0]$), a set for vertices with 1 previously visited neighbor ($A[1]$), and so on. When a vertex v is handled (lines 14 to 20), its neighbors are moved to a higher set (line 19).

This can be done in constant time per neighbor (Tarjan and Yannakakis 1984), meaning the MCS runs in linear-time. As it is performed on disjoint subgraphs, we obtain overall linear-time $O(n + m)$ for computing a DAG in $[G]$.

Theorem 10. *Algorithm 4 computes a DAG in the MEC represented by CPDAG G in time $O(n + m)$.*

B Missing Proofs and Algorithms

Section 3

We begin by giving the proof of Lemma 2:

Lemma 2. *Given a connected chordal graph $G = (V, E)$ and a sequence of visited vertices τ produced by an MCS with the current highest-label set S. Vertices $x, y \in S$ are connected in $G[S]$ iff they are connected in $G[V \setminus \tau]$.*

Proof. Trivially, if x and y are connected in $G[S]$ so they are in $G[V \setminus \tau]$. Hence, we only have to prove that if they are not connected in $G[S]$, then they are also disconnected in $G[V \setminus \tau]$. The special case that the highest label has label 0 is trivial as then we have $S = V \setminus \tau$. So let τ be non-empty and consider two arbitrary vertices $x, y \in S$ that are in different connected components of $G[S]$. For a contradiction assume that they are connected in $G[V \setminus \tau]$ via a shortest path $\pi = x - p_1 - \cdots - p_\ell - y$ with $\ell \geq 1$ and $p_1, \ldots, p_\ell \in V \setminus \tau$. Note that π is an unchorded path (i. e., one where only successive vertices are connected by an edge) due to being the shortest path between x and y. Also denote the visited neighbors of vertex v by $P(v) = \tau \cap Ne(v)$.

In π, there has to be a vertex $p_i \notin S$, else x and y are connected in $G[S]$. Hence, there is (i) a vertex $z \in P(x)$, which is not in $P(p_i)$, and (ii) a vertex $z' \in P(y)$, which is not in $P(p_i)$, because x and y have higher label than p_i, meaning $|P(x)| > |P(p_i)|$ and $|P(y)| > |P(p_i)|$.

We first consider the case that z or z' is a common neighbor of both x and y (this includes the case $z = z'$). Let this common neighbor be w.l.o.g. vertex z. Then, there is a cycle $x - p_1 - \cdots - p_\ell - y - z - x$ in G. Moreover, p_i is not a neighbor of z, vertices x and y are nonadjacent (else they would be connected in $G[S]$), and due to π being the shortest path, there are no edges between $x - p_j$ for $j > 1$, $y - p_{j'}$ for $j' < \ell$ and $p_j - p_{j'}$ for $|j - j'| > 1$. Hence, the only chords the cycle might have could be $p_j - z$ edges, with $j \neq i$. But there could only be $l - 1$ such chords and for a cycle of length $l + 3$, l chords are needed to make it chordal. A contradiction.

The remaining case is that z and z' are both not a common neighbor of x and y. We show that there can be no AMO produced by the MCS, when it chooses x or y next from S, which would be a contradiction. W.l.o.g., we show the argument for x. Let α be a topological ordering inducing an AMO of G having prefix $\tau + x$. Denote by $\alpha^{-1}(v)$ the position of v in α. It has to hold that $\alpha^{-1}(x) < \alpha^{-1}(p_1) < \alpha^{-1}(p_2) < \cdots < \alpha^{-1}(p_\ell) < \alpha^{-1}(y)$ as the first vertex p_j with $\alpha^{-1}(p_j) > \alpha^{-1}(p_{j+1})$, would induce a v-structure. As we have $\alpha^{-1}(z') < \alpha^{-1}(p_\ell) < \alpha^{-1}(y)$, vertices z' and p_ℓ have to be adjacent to avoid a v-structure. Then, the same applies to $\alpha^{-1}(z') < \alpha^{-1}(p_{\ell-1}) < \alpha^{-1}(p_\ell)$, which implies an edge between z' and $p_{\ell-1}$. This iteration can be continued only until vertex p_i, which is nonadjacent to z' by assumption. □

For the sake of completeness, we also explicitly state the EnumMCS algorithm for a CPDAG G. The enumeration task immediately reduces to finding the AMOs of the chordal components of G.

Section 4

We give the proof of Lemma 4 omitted in the main paper:

Lemma 4. *Given a bucket B and a sequence of visited vertices τ produced by the modified MCS using S^+. Vertices*

input : A CPDAG $G = (V, E)$.
output : $[G]$.

1 $U := (V, \{(u, v) \mid (u, v) \text{ and } (v, u) \in E\})$, i. e., the graph with only the undirected edges of G
2 Execute Algorithm 1 on U and each time add the directed edges of G when outputting the DAG (in line 7 of Algorithm 1).

Algorithm 5: Linear-time delay enumeration algorithm of Markov equivalent DAGs based on EnumMCS.

$x, y \in S^+$ are connected in $B[V \setminus \tau]$ iff they are connected in $B[S^+]$.

Proof. Recall that S is the set of *all* highest-label vertices including the ones with unvisited parents. As shown in the proof of Lemma 2 the vertices x and y are connected in S by the chordality of B. For a contradiction assume that x and y are connected in S but not in S^+. Then there is a shortest path $x = p_1 - p_2 - \cdots - p_{k-1} - p_k = y$ with some $p_i \in S \setminus S^+$. Consider the p_i with smallest i and assume the shortest path is chosen such that this i is maximized (in case that there are multiple shortest paths). This path is completely undirected as x or y would otherwise have an incoming edge by the same argument as in the proof of Lemma 3.

Observe that p_i has an incoming edge from an unvisited vertex z_1. By the properties of the MCS we have $z_1 \in S$ as otherwise z_1 would not be connected to a previous neighbor a of p_i implying the v-structure $a \to p_i \leftarrow c$ and violating the fact that the modified MCS returns an AMO. Therefore, z_1 needs to be connected to p_{i-1} and p_{i+1}. Furthermore, if z_1 has an incoming edge from an unvisited vertex z_2 then this vertex has to be connected to p_{i-1} and p_{i+1} as well.

This process can be iterated further and we will consider the first z_j that has no unvisited parent. Such a vertex has to exist as otherwise there would be a directed cycle. The edge between z_j and p_{i-1} is undirected by the minimality of p_i.

The edge between z_j and p_{i+1} may be directed. But if it is, there has to be an edge $z_j - p_{i+2}$. Then for some p_l there is an undirected edge $z_j - p_l$ as the edge $z_j - y$ has to be undirected. This yields the desired contradiction as either the path $x = p_1 - \cdots - p_{i-1} - z_j - p_l - \ldots p_k = y$ is shorter than the previously considered one or the first vertex with an unvisited parent has a higher index. □

We explicitly state the generalization of EnumMCS to buckets, which was only sketched in the main paper. The differences to Algorithm 1 are highlighted.

Theorem 11. *There is an algorithm that enumerates all AMOs of a given bucket with worst-case delay $O(n + m)$.*

Proof. This algorithm is given in Algorithm 6. We first show that every DAG that is output by the algorithm is an AMO. This follows from the fact that the output is produced by an MCS. Notice that the restriction to vertices with no unvisited parents leads to no violation of this fact as the algorithm still *only* chooses vertices with the highest-label. That such a vertex always exists was shown in (Wienöbst, Bannach, and Liśkiewicz 2021a).

input : A bucket $B = (V, E)$.
output : All AMOs of B.

1 $G :=$ skeleton of B
2 $A :=$ array of n initially empty sets
3 $\tau :=$ empty list
4 $A[0] := V$
5 `enumerate`$(G, A, \tau,$ B $)$

6 **function** `enumerate`$(G,$ $A,$ $\tau,$ B $)$
7 **if** $|\tau| = n$ **then**
8 Output AMO of G according to ordering τ
9 **end**
10 $i :=$ highest index of non-empty set in A
11 $S^+ := \{v \mid v \in A[i] \text{ and } \mathrm{Pa}_{B[V \setminus \tau]}(v) = \emptyset\}$
12 $v :=$ any vertex from S^+
13 $x := v$
14 **do**
15 delete x from $A[i]$
16 append x to τ
17 **foreach** $w \in (Ne(x) \setminus \tau)$ **do**
18 $j :=$ index of set in A that w is in
19 delete w from $A[j]$
20 insert w in $A[j + 1]$
21 **end**
22 `enumerate`$(G,$ $A,$ $\tau)$
23 **foreach** $w \in (Ne(x) \setminus \tau)$ **do**
24 $j :=$ index of set in A that w is in
25 delete w from $A[j]$
26 insert w in $A[j - 1]$
27 **end**
28 insert x in $A[i]$
29 pop x from τ
30 **if** $x = v$ **then**
31 $R := \{a \mid a \text{ reachable from } v \text{ in } G[S^+]\}$
32 **end**
33 **while** R *is non-empty*, $x := pop(R)$
34 **end**

Algorithm 6: Linear-time delay enumeration algorithm of consistent extensions of a bucket.

The next thing we prove is that every AMO of the given bucket is output. By Fact 2 every AMO can be represented by an MCS ordering. Which particular AMO is produced by an MCS depends on the specific choices of highest-label vertices. Our algorithm considers all possible choices at each step except (i) the ones where a vertex has unvisited parents and (ii) the ones where the vertex is unreachable from the first chosen vertex v in a recursion step.

Clearly, (i) only excludes AMOs incompatible with the given directed edges (background knowledge). For (ii) observe that if a vertex is unreachable from v, it does not matter whether it is chosen before or after v by item 2 of Lemma 3.

Finally, it remains to show that no AMO is output twice. Clearly, every sequence τ that we obtain is different. So assume for the sake of contradiction that the algorithm outputs two sequences τ_1 and τ_2 that represent the same AMO. Let x and y be the vertices in τ_1 and τ_2 at the first differing position. Then x and y are connected in the induced subgraph over the unvisited vertices (by line 29) in the skeleton of B. A contradiction since this implies that the AMOs represented by τ_1 and τ_2 differ.

For the complexity analysis we partition the algorithm into three phases (i), (ii), and (iii) as in the proof of Theorem 2. The cost of (ii) and (iii) are the same as in the analysis of Theorem 2. For (i) just observe that reachability runs in $O(d)$ due to Lemma 4. □

We also explicitly state the algorithm for enumerating PDAGs (Algorithm7). The algorithm also works for MPDAGs, in which case the highlighted line 4 does not need to be executed. The algorithm computes the buckets and runs Algorithm 6 on them.

input : A PDAG $G = (V, E)$.
output : All consistent extensions of G.

1 **if** *G has no consistent extensions* **then**
2 **return** $\emptyset$
3 **end**
4 $G :=$ the maximal orientation of G (i. e., its completion under the Meek rules)
5 $D :=$ copy of G
6 $B := (V, \{(u, v) \mid (u, v) \in E$ and u, v are connected by undirected edges in $G\})$, i. e., the graph containing only the buckets of G
7 Execute Algorithm 6 on B and each time add the remaining edges of G when outputting the DAG (in line 8 of Algorithm 6).

Algorithm 7: Linear-time delay enumeration algorithm of the consistent extensions of PDAG G based on MCS-ENUM for buckets (Algorithm 6).

Section 5

We complement the structural results in Section 5 in the main paper, by explicitly stating the algorithm (called SHD3-ENUM) described textually in the proof of Theorem 6 and analyzing its delay. The differences to CHICKERING-ENUM (Algorithm 3) are highlighted.

input : A CPDAG $G = (V, E)$.
output : $[G]$.

1 $D :=$ any DAG in $[G]$
2 vis $:= \{D\}$
3 `enumerate`$(D, \text{vis}, 0)$
4 **function** `enumerate`$(D,$ vis, $i)$
5 **if** $i \bmod 2 = 0$ **then**
6 Output D
7 **end**
8 **foreach** *DAG D' obtained by reversing a covered edge in D* **do**
9 **if** $D' \notin$ vis **then**
10 vis $:=$ vis $\cup \{D'\}$
11 `enumerate`$(D',$ vis$)$
12 **end**
13 **end**
14 **if** $i \bmod 2 = 1$ **then**
15 Output D
16 **end**
17 **end**

Algorithm 8: Enumeration algorithm with SHD ≤ 3 of Markov equivalent DAGs based on Chickerings' characterization (Chickering 1995) (SHD3-ENUM).

The main paper uses this algorithm to show that it is possible to enumerate an MEC with successive DAGs having SHD at most three. Here, we study the algorithmic properties of the algorithm.

Theorem 12. *For a CPDAG G,* SHD3-ENUM *enumerates the DAGs in $[G]$ with delay $O(m^2)$.*

Proof. The argument is similar to the proof of Theorem 9. First, it is clear that all DAGs in the MEC are output exactly once.

The analysis of the delay differs slightly in the sense that this algorithm has worst-case delay $O(m^2)$ (instead of $O(m^3)$) due to the fact that between two outputs only a constant number of recursive calls are handled. This follows from the fact that the algorithm outputs successive DAGs with SHD ≤ 3 and, hence, between the output of these DAGs, there are only constantly many steps in the traversal of the MEC. The cost per DAG can be bounded by $O(m^2)$ as in Theorem 9. □

Both Algorithm 8 as well as Theorem 12 directly apply to PDAGs as well.

C MAGs – Causal Graphs under Latent Confounding

MAGs without selection bias are mixed graphs $G = (V, E)$ with two types of edges: directed $x \to y$ and bidirected $x \leftrightarrow y$. The semantics are different compared to DAGs in that edges encode ancestral relations, i. e., a directed edge

means that x is an ancestor (a cause, but not necessarily a direct cause) of y and a bidirected edge means that neither x is a cause of y nor is y a cause of x (as simplification, think that there is a latent confounder between x and y). An excellent introduction to MAGs is given in (Zhang 2008).

Zhang and Spirtes (2005) gave the following transformational characterization of Markov equivalent MAGs without selection bias (also called *DMAGs*):

Theorem 13 (Zhang and Spirtes (2005))**.** *Two DMAGs G and G' are Markov equivalent iff there exists a sequence of single edge mark changes in G such that*

1. *after each mark change, the resulting graph is also a DMAG and is Markov equivalent to G,*
2. *after all the mark changes, the resulting graph is G'.*

It follows that the graph over Markov equivalent DMAGs with an edge between DMAGs with SHD 1 (we define the SHD similar to the DAG case, that is, the number of differing edges; one could also define it to be the number of differing edge marks, and the statement still holds) is connected and hence, there is, by the same argument as in Theorem 6, a sequence of DMAGs containing each in the MEC exactly once with successive DMAGs having SHD ≤ 3, which proves Corollary 5.

We consider it out of the scope of this work to investigate the algorithmic aspects of enumerating Markov equivalent MAGs in detail, but deem this an fascinating topic for further research. In particular, it would be highly interesting to develop enumeration algorithms based on the other strategies discussed in this paper.

Finally, we note that the transformational characterization in Theorem 13 does not hold for MAGs *with* selection bias and it follows from the example given in (Zhang and Spirtes 2005) that it is indeed not possible to find an enumeration sequence with SHD ≤ 3 in this case.

D Further Experimental Results

This section shows the full experimental results of the four algorithms (MEEK-ENUM, CHICKERING-ENUM, MCS-ENUM, SHD3-ENUM) with results for CPDAGs and PDAGs, as well as for slightly denser graphs. Fig. 7 depicts the results. The instances are generated as follows:

- For the undirected chordal graphs, a random tree was generated, to which randomly drawn edges are inserted, which do not violate chordality. This is done until the graph has $k \cdot n$ edges. In the main paper, we considered $k = 3$, here we choose slightly denser graphs with $k = \log_2 n$.
- The CPDAGs are generated by first creating a DAG D randomly, whose CPDAG representation is computed afterward. The DAG is, on the one hand, sampled analogously to the undirected case, by adding directed edges, which do not violate acyclicity until $k \cdot n$ edges are reached. On the other hand, we generate DAGs based on scale-free graphs (using the Barabási-Albert model (Albert and Barabási 2002)), which are oriented by imposing a random topological ordering.
- The PDAGs are generated by first sampling CPDAGs (in the way described above) and then orienting further edges (imitating the addition of background knowledge). For this, between three and seven undirected edges (this number is chosen uniformly at random) are oriented in the CPDAG, with the Meek rules being applied after each orientation.

For CPDAGs and PDAGs, as for the chordal graphs, we choose the parameters $k = 3$ and $k = \log_2 n$. The four algorithms were run on each of the instances for a maximum of one minute and the delay between outputs was averaged. In case of the smaller instances, the faster algorithms were often able to enumerate the whole MEC, for larger instances this is not realistically possible due to their vast size.

For the experiments, we used a single core of the AMD Ryzen Threadripper 3970X 32-core processor on a 256GB RAM machine.[7] We report the times in milliseconds and excluded averages over 32ms from the plots to keep the results visually comparable.

All results in Fig. 7 show the same pattern, namely that MCS-ENUM is superior to the other algorithms and both CHICKERING-ENUM and SHD3-ENUM outperform MEEK-ENUM. The superiority of MCS-ENUM can be easily explained by the fact that its competitors have cost at least in the size of the graph $O(n + m)$ after every (re)-orientation of an edge. While this is the *total cost* between two outputs for MCS-ENUM. The algorithms CHICKERING-ENUM and SHD3-ENUM have a very similar average delay, only the order of the output differs (in the way that SHD3-ENUM guarantees smoothly changing DAGs). MEEK-ENUM has by far the highest cost, due to the large effort arising from the repeated completion under the Meek rules.

Finally, we take a look at the distribution of delays, which we also recorded during the measurements. In Table 1, the percentage of delays d with $d \leq k \cdot \text{mean}$ is given for each algorithm and different k in every considered scenario.

It becomes evident that more than 90 percent of the measured delays are not larger than two times the mean. For the algorithms CHICKERING-ENUM, MCS-ENUM, and SHD3-ENUM, the proportion is even bigger, often between 98 and 99 percent. MEEK-ENUM has the biggest proportion of delays being larger than twice the mean, which can be explained by the fact that some sub-steps need significantly more time to apply Meek's rules than others. Outliers being larger than seven times the mean are nearly as often as outliers being larger than ten times the mean, showing that there are few (in fact, less than 0.5 percent) delays that are way above the mean. We conjecture that these infrequent outliers might be due to garbage collection during the measurements (which becomes a noticeable factor in millisecond measurements).

[7] The experiments can also be run on a desktop computer without any problems.

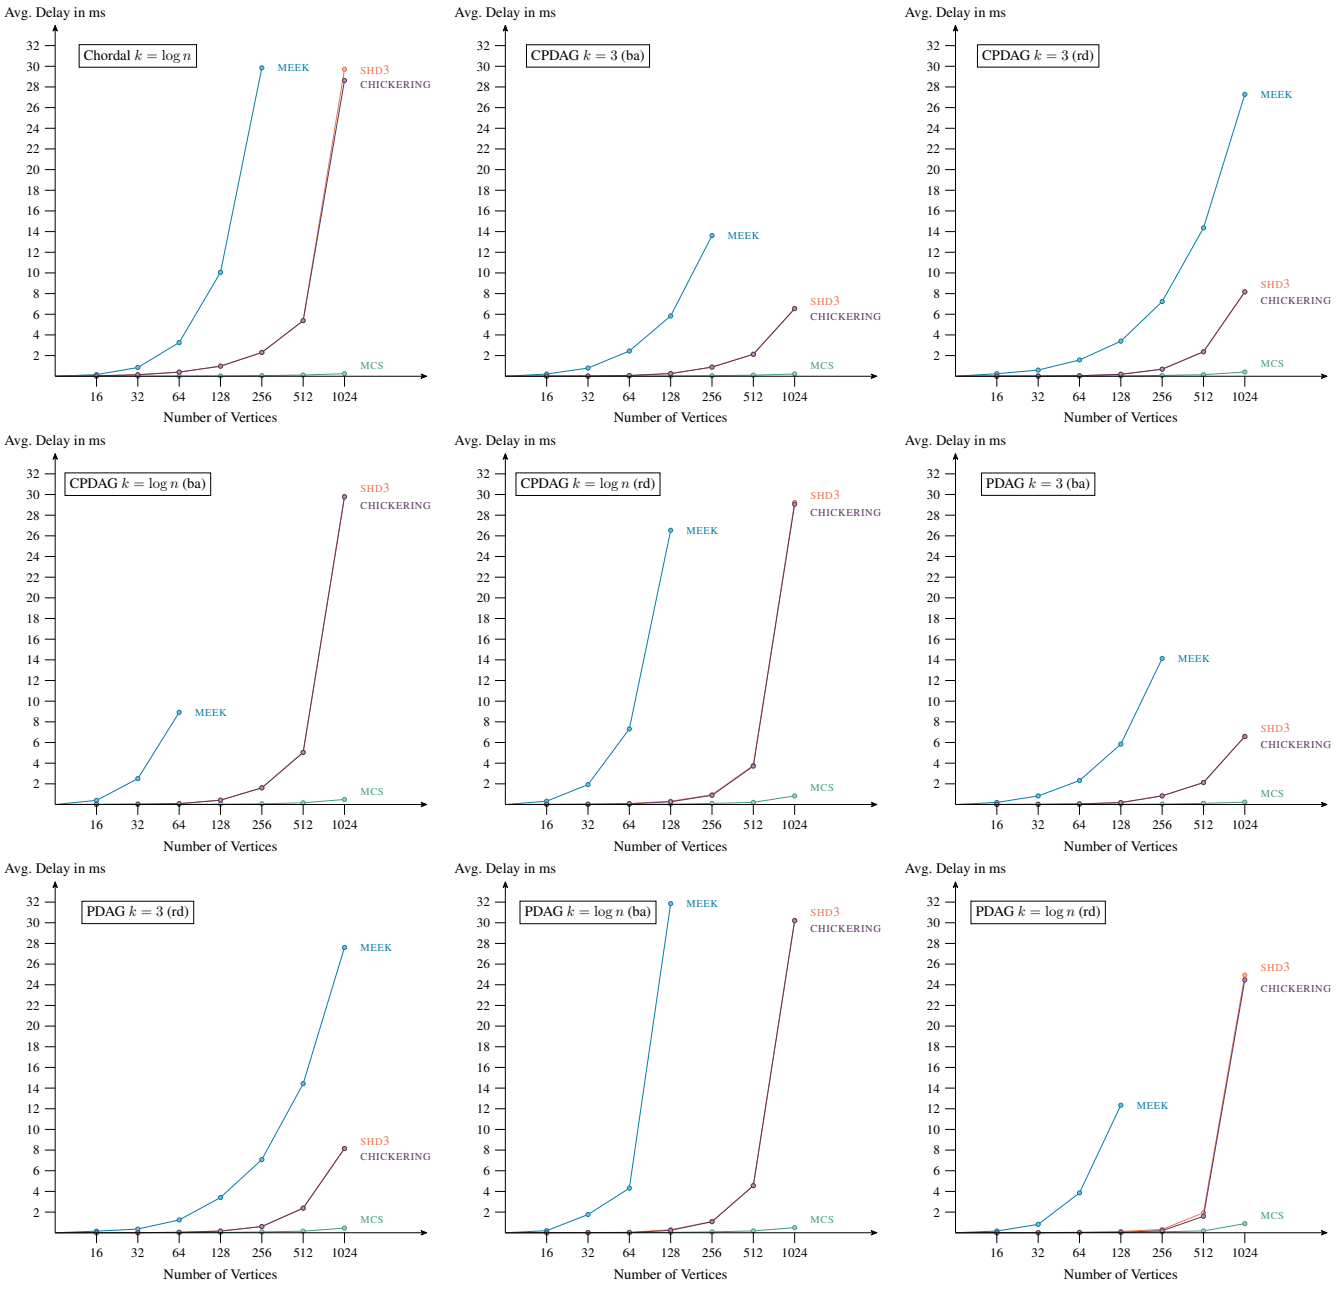

Figure 7: Comparison of the average delay between two output DAGs for the algorithms MEEK-ENUM, CHICKERING-ENUM, MCS-ENUM, and SHD3-ENUM. The graphs are generated as described in the text with density parameter k and the generation methods (ba), for scale-free graphs, and (rd), for uniformly random edge insertions.

Algorithm	Scenario	$k = 1$	$k = 2$	$k = 3$	$k = 5$	$k = 7$	$k = 10$
MEEK-ENUM	UCCG	69.56	91.26	96.69	98.86	99.25	99.68
	CPDAG	63.37	90.52	97.27	99.74	99.95	99.98
	PDAG	63.61	90.79	97.35	99.75	99.95	99.98
	All	69.35	91.24	96.71	98.89	99.27	99.69
CHICKERING-ENUM	UCCG	89.83	99.73	99.86	99.87	99.88	99.89
	CPDAG	86.04	99.77	99.79	99.80	99.81	99.83
	PDAG	81.58	99.73	99.79	99.81	99.83	99.85
	All	89.02	99.73	99.85	99.86	99.87	99.88
MCS-ENUM	UCCG	91.98	99.68	99.77	99.80	99.81	99.81
	CPDAG	76.87	98.93	99.60	99.66	99.66	99.66
	PDAG	74.29	98.42	98.49	99.63	99.63	99.63
	All	89.96	99.56	99.75	99.78	99.79	99.79
SHD3-ENUM	UCCG	77.39	99.37	99.86	99.87	99.88	99.89
	CPDAG	67.81	98.90	99.79	99.80	99.81	99.83
	PDAG	63.35	98.43	99.79	99.81	99.83	99.85
	All	75.73	99.27	99.85	99.86	99.87	99.88

Table 1: The proportion of delays (in percent) that are less or equal to multiples of the mean delay $\bar{d}$, i. e., the percentage of delays d for which $d \leq k \cdot \bar{d}$ holds. We evaluated the delays for each algorithm in a specific scenario and for all scenarios combined.