Inferring Observed Structure For Dynamic Graphs with Unobserved Variables

Isaac Davis, Erich Kummerfeld, David Danks, and Sergey Plis

November 23, 2015

Technical Report No. CMU-PHIL-193

Philosophy

Methodology

Logic

Carnegie Mellon

Pittsburgh, Pennsylvania 15213

Inferring Observed Structure For Dynamic Graphs with Unobserved Variables

Isaac Davis (CMU), Erich Kummerfeld (University of Pittsburgh),
David Danks (CMU),
and Sergey Plis (Mind Research Network, University of New Mexico)

Abstract. This paper provides theoretical analysis on the impact of unmeasured variables on causal structure learning from time series data. In particular, we provide an algorithm for transforming a dynamic causal graph over a causally sufficient set of variables into the corresponding dynamic graph over a subset of the causally sufficient variables.

1. Background and Terminology

1.1. Introduction. Many scientific investigations aim to learn the causal structure underlying some observed dynamic or time series phenomena. However, we are often unable to measure all causally relevant variables, and many structure learning algorithms fail when this causal sufficiency assumption is violated. For example, suppose the true causal structure is $V_2^t \leftarrow V_1^{t-1} \rightarrow V_3^t$, where superscripts denote time. If V_1 is unmeasured, the structure learning algorithm may infer that V_2 is a contemporaneous cause of V_3, or visa versa, when in fact neither causes the other. In this paper, we use compressed graphs to efficiently represent causal structure in time series data, and provide and algorithm for transforming a compressed graph $\mathcal{G}_V$, where V is a causally sufficient set of variables, into the corresponding compressed graph $\mathcal{G}_O$, where $O \subseteq V$. This characterizes the impact of unobserved variables when the "full" graph is known.

1.2. Assumptions. A dynamic causal graphical model consists of a graph $\mathcal{G}_V$ over nodes for random variables in V at the current time-step t, as well as every preceding time-step that contains a direct cause of a random variable at t. There is additionally a probability distribution $P(V)$ over the nodes in V, although for the purpose of this paper we will abstract away from parametric considerations. We assume that V is causally sufficient- that is, any variable which is a direct cause of a variable in V must also be in V. We impose two further assumptions on the independence relationships between variables in V: the Markov assumption, which asserts that a variable $v \in V$ is conditionally independent of its non-descendant variables conditional on its direct ancestors, and the faithfulness assumption, which asserts that the only independence relations are those arising from the Markov assumption. For theoretical simplicity, we also assume that the $\mathcal{G}_V$ has Markov order 1. That is, if V_i^{t-n} is a direct cause of V_j^t, then $n = 1$. We also assume that there is no contemporaneous causation. That is, we cannot have $V_i^t \to V_j^t$: any direct causal effect must occur after one time-step.

1.3. Mathematical Notation. A compressed graph is a compact and efficient representation of a dynamic causal graphical model. Rather than having distinct nodes for variables at each time-step, we simply write $V_i \to V_j$, and associate a set of integers to this arrow corresponding to the time-lags between variable V_i and its effect on variable V_j. Formally, a compressed graph $\mathcal{G}_W$ consists of a set W of nodes, a function $d^W : W \times W \to \mathcal{P}(\mathbb{N})$, and a function $b^W : W \times W \to \mathcal{P}(\mathbb{Z})$. The function d^W determines directed edges between nodes in W, while b^W determines bi-directed edges. If X and Y are nodes in W, then $d^W(X,Y) = S$ means that for all $n \in S$, the underlying graph contains the edge $X^{t-n} \to Y^t$. Here we will call S the set of "directed edge-lags between A and B." We will allow $d^W(X,Y) = \emptyset$, the interpretation being that the underlying graph contains no directed edges from X to Y. For notation considerations, we will write $d^W(X,Y)$ as d^W_{xy} when no confusion is possible. For bi-directed edges, we impose an arbitrary ordering $<_W$ over the nodes in W. For $X, Y \in W$, if $b^W_{XY} = S$, this means that for all $n \in S$, $X^{t-n} \leftrightarrow Y^t$ if $X <_W Y$ and $Y^{t-n} \leftrightarrow X^t$ if $Y <_W X$. For the rest of this paper, whenever bi-directed edges are being considered, we will assume that $A <_W B$ in the imposed ordering.

We will need one more bit of mathematical notation before we proceed. Suppose $S, T \subseteq \mathbb{Z}$. Then we define

$$
\begin{aligned}
S &\oplus T = \{s + t | s \in S, t \in T\} \\
S &\ominus T = \{s - t | s \in S, t \in T\}
\end{aligned}
$$

Note that $\oplus$ is a commutative operation on sets of integers while $\ominus$ is not commutative. Lastly, for a finite set $S = \{s_1, \dots, s_n\} \subset \mathbb{Z}$, we define

$$\langle S \rangle = \{\alpha_1 s_1 + \ldots + \alpha_n s_n | \alpha_i \in \mathbb{Z}^{\geq 0}\}$$

1.4. Formal Statement of The Forward Inference Problem. Suppose $\mathcal{G}_W$ is a causally sufficient, fully-sampled, compressed graph. Let $O \subset W$ and $U = W \backslash O$. We interpret O as the observed variables in W and U as the unobserved variables in W. The forward inference problem is, given $\mathcal{G}_W$, what is $\mathcal{G}_O$? That is, what is the fully-sampled compressed graph over O if we do not observe U?

The solution to the forward inference problem hinges on the following fact: Edge lag indices represent the time delays between directly correlated values in the observed variables connected by that edge. Since the values are directly correlated, these correlations can not be derivable from the values of other observed variables. This means that only unobserved variables may exist on the path between the two correlated variables (potentially none, in the case where the edge lag index is the singleton $\{1\}$ which corresponds to the observed edge between them). As such, by enumerating all the correlation- inducing paths between members of O which cross only through nodes in U, all possible ways in which the removal of U from W can induce edges in $\mathcal{G}_O$ are accounted for.

2. Simultaneous Update Equations

The goal for this section is, for any two nodes $A, B \in O$, to construct the sets d^O_{AB} and b^O_{AB}. That is, to compute the sets of directed and bi-directed edge-lags between A and B in the restricted graph $\mathcal{G}_O$. To do so, we must first introduce notation to represent the "length" of paths between A and B. To this end, we define a Cycle-Restricted Directed (CRD) path to be a directed path with no repeated cycles. Suppose the CRD-path π consists of edges $X_1 \to X_2 \to \ldots \to X_n$, where $X_1 = A$, $X_n = B$, $X_i \in U$ for all $1 < i < n$. We will define a function $L()$ which maps each path π to a set of positive integers corresponding to the set of all path lengths along this path. To do so, first let $cy(\pi)$ be the set of all sub-cycles of π. Because cycles can be iterated an arbitrary number of times along a path, each cycle will contribute infinitely many lags to the path's edge-lag set.

Now, define the set $ti(\pi)$ to be

$$ti(\pi) = \bigoplus_{i=2}^{n} d^W_{X_{i-1}X_i}$$

Next, define the set $ci(\pi)$ to be

$$ci(\pi) = \bigoplus_{p \in cy(\pi)} \langle ti(p) \rangle$$

where $\langle S \rangle$ is the set generated by elements of S. Here, $ti(\pi)$ is the set of all lengths obtainable on the path π without repeating any cycles, and $ci(\pi)$ is the additional set of lengths obtained by iterating the cycles in π. The length-set of π is defined to be

$$L(\pi) = ti(\pi) \oplus ci(\pi)$$

Now, for nodes $A, B \in O$, define $\prod_d^{W,U}(A, B)$ to be the set of all CRD-paths from A to B with intermediate nodes in U. Taking the union of the length-sets of all paths in $\prod_d^{W,U}(A, B)$ gives us the full set of directed edge-lags between A and B when all nodes in U are removed. Therefore the forward inference rule for directed edges is

$$d^O_{AB} = \bigcup_{\pi \in \prod_d^{W,U}(A,B)} L(\pi) \tag{2.1}$$

The right hand side of the above equation constitutes a complete search of all the directed paths from A to B with intermediate nodes in U. The edge lag set for $A \to B$ consists of all and only those path-lengths.

Bi-directed edges can be dealt with in a similar manner, but first we must introduce a new kind of trek. Heterogeneous treks (h-treks) are a generalization of CRD treks that account for the possible existence of bi-directed edges in $\mathcal{G}_W$, since $A \leftarrow Y \leftrightarrow X \to B$ is a legitimate (i.e. correlation-inducing) trek between A and B. An h-trek is an ordered triple $\langle \tau_1, \tau_2, e \rangle$ where τ_1 is a CRD path from some node X_1 to B, τ_2 is a CRD path from some node X_2 to A, and in the event that $X_1 \neq X_2$, e is a bi-directed edge between X_1 and X_2. At most one of τ_1, τ_2, and e can be empty. If e is empty, then $X_1 = X_2$. If τ_1 is empty, then e is a bi-directed edge between the head node of τ_2 and B. If τ_2 is empty, then e is a bi-directed edge between the head node of τ_1 and A.

We can use h-treks to employ the same strategy for bi-directed edges as we did with directed edges. Let $\prod_b^{W,U}(A, B)$ be the set of all h-treks between A and B with intermediate nodes in U. For any h-trek $\tau \in \prod_b^{W,U}(A, B)$, with $A <_W B$, define $lags(\tau) = [L(\tau_1) \ominus L(\tau_2)] \oplus ind(e)$, where $L(\emptyset) = \{0\}$ by convention and $ind(e)$ is defined as

$$ind(e) = \begin{cases} b_{AB}^W & \text{if } e \neq \emptyset \\ \{0\} & \text{otherwise} \end{cases}$$

We can now state the forward inference rule for bi-directed edges as follows:

$$b_{AB}^O = \bigcup_{\tau \in \prod_b^{W,U}(A,B)} lags(\tau) \tag{2.2}$$

Similarly to the rule for directed edges, the right hand side of this equation constitutes a complete search of all correlation-inducing h-treks between A and B with intermediate nodes in U.

With these two forward inference rules established, the forward algorithm for computing $\mathcal{G}_O$ from $\mathcal{G}_W$ is simple: for each pair of nodes $A, B \in O$, add an edge $A \to B$ with edge-lag set d_{AB}^O as computed above, then add a bi-directed edge $A \leftrightarrow B$ with edge-lag set b_{AB}^O as computed above. Then delete all nodes in U. However, it is generally not going to be feasible to implement this algorithm, as enumerating all the paths in $\prod_d^{W,U}(A,B)$ and $\prod_b^{W,U}(A,B)$ can be very computationally intensive if U is large. In practice, it is much faster to iteratively remove one variable at a time, since $\prod_d^{W,\{X\}}(A,B)$ and $\prod_b^{W,\{X\}}(A,B)$ are small enough that they can be expressed in a simple closed form and calculated quickly. For this reason, we provide a more comphrehensive set of forward equations for removing a single variable.

3. Single Variable Update Equations

Removal of a single variable can induce both edges between distinct nodes as well as self-loops (both directed and bi-directed) between a node and itself. The strategy is to consider all of the paths which will disappear when a node R is removed, and how those paths will or won't affect the edges between the remaining variables. When only a single variable is being removed, the number of paths to consider is small, so we can simply work through them one by one. This process results in a method that is certainly correct when starting from an un-marginalized compressed graph, but it must be proven that the process is correct when starting from a marginalized compressed graph as well. This will be proven in a later section. For the purpose of this section, R will denote the node that is being removed.

3.1. Directed Self-Loops. We first consider directed self-loops induced by the removal of R. In order for the removal of R to induce a directed self-loop on A, there must be a directed path that starts at A, passes only through R (though possibly more than once), and terminates again at A. The edges that could be involved in such a path are $A \to R$, $R \to A$, and $R \to R$, and the length-set of this family of paths in $d_{AR}^W \oplus d_{RA}^W \oplus \langle d_{RR}^W \rangle$. Therefore when R is removed it will add all elements of $d_{AR}^W \oplus d_{RA}^W \oplus \langle d_{RR}^W \rangle$ to the edge-lag set d_{AA}^W. We can therefore write the update rule for directed self-loops on A as

$$d_{AA}^{W \setminus \{R\}} = d_{AA}^W \cup \left(d_{AR}^W \oplus d_{RA}^W \oplus \langle d_{RR}^W \rangle \right)$$

3.2. Bi-Directed Self-Loops. For bi-directed self-loops on A, there must be a correlation-inducing trek starting and ending at A whose unique root variable is also unmeasured. Since each path in the trek must terminate at A, each path must involve either $R \to A$ or $A \leftrightarrow R$. However $A \leftrightarrow R$ cannot occur at both ends of the trek, as this would force a collider to occur along the trek. Therefore the only two types of treks that can induce bi-directed self-loops on A when R is removed are those that have $R \to A$ at both ends and those that have $A \leftrightarrow R$ on one end and $R \to A$ at the other.

When both ends of the trek terminate with $R \to A$, both the directed and bi-directed self-loops on R may participate in the trek. The directed self-loop may participate on either side of the trek, but the bi-directed self-loop can participate only once or not at all. The contribution of each trek to the edge-lag set for the induced bi-directed self-loop is equal to the difference between the path lengths on either side of the root variable. Therefore the bi-directed self-loop on A induced by these treks has edge-lag set equal to

$$[(d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \ominus [(d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \oplus (\{0\} \cup b^W_{RR})$$

In the case that the trek terminates with $R \to A$ on one end and $R \leftrightarrow A$ on the other, the bi-directed self-loop on R cannot participate on either side of the trek, as this would create a collider at R. The edge-lag set induced by these treks is equal to

$$[(d^W_{RA} \oplus \langle d^W_{RR} \rangle) \ominus (b^W_{RA} \oplus \langle d^W_{RR} \rangle)] \cup [(b^W_{RA} \oplus \langle d^W_{RR} \rangle) \ominus (d^W_{RA} \oplus \langle d^W_{RR} \rangle)]$$

Putting these sets together gives us the full update rule for bi-directed self-loops on A as:

$$\begin{aligned} b^{W\setminus\{R\}}_{AA} &= b^W_{AA} \cup [(d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \ominus [(d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \oplus (\{0\} \cup b^W_{RR}) \\ &\quad \cup [(d^W_{RA} \oplus \langle d^W_{RR} \rangle) \ominus (b^W_{RA} \oplus \langle d^W_{RR} \rangle)] \cup [(b^W_{RA} \oplus \langle d^W_{RR} \rangle) \ominus (d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \end{aligned}$$

3.3. Directed Edges Between Distinct Nodes. The removal of R can induce a directed edge from node A to node B if and only if there is a directed path from A to B through R. The only edges that can participate in such a path are $A \to R$, $R \to B$, and $R \to R$. The length-set of this family of paths is $d^W_{AR} \oplus d^W_{RB} \oplus \langle d^W_{RR} \rangle$. Therefore the update rule for directed edges from A to B is

$$d^{W\setminus\{R\}}_{AB} = d^W_{AB} \cup d^W_{AR} \oplus d^W_{RB} \oplus \langle d^W_{RR} \rangle$$

3.4. Bi-Directed Edges Between Distinct Nodes. There are three types of treks which will induce bi-directed edges between A and B when R is removed. These are: $A \leftarrow R \to B$, $A \leftarrow R \leftrightarrow B$, and $A \leftrightarrow R \to B$. For the first type of trek, both a directed self-loop at R and a bi-directed self-loop at R could participate in the trek. The edge-lag set induced by this trek is therefore (note that we assume throughout the paper that $A <_W B$):

$$(d^W_{RB} \oplus \langle d^W_{RR} \rangle) \ominus (d^W_{RA} \oplus \langle d^W_{RR} \rangle) \oplus (\{0\} \cup b^W_{RR})$$

For treks of the form $A \leftarrow R \leftrightarrow B$, the bi-directed self-loop on R cannot participate, as it would form a collider on R. Furthermore, the directed self-loop can only occur on one side of the trek, in particular the side with the directed edge (in this case the side towards A). Furthermore, we must know whether R comes before or after B in the ordering on W. To this end, let $I(X, Y)$ be an indicator function which returns 1 if $X <_W Y$ and -1 if $Y <_W X$. Then the bi-directed edge-lag set induced by treks of this form is

$$I(R, B) b^W_{RB} \ominus d^W_{RA} \ominus \langle d^W_{RR} \rangle$$

The analysis for treks of the form $A \leftrightarrow R \rightarrow B$ is symmetrical to the above analysis, and the edge-lag set induced by treks of this form is equal to

$$I(A,R)b^W_{RA} \ominus d^W_{RB} \ominus \langle d^W_{RR} \rangle$$

Putting this all together gives us the full update rule for bi-directed edges between distinct nodes:

$$\begin{aligned} b^{W\backslash\{R\}}_{AB} = b^W_{AB} &\cup [(d^W_{RB} \oplus \langle d^W_{RR} \rangle) \ominus (d^W_{RA} \oplus \langle d^W_{RR} \rangle) \oplus (\{0\} \cup b^W_{RR})] \cup [I(R,B)b^W_{RB} \ominus d^W_{RA} \ominus \langle d^W_{RR} \rangle] \\ &\cup [I(A,R)b^W_{RA} \ominus d^W_{RB} \ominus \langle d^W_{RR} \rangle] \end{aligned}$$

3.5. Summary of Update Equations. The full set of forward update equations for removal of a single node R are as follows:

$$\begin{aligned} d^{W\backslash\{R\}}_{AA} &= d^W_{AA} \cup \left(d^W_{AR} \oplus d^W_{RA} \oplus \langle d^W_{RR} \rangle\right) \\ b^{W\backslash\{R\}}_{AA} &= b^W_{AA} \cup [(d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \ominus [(d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \oplus (\{0\} \cup b^W_{RR}) \\ &\quad \cup [(d^W_{RA} \oplus \langle d^W_{RR} \rangle) \ominus (b^W_{RA} \oplus \langle d^W_{RR} \rangle)] \cup [(b^W_{RA} \oplus \langle d^W_{RR} \rangle) \ominus (d^W_{RA} \oplus \langle d^W_{RR} \rangle)] \\ d^{W\backslash\{R\}}_{AB} &= d^W_{AB} \cup d^W_{AR} \oplus d^W_{RB} \oplus \langle d^W_{RR} \rangle \\ b^{W\backslash\{R\}}_{AB} &= b^W_{AB} \cup [(d^W_{RB} \oplus \langle d^W_{RR} \rangle) \ominus (d^W_{RA} \oplus \langle d^W_{RR} \rangle) \oplus (\{0\} \cup b^W_{RR})] \cup [I(R,B)b^W_{RB} \ominus d^W_{RA} \ominus \langle d^W_{RR} \rangle] \\ &\quad \cup [I(A,R)b^W_{RA} \ominus d^W_{RB} \ominus \langle d^W_{RR} \rangle] \end{aligned}$$

With these established, we can now construct the iterative forward inference algorithm which removes one node at a time. Note that we still need to prove that deleting nodes iteratively produces the same graph as deleting nodes simultaneously, which we prove in the last section.

4. The Sequential Forward Inference Algorithm

For the purpose of this section, assume we have a compressed graph $\mathcal{G}_W$ over nodes W. We provide an algorithm for computing $\mathcal{G}_{W\backslash\{X\}}$ for some $X \in W$. For the general case, we will have some set $O \subset W$ of observed variables, with $U = W\backslash O$ representing the unobserved variables, and we wish to obtain $\mathcal{G}_O$ from $\mathcal{G}_W$. To do so, we apply this algorithm iteratively for each node $X \in U$ in any order.

4.1. Forward Inference on Directed Edges. Algorithm for constructing directed edges resulting from deletion of X from $\mathcal{G}_W$:

(1) Identify all nodes with edges into and out of X: Let $I_x = \{A \neq X \in W | d^W_{AX} \neq \emptyset\}$ and let $O_x = \{B \neq X \in W | d^W_{XB} \neq \emptyset\}$.
(2) For any pair of nodes $A \in I_X$ and $B \in O_X$, set

$$d^{W\backslash\{X\}}_{AB} = d^W_{AB} \cup \left[d^W_{AX} \oplus d^W_{XB} \oplus \langle d^W_{XX} \rangle\right]$$

(3) For any pair of nodes C, D not identified in step 2, set

$$d^{W\backslash\{X\}}_{CD} = d^W_{CD}$$

4.2. Forward Inference on Bidirected Edges. Algorithm for constructing bi-directed edges resulting from deletion of X from $\mathcal{G}_W$:

(1) Constructing bi-directed self-loops resulting from deletion of X:
   (a) Identify all nodes A with a directed edge from X: let $O_x = \{A \neq X \in V | d^W_{XA} \neq \emptyset\}$
   (b) For each $A \in O_x$, two cases follow:
      - Case 1: $b^W_{XA} = \emptyset$. Let
$$T_{XA} = \left[(d^W_{XA} \oplus \langle d^W_{XX} \rangle) \ominus (d^W_{XA} \oplus \langle d^W_{XX} \rangle)\right] \oplus (\{0\} \cup b^W_{XX})$$
        Then set
$$b^{W \setminus \{X\}}_{AA} = b^W_{AA} \cup T_{XA}$$
      - Case 2: $b^W_{XA} \neq \emptyset$. Let
$$S_{XA} = \left[(d^W_{XA} \oplus \langle d^W_{XX} \rangle) \ominus (b^W_{XA} \oplus \langle d^W_{XX} \rangle)\right] \cup \left[(b^W_{XA} \oplus \langle d^W_{XX} \rangle) \ominus (d^W_{XA} \oplus \langle d^W_{XX} \rangle)\right]$$
        and let
$$T_{XA} = \left[(d^W_{XA} \oplus \langle d^W_{XX} \rangle) \ominus (d^W_{XA} \oplus \langle d^W_{XX} \rangle)\right] \oplus (\{0\} \cup b^W_{XX})$$
        Then set
$$b^{W \setminus \{X\}}_{AA} = b^W_{AA} \cup T_{XA} \cup S_{XA}$$
   (c) For each $B \notin O_x$, set
$$b^{W \setminus \{X\}}_{BB} = b^W_{BB}$$

(2) Constructing bi-directed edges between distinct nodes resulting from deletion of X:
   (a) Identify all nodes with a directed edge from X and all nodes with a bi-directed edge with X: Let $O_x = \{A \neq X \in W | d^W_{XA} \neq \emptyset\}$, and let $B_x = \{A \neq X \in W | b^W_{XA} \neq \emptyset\}$.
   (b) In order for the bi-directed edge-lag sets to be meaningful, we must have some arbitrary ordering $<_W$ over nodes in W. For all ordered pairs of nodes $(A, B) \in (O_x \cup B_x) \times (O_x \cup B_x)$ such that $A <_W B$, set
$$L_{AB} = \begin{cases} (d^W_{XB} \oplus \langle d^W_{XX} \rangle) \ominus (d_{XA} \oplus \langle d^W_{XX} \rangle) \oplus (\{0\} \cup b^W_{XX}) & \text{if } A, B \in O_x \\ \emptyset & \text{otherwise} \end{cases}$$
      Next, let $I(S, T)$ be an indicator function that returns 1 when $S <_W T$ and -1 when $T <_W S$. Then let
$$M_{AB} = \begin{cases} I(X, B) b^W_{XA} \ominus d^W_{XB} \ominus \langle d^W_{XX} \rangle & \text{if } A \in O_x, B \in B_x \\ \emptyset & \text{otherwise} \end{cases}$$
      Lastly, let
$$N_{AB} = \begin{cases} I(A, X) b^W_{XA} \oplus d^W_{XB} \oplus \langle d^W_{XX} \rangle & \text{if } A \in B_x, B \in O_x \\ \emptyset & \text{otherwise} \end{cases}$$
      Finally, set
$$b^{W \setminus \{X\}}_{AB} = b^W_{AB} \cup L_{AB} \cup M_{AB} \cup N_{AB}$$

(c) For any pair of nodes (C, D) not identified in the previous steps, set

$$b_{AB}^{W\setminus\{X\}} = b_{AB}^{W}$$

5. Equivalence of Sequential and Simultaneous Variable Removal

Recall that the sequential forward algorithm is necessary because the simultaneous forward algorithm may be computationally intractable for large U. However, in order to use the sequential algorithm, we must prove that it yields the same graph as the simultaneous algorithm. To this end, suppose $I \subset W$, and let $Mar(\mathcal{G}_W, I)$ be the function that computes $\mathcal{G}_{W\setminus I}$ from $\mathcal{G}_W$ and I via simultaneous edge removal. We will prove the following theorem:

THEOREM 5.1. $Mar(Mar(\mathcal{G}_W, I), \{X\}) = Mar(\mathcal{G}_W, I \cup \{X\})$

Note that the left hand side of the above equation is the graph obtained by deleting $\{X\}$ from the graph obtained by deleting I from $\mathcal{G}_W$, and the right hand side of the equation is the graph obtained by deleting I and $\{X\}$ from $\mathcal{G}_W$ simultaneously. Since I, $\{X\}$, and $\mathcal{G}_W$ are arbitrary, this equation entails that iterated removal of singletons will produce the same graph as simultaneous removal of these variables. Note also that a compressed graph is defined entirely in terms of its node set W and its directed and bi-directed edge functions d^W and b^W. Since both graphs clearly have the same node set $W\setminus I \cup \{X\}$, proving the above theorem is equivalent to proving that $Mar(Mar(\mathcal{G}_W, I)\{X\})$ has the same directed and bi-directed edge lag sets as $Mar(\mathcal{G}_W, I \cup \{X\})$. We will separate the proof into two sub-proofs, one for directed edges and one for bi-directed edges.

5.1. Proof for Directed Edges.

We begin by constructing the following sets. Let:

$$\begin{aligned}
\Lambda_i^1 &= \bigcup_{\pi \in \prod_d^{W,I}(A,B)} L(\pi) \\
\Lambda_i^2 &= \bigcup_{\pi \in \prod_d^{W\setminus I,\{X\}}(A,B)} L(\pi) \\
\Lambda_b &= \bigcup_{\pi \in \prod_d^{W,I\cup\{X\}}(A,B)} L(\pi)
\end{aligned}$$

Note that Λ_i^1 is the set of lengths of all paths from A to B in $\mathcal{G}_W$ with intermediate nodes in I, Λ_i^2 is the set of lengths of all paths from A to B in $\mathcal{G}_{W\setminus I}$ with intermediate nodes in $\{X\}$, and Λ_b is the set of lengths of all paths from A to B in $\mathcal{G}_W$ with intermediate nodes in $I \cup \{X\}$. Therefore, Theorem 5.1 for directed edges is equivalent to $\Lambda_i^1 \cup \Lambda_i^2 = \Lambda_b$, which is the claim that we shall prove.

To do so, we first split up Λ_b into two components which correspond to Λ_i^1 and Λ_i^2. Let $\prod_d^{W,I}(A, B)$ be the set of all paths in $\prod_d^{W,I\cup\{X\}}(A, B)$ which do not cross X, and let $\prod_d^{W,I\cup\{X\}}(A, X, B)$ be the set of all paths in $\prod_d^{W,I\cup\{X\}}(A, B)$ which cross X at least once. Define

$$\begin{aligned}
\Lambda_b^1 &= \bigcup_{\pi \in \prod_d^{W,I}(A,B)} L(\pi) \\
\Lambda_b^2 &= \bigcup_{\pi \in \prod_d^{W,I\cup\{X\}}(A,X,B)} L(\pi)
\end{aligned}$$

Clearly we have $\Lambda_b = \Lambda_b^1 \cup \Lambda_b^2$. Furthermore, Λ_b^1 consists of all and only directed paths from A to B with intermediate nodes in I, which is exactly what Λ_i^1 is defined to be, so we have $\Lambda_i^1 = \Lambda_b^1$. Therefore we can reduce the claim $\Lambda_i^1 \cup \Lambda_i^2 = \Lambda_b$ to the claim $\Lambda_i^2 = \Lambda_b^2$.

To this end, first observe that $\prod_d^{W \setminus I, \{X\}}(A, B)$ is very small, consisting only of the following two paths:

- $\pi_a = A \to X \to B$
- $\pi_b = A \to X \to X \to B$

Therefore we can fully characterize Λ_i^2 as $L(\pi_a) \cup L(\pi_b)$. Furthermore, we know that

$$\begin{aligned} L(\pi_a) &= d_{AX}^{W \setminus I} \oplus d_{BX}^{W \setminus I} \\ L(\pi_b) &= d_{AX}^{W \setminus I} \oplus d_{BX}^{W \setminus I} \oplus \langle d_{XX}^{W \setminus I} \rangle \end{aligned}$$

Therefore we have

(5.1)

$$\Lambda_i^2 = L(\pi_a) \cup L(\pi_b) = [d_{AX}^{W \setminus I} \oplus d_{BX}^{W \setminus I}] \cup [d_{AX}^{W \setminus I} \oplus d_{BX}^{W \setminus I} \oplus \langle d_{XX}^{W \setminus I} \rangle] = d_{AX}^{W \setminus I} \oplus d_{BX}^{W \setminus I} \oplus \langle d_{XX}^{W \setminus I} \rangle$$

To calculate the indices $d_{AX}^{W \setminus I}$, $d_{XX}^{W \setminus I}$, and $d_{XB}^{W \setminus I}$ in terms of $\mathcal{G}_W$, we apply the simultaneous update equation and expand the $L()$ terms into their $ti()$ and $ci()$ components:

$$\begin{aligned} d_{AX}^{W \setminus I} &= \bigcup_{\pi \in \prod_d^{W,I}(A,X)} ti(\pi) \oplus ci(\pi) \\ d_{XX}^{W \setminus I} &= \bigcup_{\pi \in \prod_d^{W,I}(X,X)} ti(\pi) \oplus ci(\pi) \\ d_{XB}^{W \setminus I} &= \bigcup_{\pi \in \prod_d^{W,I}(X,B)} ti(\pi) \oplus ci(\pi) \end{aligned}$$

Replacing the appropriate terms in equation (5.2) yields the set Λ_i^2 expressed solely in terms of $\mathcal{G}_W$:

(5.2)

$$\Lambda_i^2 = \bigcup_{\pi \in \prod_d^{W,I}(A,X)} [ti(\pi) \oplus ci(\pi)] \oplus \bigcup_{\pi \in \prod_d^{W,I}(X,B)} [ti(\pi) \oplus ci(\pi)] \oplus \langle \bigcup_{\pi \in \prod_d^{W,I}(X,X)} [ti(\pi) \oplus ci(\pi)] \rangle$$

Finally, it will be convenient later to have the right side of the this equation expressed in only 3 terms. This result is nearly identical to equation (5.2) but will simplify later steps:

$$\begin{aligned} \Lambda_i^2(A,X) &= d_{AX}^{W \setminus I} = \bigcup_{\pi \in \prod_d^{W,I}(A,X)} ti(\pi) \oplus ci(\pi) \\ \Lambda_i^2(X,B) &= d_{XB}^{W \setminus I} = \bigcup_{\pi \in \prod_d^{W,I}(X,B)} ti(\pi) \oplus ci(\pi) \\ \Lambda_i^2(X,X) &= \langle d_{XX}^{W \setminus I} \rangle = \langle \bigcup_{\pi \in \prod_d^{W,I}(X,X)} ti(\pi) \oplus ci(\pi) \rangle \end{aligned}$$

Giving us the compact form:

$$\Lambda_i^2 = \Lambda_i^2(A,X) \oplus \Lambda_i^2(X,B) \oplus \Lambda_i^2(X,X) \tag{5.3}$$

The next step of the proof is to show that the elements of Λ_b^2 can be similarly described in terms of the summation of 3 components, which correspond usefully to those of Λ_i^2.

Expanding the $L()$ term in the definition of Λ_b^2 yields:

$$\Lambda_b^2 = \bigcup_{\pi \in \prod_d^{W, I \cup \{X\}}(A,X,B)} ti(\pi) \oplus ci(\pi)$$

In order to get a formulation of Λ_b^2 appropriate for our purposes, we need to break its $ti()$ and $ci()$ components down even further. Before we proceed, however, we introduce a bit of notation for breaking up paths into smaller sub-paths. Suppose $\pi \in \prod_d^{W,I\cup\{X\}}(A, X, B)$. We will partition π into three subpaths as follows: let $\pi(A, X)$ be the shortest subpath of π from A to X, $\pi(X, B)$ be the shortest subpath of π from X to B, and let $\pi(X)$ be the longest subpath of π from X to itself (if π only crosses X once, then $\pi(X)$ is just the node X). Clearly π is the concatenation of $\pi(A, X)$, $\pi(X)$, and $\pi(X, B)$. Furthermore, $\pi(A, X) \in \prod_d^{W,I}(A, X)$, and $\pi(X, B) \in \prod_d^{W,I}(X, B)$. The last thing to note is that $\pi(X)$ is a sequence of cycles from X to itself. Let $\pi^1(X), \ldots, \pi^n(X)$ be the sequence of cycles such that each $\pi^i(X)$ only crosses X at the head and tail, and $\pi(X)$ is the concatenation of $\pi^1(X), \ldots, \pi^n(X)$ in that order. Since π is a CRD path, we know that $\pi^i(X) \neq \pi^j(X)$ for $i \neq j$. Furthermore, for each i, we have $\pi^i(X) \in \prod_d^{W,I}(X, X)$, which means that $L(\pi^i(X)) \subseteq d_{XX}^{W\setminus I}$. Since $\pi(X)$ is the concatenation of the $\pi_i(X)$ sub-cycles, this means that $L(\pi(X)) \subseteq \langle d_{XX}^{W\setminus I}\rangle$.

With this established, we can rewrite

$$ti(\pi) = ti(\pi(A, X)) \oplus ti(\pi(X)) \oplus ti(\pi(X, B))$$

Since $\pi(X)$ is by definition a subcycle of π, we have $\pi(X) \in cy(\pi)$, therefore $ti(\pi(X)) \subseteq ci(\pi)$, which entails that $ti(\pi(X)) \oplus ci(\pi) = ci(\pi)$. We can therefore rewrite the $ti(\pi) \oplus ci(\pi)$ term as

$$ti(\pi(A, X)) \oplus ti(\pi(X, B)) \oplus ci(\pi)$$

Now, recall that $ci(\pi) = \bigoplus_{p\in cy(\pi)}\langle ti(p)\rangle$. We will break $cy(\pi)$ into three parts. Let $p \in cy_{AX}(\pi)$ iff p is a subcycle of $\pi(A, X)$, and similarly for $p \in cy_{XB}(\pi)$. Let $cy_*(\pi) = cy(\pi)\setminus(cy_{AX}(\pi) \cup cy_{XB}(\pi))$. Then $cy_*(\pi)$ consist of all and only those subcycles of π that either cross X at least once or are subcycles of $\pi(X)$. We have $cy(\pi) = cy_{AX}(\pi) \cup cy_{XB}(\pi) \cup cy_*(\pi)$, so we can break the $\bigoplus$ in $ci(\pi)$ into three smaller operations as follows:

$$\bigoplus_{p\in cy(\pi)} \langle ti(p)\rangle = \bigoplus_{p\in cy_{AX}(\pi)} \langle ti(p)\rangle \oplus \bigoplus_{p\in cy_{BX}(\pi)} \langle ti(p)\rangle \oplus \bigoplus_{p\in cy_*(\pi)} \langle ti(p)\rangle \tag{5.4}$$

For convenience let the components of the righthand side of equation (5.5) be $ci_{AX}(\pi)$, $ci_{XB}(\pi)$, and $ci_*(\pi)$ respectively. This gives us the more convenient form:

$$ci(\pi) = ci_{AX}(\pi) \oplus ci_{XB}(\pi) \oplus ci_*(\pi) \tag{5.5}$$

We can now express Λ_b^2 in the following form:

$$\Lambda_b^2 = \bigcup_{\pi\in\prod_d^{W,I\cup\{X\}}(A,X,B)} ti(\pi(A, X)) \oplus ti(\pi(X, B)) \oplus (ci_{AX}(\pi) \oplus ci_{XB}(\pi) \oplus ci_*(\pi)) \tag{5.6}$$

We can use the following terms to simplify the expression:

$$\begin{aligned} L_{A,X}(\pi) &= ti(\pi(A,X)) \oplus ci_{AX}(\pi) \\ L_{B,X}(\pi) &= ti(\pi(X,B)) \oplus ci_{XB}(\pi) \\ L_*(\pi) &= ci_*(\pi) \end{aligned}$$

We can now express Λ_b^2 in 3 terms that correspond usefully to the terms used to express Λ_i^2:

$$\Lambda_b^2 = \bigcup_{\pi \in \prod_d^{W,I\cup\{X\}}(A,X,B)} L_{A,X}(\pi) \oplus L_{X,B}(\pi) \oplus L_*(\pi) \tag{5.7}$$

With this established, we will now use the forms of equations (5.4) and (5.8) to prove that $\Lambda_b^2 \subseteq \Lambda_i^2$ and $\Lambda_i^2 \subseteq \Lambda_i^2$.

LEMMA 5.2. $\Lambda_b^2 \subseteq \Lambda_i^2$

PROOF. Suppose $\pi \in \prod_d^{W,I\cup\{X\}}(A,X,B)$. It follows immediately that $\pi(A,X) \in \prod_d^{W,I\cup\{X\}}(A,X)$ and $\pi(X,B) \in \prod_d^{W,I\cup\{X\}}(X,B)$. Therefore by their definitions, $L_{A,X}(\pi) \subseteq \Lambda_i^2(A,X)$ and $L_{X,B}(\pi) \subseteq \Lambda_i^2(X,B)$. By definition, any $\theta \in cy_*(\pi)$ is either in $\prod_d^{W,I\cup\{X\}}(X,X)$ or is a subcycle of a path in $\prod_d^{W,I\cup\{X\}}(X,X)$. Thus $\bigcup_{\theta\in cy_*(\pi)} L(\theta) \subseteq \bigcup_{\theta \in \prod_d^{W,I\cup\{X\}}(X,X)} L(\theta)$. It follows from the definitions of $\Lambda_i^2(X,X)$ and $L_*(\pi)$ that $L_*(\pi) \subseteq \Lambda_i^2(X,X)$. From the above containment relations and the properties of $\oplus$, it follows directly that for arbitrary $\pi \in \prod_d^{W,I\cup\{X\}}(A,X,B)$ that

$$L_{A,X}(\pi) \oplus L_{X,B}(\pi) \oplus L_*(\pi) \subseteq \Lambda_i^2(A,X) \oplus \Lambda_i^2(X,B) \oplus \Lambda_i^2(X,X)$$

Since π was chosen arbitrarily, this proves the lemma. □

In order to prove that $\Lambda_i^2 \subseteq \Lambda_b^2$, we will prove a slightly more general lemma that will be useful in completing the proof for bi-directed edges. The lemma is as follows:

LEMMA 5.3. Suppose $I \subseteq W$ and $D,E,F \in W$. Let $\pi_{DE} \in \prod_d^{W,I}(D,E)$, $\pi_{EE}^1,\ldots,\pi_{EE}^n \in \prod_d^{W,I}(E,E)$, and $\pi_{EF} \in \prod_d^{W,I}(E,F)$. Let π_{EE} be the concatenation of $\pi_{EE}^1,\ldots,\pi_{EE}^n$. Then for all $y \in L(\pi_{DE}) \oplus \langle L(\pi_{EE})\rangle \oplus L(\pi_{EF})$, there exists some $\pi \in \prod_d^{W,I\cup\{E\}}(D,E,F)$ such that $y \in L(\pi)$.

PROOF. Suppose $I \subseteq W$ and $D,E,F \in W$. Let $\pi_{DE} \in \prod_d^{W,I}(D,E)$, $\pi_{EE}^1,\ldots,\pi_{EE}^n \in \prod_d^{W,I}(E,E)$, and $\pi_{EF} \in \prod_d^{W,I}(E,F)$. Let π_{EE} be the concatenation of $\pi_{EE}^1,\ldots,\pi_{EE}^n$. Let $y \in L(\pi_{DE}) \oplus \langle L(\pi_{EE})\rangle \oplus L(\pi_{EF})$. Suppose by contradiction that there is no $\pi \in \prod_d^{W,I\cup\{E\}}(D,E,F)$ such that $y \in L(\pi)$. Let Z be the set of variables touched by π_{DE}, π_{EE}, and π_{EF}, and let $\mathcal{H}_Z$ be the compressed graph over Z consisting of all and only those edges included in π_{DE}, π_{EE}, and π_{EF}. Let π_c be the concatenation of π_{DE}, π_{EE}, and π_{EF}. Then π_c is a directed path in $\mathcal{H}_Z$ from D to F, and $y \in L(\pi_c)$. If all variables in Z aside from D and F are simultaneously marginalized out, then y must be in $d_{DF}^{\{D,F\}}$. However, by assumption there is no $\pi \in \prod_d^{W,I\cup\{E\}}(D,E,F)$ such that $y \in L(\pi)$. But $Z \subseteq I \cup \{D,E,F\}$, so $Z\backslash\{D,F\} \subseteq I \cup \{E\}$, so this entails that there is no $\pi \in \prod_d^{Z,Z\backslash\{D,F\}}(D,E,F)$ such that $y \in L(\pi)$. As such, applying equation (2.1) to marginalize $\mathcal{H}_Z$ would not be guaranteed to find a $d_{DE}^{\{D,F\}}$ such that $y \in d_{DE}^{\{D,F\}}$, thus contradicting the correctness of equation (2.1). Since we have derived a contradiction, there must exist some $\pi \in \prod_d^{W,I\cup\{E\}}(D,E,F)$ such that $y \in \pi$, thus proving the lemma. $\square$

Using this lemma, we can now prove that $\Lambda_i^2 \subseteq \Lambda_b^2$.

Lemma 5.4. $\Lambda_i^2 \subseteq \Lambda_b^2$

Proof. Suppose $y \in \Lambda_i^2$. Then there must exist $\pi_{AX} \in \prod_d^{W,I}(A,X)$, $\pi_{XB} \in \prod_d^{W,I}(X,B)$ and $\pi_{XX}^1, \ldots, \pi_{XX}^n \in \prod_d^{W,I}(X,X)$ such that $y \in L(\pi_{AX}) \oplus \langle L(\pi_{XX})\rangle \oplus L(\pi_{XB})$, where π_{XX} is the concatenation of $\pi_{XX}^1, \ldots \pi_{XX}^n$. By lemma 5.10, this entails that there exists some $\pi \in \prod_d^{W,I\cup\{X\}}(A,X,B)$ such that $y \in L(\pi)$. Since

$$\Lambda_b^2 = \bigcup_{\pi \in \prod_d^{W,I\cup\{X\}}(A,X,B)} L(\pi)$$

we have $L(\pi) \subseteq \Lambda_b^2$. Therefore $y \in \Lambda_b^2$, thus proving that $\Lambda_i^2 \subseteq \Lambda_b^2$. $\square$

This completes the proof that $\Lambda_i^2 = \Lambda_b^2$, thus proving the theorem for directed edges. We now move on to the proof for bi-directed edges.

5.2. Proof for Bi-Directed Edges.

Define Γ_i^1, Γ_i^2, and Γ_b as

$$\begin{aligned}
\Gamma_i^1 &= \bigcup_{\tau \in \prod_b^{W,I}(A,B)} lags(\tau) \\
\Gamma_i^2 &= \bigcup_{\tau \in \prod_b^{W\setminus I,\{X\}}(A,B)} lags(\tau) \\
\Gamma_b &= \bigcup_{\tau \in \prod_b^{W,I\cup\{X\}}(A,B)} lags(\tau)
\end{aligned}$$

Note that Γ_i^1 is the set of bi-directed edge lags induced by removing I from W, Γ_i^2 is the set of bi-directed edge lags induced by removing $\{X\}$ from $W\setminus I$, and Γ_b is the set of bi-directed edge lags induced by removing $I \cup \{X\}$ from W. Therefore, for bi-directed edges, the above theorem is equivalent to

$$\Gamma_i^1 \cup \Gamma_i^2 = \Gamma_b$$

We can partition Γ_b into two subsets as follows: by definition, $\prod_b^{W,I\cup\{X\}}(A,B)$ is the set of all correlation-inducing heterogeneous treks between A and B with intermediate nodes (i.e. all non-tail nodes) in $I \cup \{X\}$. So, let $\prod_b^{W,I}(A,B)$ be the set of all h-treks in $\prod_b^{W,I\cup\{X\}}(A,B)$ that do not cross X, and let $\prod_b^{W,I\cup\{X\}}(A,X,B)$ be the set of all h-treks in $\prod_b^{W,I\cup\{X\}}(A,B)$ that cross X at least once. Clearly we have

$$\prod_b^{W,I\cup\{X\}}(A,B) = \prod_b^{W,I}(A,B) \cup \prod_b^{W,I\cup\{X\}}(A,X,B)$$

and

$$\prod_b^{W,I}(A,B) \cap \prod_b^{W,I\cup\{X\}}(A,X,B) = \emptyset$$

Now, define

$$\Gamma_b^1 = \bigcup_{\tau \in \prod_b^{W,I}(A,B)} lags(\tau)$$
$$\Gamma_b^2 = \bigcup_{\tau \in \prod_b^{W,I\cup\{X\}}(A,X,B)} lags(\tau)$$

Clearly $\Gamma_b = \Gamma_b^1 \cup \Gamma_b^2$. Furthermore, Γ_b^1 is identical to Γ_i^1, as both unions range over correlation inducing h-treks between A and B with intermediate nodes in I. Therefore the proof of the theorem is reduced to showing that $\Gamma_i^2 = \Gamma_b^2$.

To this end, we first consider Γ_i^2. In particular, we examine what treks are in the set $\prod_b^{W\setminus I,\{X\}}(A,B)$. This is the set of correlation-inducing h-treks in $\mathcal{G}_{W\setminus I}$ between A and B with intermediate nodes in $\{X\}$, and consists of only three types of treks. These are

(1) $\tau_a : A \leftarrow X \rightarrow B$
(2) $\tau_b : A \leftarrow X \leftrightarrow B$
(3) $\tau_c : A \leftrightarrow X \rightarrow B$

For τ_a, both a directed self-loop at X and a bi-directed self-loop at X could participate in the trek. Therefore the edge-lag index of this h-trek is

$$lags(\tau_a) = (d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \ominus (d_{XA}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \oplus (\{0\} \cup b_{XX}^{W\setminus I})$$

For τ_b, a bi-directed self-loop at X cannot participate in the trek, as it would create a collider, and a directed self-loop at X could only participate on one side of the trek, namely the side leading to A (otherwise there would be a collider at X). Furthermore, the edge-lag index of this trek depends on whether $X <_W B$ or $B <_W X$. Let $I(S,T)$ be an indicator function that returns 1 if $S <_W T$ and -1 if $T <_W S$. Then the edge-lag index of this h-trek is

$$lags(\tau_b) = I(X,B)b_{XB}^{W\setminus I} \ominus d_{XA}^{W\setminus I} \ominus \langle d_{XX}^{W\setminus I}\rangle$$

For τ_c, the analysis is similar to that of τ_b, but with A taking the place of B. The edge-lag index for this h-trek is

$$lags(\tau_c) = I(A,X)b_{XA}^{W\setminus I} \oplus d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle$$

This exhausts all correlation-inducing h-treks between A and B with intermediate nodes in $\{X\}$, so we can fully characterize Γ_i^2 as

$$\Gamma_i^2 = lags(\tau_a) \cup lags(\tau_b) \cup lags(\tau_c)$$

Using this decomposition, we now show that $\Gamma_i^2 = \Gamma_b^2$ by proving that $\Gamma_b^2 \subseteq \Gamma_i^2$ and $\Gamma_i^2 \subseteq \Gamma_b^2$.

First, we handle the $\Gamma_b^2 \subseteq \Gamma_i^2$ direction. So, suppose $y \in \Gamma_b^2$. Then there exists an h-trek $\tau \in \prod_b^{W,I\cup\{X\}}(A,X,B)$ such that $y \in lags(\tau)$. By definition, an h-trek τ is an ordered triple $\langle \tau_1, \tau_2, e\rangle$, where τ_1 is a directed path from some $H_1 \in I$ to B, τ_2 is a directed path from some $H_2 \in I$ to A, and in the case that $H_1 \neq H_2$, e is a bi-directed edge between H_1 and H_2. Furthermore, because $\tau \in \prod_b^{W,I\cup\{X\}}(A,X,B)$, at least one of τ_1, τ_2 crosses X at least once. Four cases follow:

(1) $e = \emptyset$: in this case, $H_1 = H_2 = H$. Then $lags(\tau) = L(\tau_1) \ominus L(\tau_2)$. Suppose first that τ_1 crosses X and τ_2 does not. Let $\tau_1(H,X)$, $\tau_1(X)$, $\tau_1(X,B)$ be the partition of τ_1 introduced above. Clearly $L(\tau_1) = L(\tau_1(H,X)) \oplus$

$L(\tau_1(X)) \oplus L(\tau_1(X,B))$, so we have $y \in (L(\tau_1(H,X)) \oplus L(\tau_1(X)) \oplus L(\tau_1(X,B))) \ominus L(\tau_2)$. Furthermore, $\tau_1(X,B) \in \prod_d^{W,I}(X,B)$, therefore $L(\tau_1(X,B)) \subseteq d_{XB}^{W\setminus I}$. Note also that $\tau_1(H,X)$ and τ_2 form a correlation-inducing trek between X and A, so $\langle \tau_1(H,X), \tau_2, \emptyset \rangle \in \prod_b^{W,I}(X,A)$. Therefore we have $L(\tau_1(H,X)) \ominus L(\tau_2) \subseteq b_{XA}^{W\setminus I}$ if $A <_W X$ and $L(\tau_2) \ominus L(\tau_1(H,X)) \subseteq b_{XA}^{W\setminus I}$ if $X <_W A$. We can express this more compactly with $L(\tau_1(H,X)) \ominus L(\tau_2) \subseteq I(A,X) b_{XA}^{W\setminus I}$. Lastly, $L(\tau_1(X)) \subseteq \langle d_{XX}^{W\setminus I} \rangle$. Putting this all together gives us the following:

$$\begin{aligned}
y &\in [L(\tau_1(H,X)) \oplus L(\tau_1(X)) \oplus L(\tau_1(X,B))] \ominus L(\tau_2) \\
&= L(\tau_1(X,B)) \oplus L(\tau_1(X)) \oplus (L(\tau_1(H,X) \ominus L(\tau_2)) \\
&\subseteq d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I} \rangle \oplus I(X,A) b_{XA}^{W\setminus I} \\
&= lags(\tau_c) \subseteq \Gamma_i^2
\end{aligned}$$

Therefore in the case that τ_1 crosses X but τ_2 does not, we have shown that $y \in \Gamma_i^2$. The case in which τ_2 crosses X but τ_1 does not is analogous, and we end up showing that $y \in lags(\tau_b)$, rather than $lags(\tau_c)$. Lastly, we consider the case in which both τ_1 and τ_2 cross X. In this case, we split up both τ_1 and τ_2 into three parts each, $L(\tau_1) = L(\tau_1(H,X)) \oplus L(\tau_1(X)) \oplus L(\tau_1(X,B))$, and $L(\tau_2) = L(\tau_2(H,X)) \oplus L(\tau_2(X)) \oplus L(\tau_2(X,A))$. Furthermore, we have

$$\begin{aligned}
L(\tau_1(X,B)) &\subseteq d_{XB}^{W\setminus I} \\
L(\tau_2(X,A)) &\subseteq d_{XA}^{W\setminus I} \\
L(\tau_1(X)) &\subseteq \langle d_{XX}^{W\setminus I} \rangle \\
L(\tau_2(X)) &\subseteq \langle d_{XX}^{W\setminus I} \rangle \\
\langle \tau_1(H,X), \tau_2(H,X), \emptyset \rangle &\in \prod_b^{W,I}(X,X) \Rightarrow L(\tau_1(H,X)) \ominus L(\tau_2(H,X)) \subseteq b_{XX}^{W\setminus I}
\end{aligned}$$

Putting this all together gives us

$$\begin{aligned}
y &\in L(\tau_1) \ominus L(\tau_2) \\
&= [L(\tau_1(H,X)) \oplus L(\tau_1(X)) \oplus L(\tau_1(X,B))] \ominus [L(\tau_2(H,X)) \oplus L(\tau_2(X)) \oplus L(\tau_2(X,A))] \\
&= [L(\tau_1(X,B) \oplus L(\tau_1(X))] \ominus [L(\tau_2(X,A)) \oplus L(\tau_2(X))] \oplus [L(\tau_1(H,X)) \ominus L(\tau_2(H,X))] \\
&\subseteq (d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I} \rangle) \ominus (d_{XA}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I} \rangle) \oplus (b_{XX}^{W\setminus I}) \\
&\subseteq (d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I} \rangle) \ominus (d_{XA}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I} \rangle) \oplus (\{0\} \cup b_{XX}^{W\setminus I}) \\
&= lags(\tau_a) \subseteq \Gamma_i^2
\end{aligned}$$

Therefore in the case that both τ_1 and τ_2 cross X, $y \in \Gamma_i^2$. This completes the proof for the case that $e = \emptyset$.

(2) $\tau_1 = \emptyset$: in this case, τ_2 must pass through X, and e must be a bi-directed edge between H and B, where H is the head node of τ_2. Since τ_2 is a directed path to A and $A <_W B$, the lag set of this trek is

$$lags(\tau) = (\{0\} \ominus L(\tau_2)) \oplus ind(e) = ind(e) \ominus L(\tau_2)$$

So, suppose $y \in lags(\tau) = ind(e) \ominus L(\tau_2)$. We can split τ_2 into three components $\tau_2(H,X)$, $\tau_2(X)$, and $\tau_2(X,A)$. Then $\langle \emptyset, \tau_2(H,X), e \rangle$ is a correlation inducing h-trek between X and B with intermediate nodes in I, therefore

$$\langle \emptyset, \tau_2(H,X), e\rangle \in \prod_b^{W,I}(X,B)$$

This implies that the lag set of this h-trek, which is $[I(X,B)ind(e)] \ominus [I(X,B)L(\tau_2(H,X))]$, is a subset of $b_{XB}^{W\setminus I}$. We can write this more compactly as

$$ind(e) \ominus L(\tau_2(H,X)) \subseteq I(X,B)b_{XB}^{W\setminus I}$$

Furthermore, we have $L(\tau_2(X)) \subseteq \langle d_{XX}^{W\setminus I}\rangle$, and $\tau_2(X,A) \in \prod_d^{W,I}(X,A)$, so $L(\tau_2(X,A)) \subseteq d_{XA}^{W\setminus I}$. Putting this all together gives us

$$\begin{aligned} y &\in ind(e) \ominus L(\tau_2) \\ &= ind(e) \ominus [L(\tau_2(H,X) \oplus L(\tau_2(X,A) \oplus L(\tau_2(X))] \\ &= [ind(e) \ominus L(\tau_2(H,X))] \ominus L(\tau_2(X,A)) \ominus L(\tau_2(X)) \\ &\subseteq \left[I(X,B)b_{XB}^{W\setminus I}\right] \ominus d_{XA}^{W\setminus I} \ominus \langle d_{XX}^{W\setminus I}\rangle \\ &= lags(\tau_b) \subseteq \Gamma_i^2 \end{aligned}$$

Therefore in the case that $\tau_1 = \emptyset$, we have shown that $y \in \Gamma_i^2$.

(3) $\tau_2 = \emptyset$: this case is quite similar to the previous case. In this case, τ_1 is a directed path from H to B, and e is a bi-directed edge between H and A. The lag set of τ is therefore

$$lags(\tau) = (L(\tau_1) \ominus \{0\}) \oplus ind(e) = L(\tau_1) \oplus ind(e)$$

As in the previous case, we separate τ_1 into three components $\tau_1(H,X)$, $\tau_1(X)$ and $\tau_1(X,B)$. Then $\langle \tau_1(H,X), \emptyset, e\rangle \in \prod_b^{W,I}(X,A)$, and the lag set of this h-trek is $L(\tau_1(H,X)) \oplus ind(e)$, so we have

$$L(\tau_1(H,X)) \oplus ind(e) \subseteq I(A,X)b_{XA}^{W\setminus I}$$

Furthermore, $L(\tau_1(X)) \subseteq \langle d_{XX}^{W\setminus I}\rangle$ and $L(\tau_1(X,B)) \subseteq d_{XB}^{W\setminus I}$. Putting this all together gives us

$$\begin{aligned} y &\in L(\tau_1) \oplus ind(e) \\ &= [L(\tau_1(H,X)) \oplus L(\tau_1(X,B)) \oplus L(\tau_1(X))] \oplus ind(e) \\ &= [L(\tau_1(H,X) \oplus ind(e)] \oplus L(\tau_1(X,B)) \oplus L(\tau_1(X)) \\ &\subseteq I(A,X)b_{XA}^{W\setminus I} \oplus d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle \\ &= lags(t_c) \subseteq \Gamma_i^2 \end{aligned}$$

So in the case that $\tau_2 = \emptyset$, we have shown that $y \in \Gamma_i^2$

(4) : $\tau_1, \tau_2, e \neq \emptyset$: in this case, τ_1 is a directed path from H_1 to B, τ_2 is a directed path from H_2 to A, and e is a bi-directed edge between H_1 and H_2, with $H_1 \neq H_2$. The lag set of this h-trek is

$$lags(\tau) = [L(\tau_1) \ominus L(\tau_2)] \oplus ind(e)$$

Furthermore, at least one of τ_1, τ_2 must pass through X. Suppose first that only τ_1, but not τ_2, passes through X. Then $\langle \tau_1(H_1,X), \tau_2, e\rangle$ is a correlation-inducing h-trek between X and A, so

$$[L(\tau_1(H_1,X)) \ominus L(\tau_2)] \oplus ind(e) \subseteq I(A,X)b_{XA}^{W\setminus I}$$

Furthermore $L(\tau_1(X)) \subseteq \langle d_{XX}^{W\setminus I}\rangle$, and $L(\tau_1(X,B)) \subseteq d_{XB}^{W\setminus I}$. Putting this all together gives us

$$\begin{aligned}
y &\in [L(\tau_1) \ominus L(\tau_2)] \oplus ind(e)\\
&= [(L(\tau_1(H_1,X)) \oplus L(\tau_1(X)) \oplus L(\tau_1(X,B))) \ominus L(\tau_2)] \oplus ind(e)\\
&= [(L(\tau_1(H_1,X)) \ominus L(\tau_2)) \oplus ind(e)] \oplus L(\tau_1(X,B)) \oplus L(\tau_1(X))\\
&\subseteq I(A,X)b_{XA}^{W\setminus I} \oplus d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle\\
&= lags(\tau_c) \subseteq \Gamma_i^2
\end{aligned}$$

For the case in which τ_2, but not τ_1, crosses X, the analysis is analogous, and we show that in this case $y \in lags(\tau_b) \subseteq \Gamma_i^2$. Lastly, we consider the case in which both τ_1 and τ_2 cross X. In this case, $\langle \tau_1(H_1,X), \tau_2(H_2,X), e\rangle$ is a correlation-inducing h-trek between X and itself. Therefore the lag set of this h-trek, $[L(\tau_1(H_1,X)) \ominus L(\tau_2(H_2,X))] \oplus ind(e)$, is a subset of $b_{XX}^{W\setminus I}$. Furthermore, we have

$$\begin{aligned}
L(\tau_1(X)) &\subseteq \langle d_{XX}^{W\setminus I}\rangle\\
L(\tau_2(X)) &\subseteq \langle d_{XX}^{W\setminus I}\rangle\\
L(\tau_1(X,B)) &\subseteq d_{XB}^{W\setminus I}\\
L(\tau_2(X,A) &\subseteq d_{XA}^{W\setminus I}
\end{aligned}$$

Putting this all together gives us

$$\begin{aligned}
y &\in [L(\tau_1) \ominus L(\tau_2)] \oplus ind(e)\\
&= (L(\tau_1(H_1,X)) \oplus L(\tau_1(X)) \oplus L(\tau_1(X,B))) \ominus (L(\tau_2(H_2,X)) \oplus L(\tau_2(X,A)) \oplus L(\tau_2(X))) \oplus ind(e)\\
&= (L(\tau_1(X,B)) \oplus L(\tau_1(X))) \ominus (L(\tau_2(X,A)) \oplus L(\tau_2(X))) \oplus (L(\tau_1(H_1,X)) \ominus L(\tau_2(H_2,X))) \oplus ind(e)\\
&\subseteq (d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \ominus (d_{XA}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \oplus b_{XX}^{W\setminus I}\\
&\subseteq (d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \ominus (d_{XA}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \oplus (\{0\} \cup b_{XX}^{W\setminus I})\\
&= lags(\tau_a) \subseteq \Gamma_i^2
\end{aligned}$$

Therefore we have shown in all possible cases that if $y \in \Gamma_b^2$, we must also have $y \in \Gamma_i^2$. This completes the proof that $\Gamma_b^2 \subseteq \Gamma_i^2$.

Now we show that $\Gamma_i^2 \subseteq \Gamma_b^2$. To this end, suppose $y \in \Gamma_i^2$. Since $\Gamma_i^2 = lags(\tau_a) \cup lags(\tau_b) \cup lags(\tau_c)$, three cases follow:

(1) Case 1: $y \in lags(\tau_a)$. In this case, we have

$$y \in (d_{XB}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \ominus (d_{XA}^{W\setminus I} \oplus \langle d_{XX}^{W\setminus I}\rangle) \oplus b_{XX}^{W\setminus I}$$

Therefore there must exist

$$\begin{aligned}\pi_{XA} &\in \prod_d^{W,I}(X,A)\\ \pi_{XB} &\in \prod_d^{W,I}(X,B)\\ \pi^1_{XX},\ldots,\pi^n_{XX} &\in \prod_d^{W,I}(X,X)\\ \tau_{XX} &\in \prod_b^{W,I}(X,X)\end{aligned}$$

such that

$$y \in (L(\pi_{XB}) \oplus \langle L(\pi_{XX})\rangle) \ominus (L(\pi_{XA}) \oplus \langle L(\pi_{XX})\rangle) \oplus lags(\tau_{XX})$$

where π_{XX} is the concatenation of $\pi^1_{XX},\ldots,\pi^n_{XX}$. By definition, τ_{XX} consists of a path $\pi_{H_1X} \in \prod_d^{W,I}(H_1,X)$, a path $\pi_{H_2X} \in \prod_d^{W,I}(H_2,X)$, and if $H_1 \neq H_2$, a bi-directed edge e between H_1 and H_2. Then we have $lags(\tau_{XX}) = L(\pi_{H_2X}) \ominus L(\pi_{H_1X}) \oplus ind(e)$. Therefore

$$y \in (L(\pi_{XB}) \oplus \langle L(\pi_{XX})\rangle) \ominus (L(\pi_{XA}) \oplus \langle L(\pi_{XX})\rangle) \oplus (L(\pi_{H_2X}) \ominus L(\pi_{H_1X}) \oplus ind(e))$$

By definition, there must exist integers $a \in L(\pi_{XA})$, $b \in L(\pi_{XB})$, $c_1, c_2 \in \langle L(\pi_{XX})\rangle$, $h_1 \in L(\pi_{H_1X})$, $h_2 \in L(\pi_{H_2X})$, and $\epsilon \in ind(e)$, such that

$$y = (b + c_1) - (a + c_2) + (h_2 - h_1 + \epsilon) = (h_2 + c_1 + b) - (h_1 + c_1 + a) + \epsilon$$

Note that $h_2 + c_2 + b \in L(\pi_{H_2X}) \oplus \langle L(\pi_{XX})\rangle \oplus L(\pi_{XB})$, therefore by lemma 5.10 there exists some $\pi_2 \in \prod_d^{W,I\cup\{X\}}(H_2,X,B)$ such that $h_2 + c_2 + b \in L(\pi_2)$. Similarly, $h_1 + c_1 + a \in L(\pi_{H_1X}) \oplus \langle L(\pi_{XX})\rangle \oplus L(\pi_{XA})$, so there exists $\pi_1 \in \prod_d^{W,I\cup\{X\}}(H_1,X,A)$ such that $h_1 + c_1 + a \in L(\pi_1)$. Furthermore, $\tau = \langle \pi_1, \pi_2, e\rangle \in \prod_b^{W,I\cup\{X\}}(A,X,B)$, and $lags(\tau) = L(\pi_2) \ominus L(\pi_1) \oplus ind(e)$. Putting this all together gives us

$$y = (h_2 + c_2 + b) - (h_1 + c_1 + a) + \epsilon \in (L(\pi_2) \ominus L(\pi_1)) \oplus ind(e) = lags(\tau) \subseteq \Gamma_b^2$$

Therefore in the case that $y \in lags(\tau_a)$, we have $y \in \Gamma_b^2$.

(2) Case 2: $y \in lags(\tau_b)$. In this case, we have

$$y \in I(X,B) b_{XB}^{W\setminus I} \ominus d_{XA}^{W\setminus I} \ominus \langle d_{XX}^{W\setminus I}\rangle$$

Therefore there must exist

$$\begin{aligned}\tau_{XB} &\in \prod_b^{W,I}(X,B)\\ \pi_{XA} &\in \prod_d^{W,I}(X,A)\\ \pi^1_{XX},\ldots,\pi^n_{XX} &\in \prod_d^{W,I}(X,X)\end{aligned}$$

such that

$$y \in I(X,B)lags(\tau_{XB}) \ominus L(\pi_{XA}) \ominus \langle L(\pi_{XX})\rangle$$

where π_{XX} is the concatenation of $\pi^1_{XX}, \ldots, \pi^n_{XX}$. By definition, τ_{XB} consists of some path $\pi_{H_1X} \in \prod_d^{W,I}(H_1, X)$, some path $\pi_{H_2X} \in \prod_d^{W,I}(H_2, B)$, and if $H_1 \neq H_2$, some bi-directed edge e between H_1 and H_2. So we have

$$lags(\tau_{XB}) = I(X,B)(L(\pi_{H_2B}) \ominus L(\pi_{H_1X})) \oplus ind(e)$$

Therefore we can write

$$y \in (L(\pi_{H_2B}) \ominus L(\pi_{H_1X})) \oplus ind(e) \ominus L(\pi_{XA}) \ominus \langle L(\pi_{XX})\rangle$$

By definition, this means that there exist integers $h_1 \in L(\pi_{H_1X})$, $h_2 \in L(\pi_{H_2B})$, $a \in L(\pi_{XA})$, $c \in \langle L(\pi_{XX})\rangle$, and $\epsilon \in ind(e)$ such that

$$y = (h_2 - h_1) + \epsilon - a - c = h_2 - (h_1 + c + a) + \epsilon$$

Note that $h_1 + c + a \in L(\pi_{H_1X}) \oplus \langle L(\pi_{XX})\rangle \oplus L(\pi_{XA})$, so by lemma 5.10 there exists some $\pi \in \prod_d^{W,I\cup\{X\}}(H_1, X, A)$ such that $h_1 + c + a \in L(\pi)$. Furthermore, $\tau = \langle \pi_{H_2B}, \pi, e\rangle \in \prod_b^{W,I\cup\{X\}}(A, X, B)$, so $lags(\tau) = L(\pi_{H_2B}) \ominus L(\pi) \oplus ind(e) \subseteq \Gamma_b^2$. Putting this all together gives us

$$y = h_2 - (h_1 + c + a) + \epsilon \in L(\pi_{H_2B}) \ominus L(\pi) \oplus ind(e) = lags(\tau) \subseteq \Gamma_b^2$$

Therefore in the case that $y \in lags(\tau_b)$, we have $y \in \Gamma_b^2$.

(3) Case 3: $y \in lags(\tau_c)$. This case is analogous to the previous case, but with τ_{XA} taking the place of τ_{XB}. This exhausts all possible cases, therefore we have shown that for any $y \in \Gamma_i^2$, we must have $y \in \Gamma_b^2$. Thus we have $\Gamma_i^2 \subseteq \Gamma_b^2$, thereby completing the proof that $\Gamma_i^2 = \Gamma_b^2$.

Therefore we have proven that the graphs obtained by $Mar(Mar(\mathcal{G}_W, I), \{X\})$ and $Mar(\mathcal{G}_W, I \cup \{X\})$ have the same directed and bi-directed edge-lag sets, and clearly they have the same node-sets. Therefore the two graphs are equivalent, and we have proven that simultaneous edge removal yields the same graph as iterated edge removal.