

Worlds of Events

Deduction with Partial Knowledge about Causality

Seyed Hossein Haeri

Peter Van Roy

Carlos Baquero

Université catholique de Louvain, Belgium

Universidade do Minho, Portugal

Christopher Meiklejohn

Université catholique de Louvain, Belgium

Interactions between internet users are mediated by their devices and the common support infrastructure in data centres. Keeping track of causality amongst actions that take place in this distributed system is key to provide a seamless interaction where effects follow causes. Tracking causality in large scale interactions is difficult due to the cost of keeping large quantities of metadata; even more challenging when dealing with resource-limited devices. In this paper, we focus on keeping partial knowledge on causality and address deduction from that knowledge.

We provide the first proof-theoretic causality modelling for distributed partial knowledge. We prove computability and consistency results. We also prove that the partial knowledge gives rise to a weaker model than classical causality. We provide rules for offline deduction about causality and refute some related folklore. We define two notions of forward and backward bisimilarity between devices, using which we prove two important results. Namely, no matter the order of addition/removal, two devices deduce similarly about causality so long as: (1) the same causal information is fed to both. (2) they start bisimilar and erase the same causal information. Thanks to our establishment of forward and backward bisimilarity, respectively, proofs of the latter two results work by simple induction on length.

1 Introduction

Causality [16, 20] is an essential for our perception of the physical world, and of our relations to other entities. If one puts a cup on a table, and looks back at it, one expects it to be there. One also expects to get a reply to one's postcards, **after** they were sent, and not before.

Given the fault-tolerance and high availability expected of internet-based services today, distributed algorithms have become ubiquitous. One duty of these algorithms is to order events totally across multiple replicas of a service. This total order is required to ensure computation determinism; given the requirement of having multiple replicas appear as a single system, each replica must implement a state machine which observes the same events in the same order [19]. However, because of the amount of coordination required, a total order in the entire distributed system is not always feasible while maintaining availability [12].

Given the intractability of a total order, techniques that favour a partial order based on causality are explored for they express user's **intent**. For a key-value store, one's writes may, e.g., be directed to one replica, and, subsequent reads served from a replica which has not yet received those writes. If we consider the canonical example of an access control list for viewing photos [17], one would expect that a write operation removing Eve from having access to Alice's photos prior to Alice uploading a photo she did not want Eve to see, would be observed in this order by Eve when performing read operations.

However, tracking causality can be very expensive in terms of metadata size; more so when interactions amongst many distinct entities are targeted [7]. Devising scalable solutions to causality tracking is

a demanding problem [17, 21] to the extent that some solutions even accept to lose causal information by pruning metadata [8]. Be it because available resources are limited (say to an edge device) or because a replica (say in a data centre) is temporarily out-of-sync, only a partial view of the system causality might be available. There is not much study on dealing with that partiality of knowledge, however. In this paper, we address that problem via a proof-theoretic modelling for partial causality knowledge of distributed systems. Partiality is not a loss in our model: What is not stored might be deducible – acting in favour of metadata size reduction.

Contributions of this paper are as follows:

1. We model distributed causality such that the holistic system and the partial causal knowledge of a device are categorically distinct (Definitions 2.1 and 2.3).
2. We offer rules for deducing causality when a device is online (Definition 3.1) and prove its computability (Theorem 3.4) and consistency (Theorem 3.5).
3. We show that deduction of causality with partial knowledge is strictly less accurate than the holistic causal knowledge (Lemma 3.7) and that the deductions of different devices do not conflict (Corollary 3.8).
4. We offer rules for a device to deduce causal information independent of new causal data from outside, e.g., when offline (Definition 4.1) and prove its consistency with the online rules (Lemmata 4.2 and 4.3). We also prove a related folklore wrong (Lemma 4.4) using the latter machinery.
5. We craft a notion of bisimilarity (Definition 5.1) and prove that the order of arrival of new causal data is insignificant for bisimilar devices (Theorem 5.8).
6. We craft another notion of bisimilarity (Definition 6.2) to prove that the order of removal of causal data is also insignificant for bisimilar devices (Theorem 6.4).

Unlike traditional approaches to causality modelling that store a partial order of known causally related events and consider non related events as concurrent events, we explicitly model concurrency information and provide a broader spectrum of relations amongst events.

Real-World Benefits The technical developments of this paper are beneficial in the following ways:

Firstly, whilst being more general, our forward bisimilarity (Definition 5.6) captures replication: like-stated replicas are bisimilar. Replication in distributed systems serves fault-tolerance in that, for example, a like-stated replica will cover for a crashed replica. The idea is that, because the replica providing the cover was like-stated, the crash will go **unobserved**. Our forward bisimilarity serves that by its formalism for observational equivalence. Secondly, offline decision making (Definition 4.1) entails that, in presence of network partitions, a device gone offline will still be able to make (useful) new deductions (e.g., Lemma 4.4). It only is that the new deductions may not be at the same level of accuracy as those of its bisimilar devices that are still connected or when the device itself retrieves connection (Lemma 3.7). That is the service offline decision making provides to fault-tolerance. Thirdly, Theorems 5.8 and 6.4 are formal characterisations for strong eventual consistency—so long as all the correspondences arrive/leave, the replicas are causally consistent, i.e., forward/backward bisimilar—serving key-value stores.

2 Worlds of Events and Microcosms

Call a binary relation R a strict partial order on a set P when R is irreflexive, asymmetric, and transitive. Then, we say that (P, R) is a strict poset. For a strict poset (T, R) , when R is also total, call (T, R) a strict

chain. Let R be a relation on a set S . For a subset U of S , the symbol $R|_U$ denotes R restricted to U . We use “ $\underline{\vee}$ ” for the exclusive or of mathematical logic. For a set S , write $|S|$ for the cardinality of S . As is common in Set Theory, \aleph_0 denotes the cardinality of Natural Numbers (\mathbb{N}). Throughout this paper, “ $_$ ” is our wild card; its usage expresses our lack of interest in the exact details of what “ $_$ ” has replaced.

Definition 2.1 Call $W(\langle, \parallel)$ a world of events when:

(W1) W is an infinitely countable set (i.e., $|W| = \aleph_0$) of events that are ranged over by $e_1, e_2, \dots, e, e', \dots$,

(W2) \langle and \parallel are binary relations defined on W that are ranged over by $r_1, r_2, \dots, r, r', \dots$,

(W3) (W, \langle) is a strict poset,

(W4) \parallel is irreflexive and symmetric, and

(W5) $e_1 \neq e_2$ iff $e_1 \parallel e_2 \vee e_1 \langle e_2 \vee e_2 \langle e_1$.

For $e_1, e_2 \in W$, when $(e_1, e_2) \in r$ for $r \in \{\langle, \parallel\}$, we write $W \models e_1 r e_2$ and say $e_1 r e_2$ holds for W . \square

The relations \langle and \parallel denote the familiar happens-before and is-concurrent-with, respectively [16]. Notice that here we define \parallel explicitly – whilst the usual derived definition for non strict posets covers the elements that are not related in the order by stating $e_1 \parallel e_2$ iff $e_1 \not\prec e_2 \wedge e_2 \not\prec e_1$. Notice also that, in line with the traditional understanding about it [20, Observation 1.3], (W4) does not define \parallel transitive.

For a world of events, we take the relation $\langle \rangle$ (read is-causally-related-to)¹ as a syntactic sugar for $\langle \cup \langle^{-1}$, namely, $e_1 \langle \rangle e_2 \stackrel{\text{def}}{=} e_1 \langle e_2 \vee e_2 \langle e_1$. Hence, $\langle \rangle$ is symmetric. We can also observe that every distinct pair of events are attributed to exactly one of the basic relations, and that “ \langle ” \cap “ \langle^{-1} ” \cap “ $=$ ” \cap “ \parallel ” is always \emptyset .

Fix the set of **accurate** relations $R = \{\langle, \parallel\}$. The relation $\langle \rangle$ is an **inaccurate** relation in that it does not expose the exact known direction of \langle . We now extend \models to \models^* for when the inaccurate relation $\langle \rangle$ is also needed to be taken into consideration. Write $W \models^* e_1 r e_2$ iff: $W \models e_1 r e_2$; or, $r = \langle \rangle$ and either $W \models e_1 \langle e_2$ or $W \models e_2 \langle e_1$. Note that, unlike \models , not every distinct pair of events are attributed to a unique relation by \models^* . In particular, for every e_1 and e_2 such that $W \models e_1 \langle e_2$, by definition, it is both the case that $W \models^* e_1 \langle e_2$ and $W \models^* e_1 \langle \rangle e_2$. We call $e_1 r e_2$ a correspondence, ranged over by $c_1, c_2, \dots, c, c', \dots$. For a world of events W , we also fix $\mathcal{C}_W^* = \{c \mid W \models^* c\}$. For $c = e_1 r e_2$, we say c is an accurate correspondence when r is accurate. We call c inaccurate otherwise.

Proposition 2.2 Every world of events W is consistent: $W \models e_1 r e_2$ and $W \models e_1 r' e_2$ imply $r = r'$.

Definition 2.3 Let $W(\langle, \parallel)$ be a world of events. Call $M(I, E)$ a microcosm of W (write $M \triangleleft W$) when:

(M1) $I \subset W$ and $|I| < \aleph_0$,

(M2) $(I, \langle|_I)$ is a strict chain,

(M3) $E \subset \mathcal{C}_W^*$ and $|E| < \aleph_0$,

(M4) $e_1 r e_2 \in E$ implies that there is no chain of events e'_1, \dots, e'_n in M such that $e_1 = e'_1 \langle|_M \dots \langle|_M e'_n = e_2$ or $e_2 = e'_1 \langle|_M \dots \langle|_M e'_n = e_1$.²

Accordingly, call W the enclosing world of M and let $\mathcal{M}_W = \{M \mid M \triangleleft W\}$. \square

¹For a use of $\langle \rangle$ in reality, see the CISE proof system [14]. In a valid CISE execution, when a pair of events e and e' possess conflicting tokens, it is required that $e \langle \rangle e'$.

²More on the motivation behind (M4) to follow.

The difference between the notation we use for worlds of events and the one we use for microcosms might cause confusion at the first glance. In addition to being the world of events, the W in $W(<, \parallel)$ is a set, $<$ and \parallel are relations on which. To the contrary, the M in $M(I, E)$ is only a name for the pair (I, E) . Furthermore, I is a set of events, whilst E is a set of correspondences; they are not of the same sort.

For an $M(I, E)$, we refer to I as the **internal** events of M , and, to E as the set of **external** correspondences known to it. When appropriate, we use the alternative notions $I(M)$ and $E(M)$, respectively. Write $e_1 < e_2 \in I$ when $e_1, e_2 \in I$ and $W \models e_1 < e_2$. Besides, write $e_1 r e_2 \in M$, when $e_1 r e_2 \in I$ or $e_1 r e_2 \in E$. Write $e \in E$ when $\exists e' \in W. e _ e' \in E \vee e' _ e \in E$. Finally, write $e \in M$ when $e \in I$ or $e \in E$. That is how we formalise the notion of microcosm membership informally used in (M4).

A microcosm is our abstraction for a single state – out of the possibly many states – of a generic device. The enclosing world of events of a microcosm is the abstraction we use for the ecosystem in which a device lives. Certain events can take place locally for a device; in which case, they are stored in the internal events of the respective microcosm. The correspondence between certain events can also be disclosed to a device by the ecosystem; in which case, they are stored in the external correspondences of the respective microcosm. In our model, devices do not get to communicate directly with one another. The ecosystem sits between devices in that news from other devices in the same ecosystem arrives via the ecosystem (as opposed to the other devices themselves).

Note that, unlike a world of events, for a microcosm, the relation $<>$ is not a syntactic sugar. To the latter, an $<>$ instance is all the information that is given for the respective pair of events. In that case, whilst no stronger information about the given pair is provided to the microcosm, the enclosing world of events is aware of the exact $<$ direction between the pair. It, nevertheless, follows from irreflexiveness of $<$ that $<>$ is irreflexive too – both for worlds of events and their microcosms.

Example 2.4 For a microcosm M such that $I(M) = e_1 < e_2 < e_3$, no correspondence $e_1 r e_2$ can exist in $E(M)$ or (M4) will be violated. There is no need for any “order” to exist between all the events a microcosm knows of – be it partial or total. $M' = (\emptyset, \{e_1 <> e_2, e_2 <> e_3\})$ is an entirely fine microcosm (perhaps of the same world of events as M), in which there is neither a total order nor a partial order between e_1 , e_2 , and e_3 . $M'' = (M' + e_3 < e_4)^3$ is another permissible microcosm – regardless of whether $e_3 < e_4 \in I(M'')$ or $e_3 < e_4 \in E(M'')$. Note that M'' does contain a partial order but no total one. Finally, $M''' = (I(M), \emptyset)$ is yet another valid microcosm, in which there truly is a total order.

We now introduce the first sort of deduction for microcosms (Definition 2.5). The idea is that such a deduction is for a microcosm to decree on the correspondences it does know of. Later in Section 3, we will generalise deduction for a microcosm to also conclude that it does not know the correspondence between a given pair of events.

Definition 2.5 Let M be a microcosm. Judgements of the form $M \overset{\circ}{\vdash} e_1 r e_2$ are called the initial judgements of M when they are derived using the rules in Fig. 1. Write $M \overset{\circ}{\not\vdash} e_1 r e_2$ when $M \overset{\circ}{\vdash} e_1 r e_2$ is not true. \square

Note that with “ $_$ ” being existential in nature, the negation acts universally. In particular, $M \overset{\circ}{\not\vdash} e_1 _ e_2$ stipulates the lack of *any* initial correspondence between e_1 and e_2 in M .

³Notation defined at the end of section.

$$\begin{array}{c}
 \boxed{M \overset{\circ}{\vdash} e_1 r e_2} \quad \text{where } r \in R \cup \{<>\} \\
 \\
 \frac{e_1 r e_2 \in M}{M \overset{\circ}{\vdash} e_1 r e_2} \text{ (INIT)} \\
 \\
 \frac{M \overset{\circ}{\vdash} e_1 < e_2 \quad M \overset{\circ}{\vdash} e_2 < e_3}{M \overset{\circ}{\vdash} e_1 < e_3} \text{ (IN-TR)} \\
 \\
 \frac{M \overset{\circ}{\vdash} e_1 \parallel e_2}{M \overset{\circ}{\vdash} e_2 \parallel e_1} \text{ (CO-SYM)} \\
 \\
 \frac{M \overset{\circ}{\vdash} e_1 <> e_2}{M \overset{\circ}{\vdash} e_2 <> e_1} \text{ (CR-SYM)}
 \end{array}$$

Figure 1: Microcosm Initial Judgements

Here is an informal account of the rules in Fig. 1: (INIT) states that every piece of information that is initially provided to a microcosm is reliable in the initial judgements made inside that microcosm. (IN-TR) legislates transitivity of $<$ for initial judgements (regardless of whether the premises come from internal or external knowledge of a microcosm or a combination of those). (CO-SYM) and (CR-SYM) are routine and legislate symmetry for \parallel and $<>$.

There are two possible ways a microcosm can evolve upon receipt of new information: (Section 3 gives more details about the intuition and the semantics of the two evolution mechanisms.) For a microcosm M , when $M \not\vdash e_1 - e_2$, write $(M + e_1 r e_2)$ for the microcosm M with the additional information $e_1 r e_2$. We assume that $e_1 r e_2$ is known to be internal or external to the resulting microcosm. Likewise, when $M \not\vdash e_1 <> e_2$ (or $M \not\vdash e_2 <> e_1$), define $M[e_1 < e_2]$ for the microcosm M in which $e_1 < e_2$ replaces $e_1 <> e_2$ (or $e_2 <> e_1$).

For both addition – namely, $(M + e_1 r e_2)$ – and update – namely, $M[e_1 < e_2]$ – we assume that the change will not violate (M4). See Fig. 2a and 2b for when careless addition and update violate (M4). That can be considered a limitation in our model: Once a microcosm reaches either state, further evolution via the given correspondence is banned by (M4) forever. We believe a batch addition (and update) can circumvent that limitation; of course, subject to sanity checks. Take Fig. 2b for example: A batch update $M[e_1 < e_2, e'_1 < e'_2]$ is the simultaneous evolution of M with both $e_1 < e_2$ and $e'_1 < e'_2$, in which (M4) is also maintained when $e'_1 < e'_2$ is removed afterwards. In this paper, we disregard batch addition and update and leave them to future.

The above discussion gives us context for explaining our design choice on including (M4) in Definition 2.3. Note that, without (M4) outlawing it, after the addition of $e_1 < e_2$ to M in Fig. 2a, for instance, the initial judgements of the resulting microcosm would have become inconsistent: On the one hand, $e'_1 <> e'_2$ would have been given; on the other hand, $e'_1 < e'_2$ would have been deducible by transitivity.

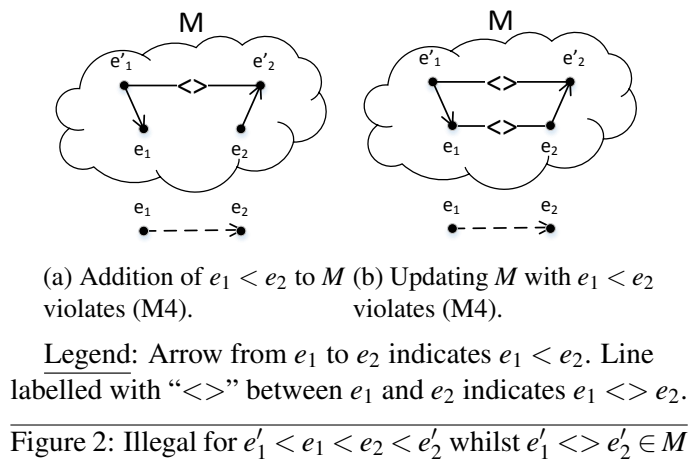


Figure 2: Illegal for $e'_1 < e_1 < e_2 < e'_2$ whilst $e'_1 <> e'_2 \in M$

3 Online Decision Making

This section provides an algorithm for a microcosm to issue its verdict on the relation it can deduce, to the best of its knowledge, to hold between a queried pair of (distinct) events. This algorithm (manifested in Fig. 3) is called the **online** decision making procedure. The idea is that the decision accuracy keeps improving using this procedure upon the inflow of the new or updated correspondences. In crude terms, this is the situation where the device is connected and thus online. Contrast this with what comes in Section 4. We prove computability (Theorem 3.4) and consistency (Theorem 3.5) of the online decision making. We show that the causal knowledge of a microcosm is strictly less than its enclosing world of events (Lemma 3.7), there is no conflict between the verdict of two microcosms of the same world of events – even when they do not issue the exact same correspondence (Corollary 3.8).

Definition 3.1 Define the online decision making procedure of a microcosm using the rules in Fig. 3. \square

In Fig. 3, a judgement $M \vdash e_1 ? e_2$ stipulates the lack of knowledge “in M ” about the correspondence between the pair of events e_1 and e_2 . As such, $?$ (read is-unknown-to) is another inaccurate relation. Note

$$\begin{array}{c}
\boxed{M \vdash e_1 r e_2} \quad \text{where } r \in R \cup \{<>, ?\} \\
\\
\frac{M \overset{\circ}{\vdash} e_1 r e_2}{M \vdash e_1 r e_2} \text{ (IN-OK)} \quad \frac{M \vdash e_1 ? e_2}{M \vdash e_2 ? e_1} \text{ (UN-SYM)} \\
\\
\frac{M \vdash e_1 r e_2 \quad M \vdash e_2 r e_3 \quad r \in \{||, <>, ?\}}{M \vdash e_1 ? e_3} \text{ (UN-1)} \quad \frac{M \vdash e_1 r e_2 \quad M \vdash e_2 ? e_3 \quad r \in R \cup \{<>\}}{M \vdash e_1 ? e_3} \text{ (UN-2)} \\
\\
\frac{M \overset{\circ}{\vdash} e_1 - e_2 \quad \nexists e' \in M. [(M \vdash e_1 - e') \wedge (M \vdash e' - e_2)]}{M \vdash e_1 ? e_2} \text{ (UN-3)} \quad \frac{e \notin M \quad e \neq e'}{M \vdash e ? e'} \text{ (UN-4)} \\
\\
\frac{M \vdash e_1 ? e_2 \quad M \triangleleft W \quad W \vDash^* e_1 r e_2}{(M + e_1 r e_2) \vdash e_1 r e_2} \text{ (STRNG)} \quad \frac{M \vdash e_1 r e_2 \quad r \neq ? \quad M \triangleleft W \quad W \vDash^* e'_1 r' e'_2}{(M + e'_1 r' e'_2) \vdash e_1 r e_2} \text{ (WEAK)} \\
\\
\frac{M \triangleleft W \quad M \vdash e_1 <> e_2 \quad W \vDash e_1 < e_2}{M[e_1 < e_2] \vdash e_1 < e_2} \text{ (UP-S)} \\
\\
\frac{M \triangleleft W \quad W \vDash e_1 < e_2 \quad M \vdash e_1 <> e_2 \quad M \vdash e'_1 r' e'_2 \quad r' \neq ?}{M[e_1 < e_2] \vdash e'_1 r' e'_2} \text{ (UP-W)}
\end{array}$$

Figure 3: Online Decision Making

that, unlike $<>$, the relation $?$ is only available for microcosms. Recall that, as axiomatised by (W5), the correspondence between every two distinct pair of events is known to their enclosing world of events.

Here is an informal account of the rules in Fig. 3:

(IN-OK) says online decision making approves of initial judgements. (UN-SYM) legislates symmetry of $?$. The next two rows concern when a microcosm judges two events as unknown to one another. (UN-1) decrees so when there is an intermediate event e_2 that has the same correspondence r with both e_1 and e_3 (in different orders albeit). Of course, given the transitivity of $<$, in the case of (UN-1), r cannot be $<$. (UN-2) is similar except that, in the microcosm of discourse, the intermediate event e_2 is unknown to e_3 . Then, (UN-3) decrees for e_1 and e_2 to be unknown to one another when there is no intermediate event in the microcosm that is in correspondence with both e_1 and e_2 . The last rule of the group, i.e., (UN-4) declares the correspondence between an event that is not in a microcosm to be unknown with any other event. Note that all the (UN-*) rules except (UN-4) assume that the microcosm has no initial judgements between the two events.

Next, the rules in the fourth row concern when a microcosm is supplied with new event information. With such a supply, the microcosm of discourse evolves into a new one. To this latter microcosm, one (and only one) more initial correspondence is available than the old microcosm. Evolution happens either by strengthening or weakening. (STRNG) states that, when two events are judged to be unknown to one another by a microcosm, the judgement will be changed accordingly when the respective information from the enclosing world of event evolves the microcosm. (WEAK) says the supply of new event information from the enclosing world of events preserves every event correspondence decreed earlier not to be unknown. Note that the supply of new information is only possible via the enclosing world of event.

Finally, the last two rules are on update of $<>$ instances. (UP-S) (for strengthening) and (UP-W) (for weakening) are the update counterparts (STRNG) and (WEAK). The difference is that, for the former pair of rules, the total number of correspondences initially known to the old microcosm and the new one are equal. Yet, in (UP-S) and (UP-W), one and only one $<>$ in the old microcosm is replaced by exactly

one $<$ in the new microcosm. The respective microcosm judgements are updated consequently.

The following lemma will later be used in Lemma 4.4.

Lemma 3.2 *Suppose that $M \vdash e_1 r e_2$, where $r \in R \cup \{<>\}$. Then, $M \stackrel{\circ}{\vdash} e_1 r e_2$.*

Here are some notational conventions to be used shortly and thereafter: We let $\Pi, \Pi', \dots, \Pi_1, \Pi_2, \dots$ range over derivation trees. For a derivation tree Π , we write $\text{lr}(\Pi)$ for the last rule used in Π . Additionally, for a derivation Π of the form

$$\frac{\Pi_1 \quad \Pi_2 \quad \dots \quad \Pi_n}{M _ c'}$$

we write $c \notin \Pi$ when $c \neq c'$ and $c \notin \Pi_1, c \notin \Pi_2, \dots, c \notin \Pi_n$. (The “ $_$ ” in “ $M _ c'$ ” above can be “ $\stackrel{\circ}{\vdash}$ ”, “ \vdash ”, and “ \vdash^* .” See Definition 4.1 for the latter.) Intuitively, $c \notin \Pi$ means that ‘ c does not appear in Π .’

Lemma 3.3 *Let $M \stackrel{\circ}{\not\vdash} c$ and $\Pi = (M + c) \vdash e_1 r e_2$ but $c \notin \Pi$. Then, $M \vdash e_1 r e_2$.*

Informally, the above lemma gives a criterion for shrinking the size of a derivation tree of online decisions. Lemma 3.3 will be used in the proof of Lemma 5.4.

Fundamental results about online decision making follow. Theorem 3.4 is on its computability. Then, Theorem 3.5 proves consistency. At last, Lemma 3.7 and Corollary 3.8 focus on the relative accuracy of online decision making.

Theorem 3.4 *Online decision making is computable: For every distinct pair of events e_1 and e_2 and microcosm M , in finite number of steps, the relation r for which $M \vdash e_1 r e_2$ can be found, if any.*

Proof. Let $p(e_1, e_2) = \exists r \in R \cup \{<>, ?\}. M \vdash e_1 r e_2$. The result is trivial when no rule applies because, then, $p(e_1, e_2) = \perp$ in zero steps. Otherwise, we assume availability of a mechanism for preventing infinite trial of the symmetry rules. Similarly, we assume a book-keeping to prevent self-lookup over seeking an intermediate event in the case of (UN- n), where $n \in \{1, 2, 3\}$. The proof is by rule-based induction on $M \vdash e_1 r e_2$. ■

Here is a note on the computational complexity of the online decision making. Let m be the size of a microcosm. Taken naïvely, the rules in Fig. 3 give rise to exponential complexity w.r.t. m . Using a simple $m \times m$ memoisation matrix, however, one can reduce that complexity to quadratic w.r.t. m . Note that, having only a partial knowledge, being quadratic w.r.t. the size of a microcosm is far less than quadratic w.r.t. the size of a world of events as a whole. This proves our model practically more useful than the classical holistic models. In the presence of the above matrix, furthermore, maintaining (M4) upon arrival of new correspondences is DLOG-Complete w.r.t. m [22, §5.7].

Theorem 3.5 *Online decision making is consistent: $M \vdash e_1 r e_2$ and $M \vdash e_1 r' e_2$ imply $r = r'$.*

Proof. Let $\Pi = M \vdash e_1 r e_2$ and $\Pi' = M \vdash e_1 r' e_2$. The proof is by rule-based induction on Π , namely, by case analysis of $\text{lr}(\Pi)$:

- (UN- n) for $n \in \{1, 2, 3\}$. In all those cases, as a part of the hypotheses, $M \stackrel{\circ}{\not\vdash} e_1 _ e_2$. Hence, $\text{lr}(\Pi') \neq (\text{IN-OK})$. Furthermore, $\text{lr}(\Pi') \neq (\text{UP-S})$ (because, then, $r' = <$ and $M \stackrel{\circ}{\vdash} e_1 < e_2$) and $\text{lr}(\Pi') \neq (\text{UP-W})$ (because, then, $r' \neq ?$, and, by Lemma 3.2, $M \stackrel{\circ}{\vdash} e_1 r' e_2$). Likewise, $\text{lr}(\Pi') \neq (\text{STRNG})$ either because, then, $M \stackrel{\circ}{\vdash} e_1 r' e_2$ for $M = (_ + e_1 r' e_2)$. We claim that the last rule in $M \vdash e_1 r' e_2$ cannot be (WEAK) either, and, the result follows because all the remaining rules imply that $r' = ?$.

We now prove our last claim. If the last rule in $M \vdash e_1 r' e_2$ is to be (WEAK), there exists a microcosm M' such that $M = (M' + e'_1 _ e'_2)$ and $M' \vdash e_1 r' e_2$. Besides, $r' \neq ?$, which, by Lemma 3.2, implies $M' \stackrel{\circ}{\vdash} e_1 r' e_2$. This is, however, a contradiction because, then $M \stackrel{\circ}{\vdash} e_1 r' e_2$.

- (UN-4). When $e \notin M$ and $e \neq e'$, there essentially is no other rule that can apply than (UN-4). That is, the last rule for $M \vdash e_1 r' e_2$ too needs to be (UN-4) and $r' = ?$.

We drop the remaining cases due to space restrictions. ■

Definition 3.6 For a pair of relations $r, r' \in R \cup \{<>, ?\}$, write $r \sqsubset r'$ –for r is at most as accurate as r' – when: (i) $r' \neq ?$ and $r = ?$, (ii) $r' = <$ and $r = <>$, (iii) $r' = <^{-1}$ and $r = <>$, (iv) $r = r' = <$, and (v) $r = r' = \parallel$. Write \sqsubseteq for the reflexive closure of \sqsubset .

The following result states that a microcosm always **approximates** its enclosing world of event: For every pair of events, when the relation a microcosm attributes to the pair does not exactly coincide with that of its enclosing world of events, the microcosm is only less accurate. This is the essence of our model being weaker than the mainstream practice where every device is exactly as accurate as its enclosing ecosystem.

Lemma 3.7 Let $M \triangleleft W$. Suppose also that $W \vDash e_1 r_W e_2$ and $M \vdash e_1 r_M e_2$. Then, $r_M \sqsubset r_W$.

The next result says: When two microcosms of the same world of events do not agree on a given pair of events, it only is that one of the two is more accurate than the other. In other words, two microcosms of the same world of events will never attribute conflicting relations to any given pair of events.

Corollary 3.8 Let $M \triangleleft W$ and $M' \triangleleft W$ with $M \vdash e_1 r e_2$ and $M' \vdash e_1 r' e_2$. Then, $r \sqsubseteq r'$ or $r' \sqsubseteq r$.

4 Offline Decision Making

The algorithm presented in this section enables a microcosm to make new decisions without depending on new supply from the enclosing world of events. As such, it suits a device required to perform offline computation. Hence, the naming “offline.” Unlike our online algorithm that exclusively proves correspondences, our offline algorithm is based on cancelling possibilities. That is, deducing it that certain correspondences cannot possibly hold between the given pair of events. We say that the online decision making *confirms*, whereas the offline one (mostly) *refutes*.

Sometimes, cancelling enough possibilities out will prove the only remaining correspondence (e.g., Fig. 4b). But, even if that is not quite the case, cancelling one or more correspondences out is still useful (e.g., Fig. 4a): It conveys the information that the given pair of events are **not** unknown to one another. (See Lemma 4.4.) Most particularly, in such a scenario, it would be wrong to consider the pair concurrent. That is in exact contrast with the common causality folklore that: ‘when one cannot confirm any correspondence between two events, one can safely [sic] consider them concurrent.’

In Fig. 4a, given that $e_1 \parallel e_2$ and $e_2 < e_3$, it cannot be that $e_3 < e_1$. This is because, then, by transitivity of happens-before, $e_2 < e_3$ and $e_3 < e_1$, imply $e_2 < e_1$, contradicting $e_1 \parallel e_2$. Fig. 4b rules $e_3 < e_1$ out similarly. But, then, given that $e_1 <> e_3$, the implication is $e_1 < e_3$. Note that the only correspondences that were available prior to concluding $e_3 \not< e_1$ (in Fig. 4a) and $e_1 < e_3$ (in Fig. 4b) were the black lines between e_1 , e_2 , and e_3 . No new correspondence was supplied over the arguments either. The important observation to make, hence, is that such arguments do not depend on new supply from the enclosing world of events. Offline decision making (Definition 4.1) enables such arguments.

Before we can delve into offline decision making itself, we need to introduce a couple of notations. For a microcosm M and a pair of distinct events e_1 and e_2 such that $M \not\vdash e_1 - e_2$, write $(M +_? e_1 r e_2)$ for a microcosm that is *structurally* the same as $(M + e_1 r e_2)$, namely, contains the exact same correspondences. Despite their same structure, the former is meant to be used only when $e_1 r e_2$ is not supplied by




Figure 4: Two Useful Offline Deductions

$M \vdash^* e_1 \tilde{r} e_2$ where $\tilde{r} ::= r \mid \not r$		
$\frac{M \vdash e_1 r e_2}{M \vdash^* e_1 r e_2}$ (ONL-OK)	$\frac{M \vdash^* e_1 \not r e_2 \quad r \in R}{M \vdash^* e_1 \not r e_2}$ (NOT-R)	
$\frac{M \vdash^* e_1 \ll e_2}{M \vdash^* e_1 \parallel e_2}$ (NOT-CR)	$\frac{M \vdash^* e_1 \not\parallel e_2}{M \vdash^* e_1 <> e_2}$ (NOT-CO)	
$\frac{(M +_? e_1 r e_2) \vdash^* e'_1 r'_1 e'_2 \quad (M +_? e_1 r e_2) \vdash^* e'_1 r'_2 e'_2 \quad r_1 \neq r_2}{M \vdash^* e_1 \not r e_2}$ (CNTRD)		
$\frac{M[e_1 r e_2]_? \vdash^* e'_1 r'_1 e'_2 \quad M[e_1 r e_2]_? \vdash^* e'_1 r'_2 e'_2 \quad r_1 \neq r_2}{M \vdash^* e_1 \not r e_2}$ (UP-CNTRD)		
$\frac{M \vdash^* e_1 <> e_2 \quad M \vdash^* e_1 \not< e_2}{M \vdash^* e_2 < e_1}$ (NOT-HB)	$\frac{M \vdash^* e_1 \not< e_2 \quad M \vdash^* e_2 \not< e_1}{M \vdash^* e_1 \ll e_2}$ (NO-HBS)	
$\frac{M \vdash^* e_1 \not r e_2}{M \vdash^* e_2 \not r e_1}$ (NU-SYM)	$\frac{M \vdash^* e_1 \ll e_2}{M \vdash^* e_2 \ll e_1}$ (NCR-SYM)	$\frac{M \vdash^* e_1 \not\parallel e_2}{M \vdash^* e_2 \not\parallel e_1}$ (NCO-SYM)

Figure 5: Offline Decision Making

the enclosing world of events; it rather is M with the *hypothesis* that $e_1 r e_2$ was also known by M . That is, “+?” is like the blue arrow in Fig. 4a. Note the additional requirement of the former over the latter. The latter only requires that $M \not\vdash e_1 - e_2$. In contrast, the former requires that $M \not\vdash e_1 - e_2$. (By definition, the requirement for $(M +_? e_1 r e_2)$ implies the requirement of $(M + e_1 r e_2)$ too. Hence, $(M +_? e_1 r e_2)$ is well-defined.) Note also that, by Theorem 3.4, satisfiability of $M \not\vdash e_1 - e_2$ is computable. When $M \not\vdash e_1 - e_2$, define $M[e_1 r e_2]_?$, similarly for a microcosm that is structurally like $M[e_1 r e_2]$; yet $e_1 r e_2$ is not supplied by the enclosing world of events but is only hypothesised. That is, “[.]” is like the blue arrow in Fig. 4b.

Definition 4.1 Define the offline decision making procedure of a microcosm using the rules in Fig. 5, where the judgements take the form $M \vdash^* e_1 \tilde{r} e_2$, and $\tilde{r} ::= r \mid \not r$.

The rules in Fig. 5 are fairly self-explanatory and we drop explanation to save space, except for the two key rules: (CNTRD) and (UP-CNTRD). If a hypothetical correspondence between a pair of events leads to two different conclusions about a single pair of distinct events, we have come to a contradiction, and, conclude the hypothesis to be false. (CNTRD) manifests that for additions whilst (UP-CNTRD) does so for updates.

The online and offline decisions on the same pair of events will not conflict. That is, online and offline decision making are consistent:

$\frac{M \vdash e_1 \parallel e_2}{(M \ +? \ e_3 < e_1) \vdash e_1 \parallel e_2} \text{ (WEAK)}$	$\frac{M \vdash e_2 < e_3}{M \vdash^{\circ} e_2 < e_3} \text{ (*)}$	$\frac{e_3 < e_1 \in (M \ +? \ e_3 < e_1)}{(M \ +? \ e_3 < e_1) \vdash^{\circ} e_3 < e_1} \text{ (INIT)}$
$\frac{(M \ +? \ e_3 < e_1) \vdash e_1 \parallel e_2}{(M \ +? \ e_3 < e_1) \vdash e_2 \parallel e_1} \text{ (CO-SYM)}$	$\frac{(M \ +? \ e_3 < e_1) \vdash^{\circ} e_2 < e_1}{(M \ +? \ e_3 < e_1) \vdash e_2 < e_1} \text{ (IN-OK)}$	$\frac{(M \ +? \ e_3 < e_1) \vdash^{\circ} e_3 < e_1}{(M \ +? \ e_3 < e_1) \vdash e_2 < e_1} \text{ (IN-TR)}$
$\frac{(M \ +? \ e_3 < e_1) \vdash e_2 \parallel e_1}{(M \ +? \ e_3 < e_1) \vdash^* e_2 \parallel e_1} \text{ (ONL-OK)}$	$\frac{(M \ +? \ e_3 < e_1) \vdash e_2 < e_1}{(M \ +? \ e_3 < e_1) \vdash^* e_2 < e_1} \text{ (ONL-OK)}$	$\frac{(M \ +? \ e_3 < e_1) \vdash^* e_2 < e_1}{(M \ +? \ e_3 < e_1) \vdash^* e_2 < e_1} \text{ (CNTRD)}$
$\frac{M \vdash^* e_3 \not< e_1}{M \vdash^* e_3 \not\geq e_1} \text{ (NOT-R)}$		

Figure 6: Mechanical Proof of Lemma 4.4

Lemma 4.2 *Let e_1 and e_2 be a pair of distinct events in M . Then: (i) $M \vdash e_1 r e_2$ implies $M \not\vdash^* e_1 \not r e_2$, and (ii) $M \vdash^* e_1 \not r e_2$ implies $M \not\vdash e_1 _ e_2$, in particular, $M \not\vdash e_1 r e_2$.*

Lemma 4.3 *If $M \vdash e_1 r e_2$ and $M \vdash^* e_1 r' e_2$, then $r = r'$.*

The offline decision making can be used, for example, to mechanically conclude in the case of Fig. 4a that $M \vdash^* e_3 \not\geq e_1$:

Lemma 4.4 *Let $M \vdash e_1 \parallel e_2$ and $M \vdash e_2 < e_3$ but $M \not\vdash e_3 _ e_1$. Then, $M \vdash^* e_3 \not\geq e_1$.*

Proof. The mechanical proof comes in Fig. 6, where the derivation labelled (*) is Lemma 3.2. ■

Despite its merits, offline decision making is confronted with two problems: Firstly, getting to refute the right correspondence may only be possible using human intelligence. Although mechanical proofs like Fig. 6 help a human-being legislate informal reasoning such as Fig. 4, how likely that is for a machine to produce that is not yet known. Secondly, the search space for getting to a contradiction (and hence a refute) is exponential in the number of events known to a microcosm. We are not aware of any technique for reducing that space.

5 Forward Bisimilarity

In this section, we present our first notion of microcosm bisimilarity. We start by defining microcosm analogy (Definition 5.1), namely, what exactly we mean when we say two microcosms agree on every correspondence. Then, we show that such microcosms will evolve likewise when supplied with the exact same new single correspondence (Theorem 5.7), i.e., they are forward bisimilar (Definition 5.6). The most important result of this section is Theorem 5.8, which proves it that the order of arrival of causal information is irrelevant so long as the same correspondences are available to a pair of bisimilar microcosms. Finally, Theorem 5.9 establishes the bisimilarity of analogy. We call the bisimilarity of this section forward to contrast it with that of next section (Definition 6.2), which we call backward.

Definition 5.1 *Call microcosms M and M' analogous – write $M \approx M'$ – when: $\forall e_1, e_2. M \vdash e_1 r e_2 \Leftrightarrow M' \vdash e_1 r e_2$.*

In words, two microcosms are analogous when they ‘agree on the correspondence between every pair of events.’ That can, for instance, be two replicas of a single data centre that are in the same state. As another example consider a copy taken from a device before it temporarily dies. As soon as the original device comes back to life, the original device and the copy would be analogous. Interestingly enough, the order of arrival of the causal information to the original device is completely sporadic to the copy. Note that Definition 5.1 has even no explicit mention of the enclosing worlds of events of the two microcosms.

Definition 5.2 *Define \xrightarrow{c} for the transition system $\mathcal{T}_F(W) = (\mathcal{M}_W, \mathcal{C}_W^*, \rightarrow)$ such that $M \xrightarrow{c} M'$ when $M' = (M + c)$ for some $c \in \mathcal{C}_W^*$. Call $\mathcal{T}_F(W)$ the forward transition system of W .*

The above definition formalises our understanding of a microcosm evolving *forward* with the arrival of new supply to it. We now present a technical lemma for later use.

Lemma 5.3 *Suppose that $M \approx M'$. Then, $\Pi = (M + c) \stackrel{\circ}{\vdash} e_1 r e_2$ implies $(M' + c) \stackrel{\circ}{\vdash} e_1 r e_2$ when $c \in \Pi$.*

Proof. By case distinction on $\text{lr}(\Pi)$. ■

The following two lemmata explore two different scenarios for forward evolution: when the new supply is not used for deriving the correspondence between a given pair of events (Lemma 5.4) and when it is (Lemma 5.5). Those two pave the road for Theorem 5.7.

Lemma 5.4 *Suppose that $M \approx M'$ and $M \xrightarrow{c} (M + c)$. Then, $\Pi = (M + c) \vdash e_1 r e_2$ implies $(M' + c) \vdash e_1 r e_2$ when $c \notin \Pi$.*

Proof. By Lemma 3.3, $(M + c) \vdash e_1 r e_2$ and $c \notin \Pi$ imply $M \vdash e_1 r e_2$. Given that $M \approx M'$, thus, $M' \vdash e_1 r e_2$. Note, then, that $M' \not\vdash e_1 _ e_2$ is not true. We claim that the result follows, i.e., $(M' + c) \vdash e_1 r e_2$.

Here is the proof of our claim. Suppose otherwise, namely, that $(M' + c) \vdash e_1 r' e_2$ where $r \neq r'$. Then, the only rule that can take $M' \vdash e_1 r e_2$ to $(M' + c) \vdash e_1 r' e_2$ is (STRNG), in which case, $r = ?$ and $r' \in R \cup \{<>\}$. But, then, $M \vdash e_1 ? e_2$ for $M \approx M'$. It follows by (STRNG) that $(M + c) \vdash e_1 r' e_2$. That, however, is a contradiction because, according to Theorem 3.5, \vdash is consistent. ■

Lemma 5.5 *Suppose that $M \approx M'$ and $M \xrightarrow{c} (M + c)$. Then, $\Pi = (M + c) \vdash e_1 r e_2$ implies $(M' + c) \vdash e_1 r e_2$ when $c \in \Pi$.*

Proof. Induction on the size of Π by case distinction on $\text{lr}(\Pi)$.

- (IN-OK). In this case, $(M + c) \stackrel{\circ}{\vdash} e_1 r e_2$. By Lemma 5.3, $(M' + c) \stackrel{\circ}{\vdash} e_1 r e_2$. Using an application of (IN-OK), one derives the desirable.
- (UN-1). In this case, $r = ?$, and, there exist $r' \in R \cup \{<>\}$ and an intermediate event e such that $\Pi_1 = (M + c) \vdash e_1 r' e$ and $\Pi_2 = (M + c) \vdash e r' e_2$. When $c \notin \Pi_1$, by Lemma 5.4, $(M' + c) \vdash e_1 r' e$. Otherwise, the same result is obtained by the inductive hypothesis. Based on whether $c \in \Pi_2$ or not, one obtains $(M' + c) \vdash e r' e_2$ similarly. The other hypothesis of this rule is $(M + c) \stackrel{\circ}{\vdash} e_1 _ e_2$. We claim that $(M' + c) \stackrel{\circ}{\vdash} e_1 _ e_2$. One derives the desirable using an application of (UN-1).
Here is the proof of our claim. Suppose otherwise. Then, $\Pi_o = (M' + c) \stackrel{\circ}{\vdash} e_1 r_o e_2$ for some $r_o \in R \cup \{<>\}$. When $c \in \Pi_o$, by Lemma 5.3, $(M + c) \stackrel{\circ}{\vdash} e_1 r_o e_2$. When $c \notin \Pi_o$, using an application of (IN-OK), one first gets $(M' + c) \vdash e_1 r_o e_2$. Next, one uses Lemma 5.4 to conclude $(M + c) \vdash e_1 r_o e_2$. Both cases, however, contradict consistency of \vdash (c.f., Theorem 3.5).
- (UN-4). This case is not applicable. Here is why. Suppose otherwise. Then, $r = ?$. Furthermore, the only way for $c \in \Pi$ is that $c = e_1 ? e_2$. But, that is not possible because, by Definition 5.2, $M \xrightarrow{c} (M + c)$ is only defined when $c \in \mathcal{C}_W^*$. That is, $W \vDash^* e_1 ? e_2$, which cannot be.
- (STRNG). In this case, $M \vdash e_1 ? e_2$. Given that $M \approx M'$, it follows that $M' \vdash e_1 ? e_2$. Using an application of (STRNG), then, $(M' + c) \vdash e_1 r e_2$.
- (WEAK). In this case, $M \vdash e_1 r e_2$. Given that $M \approx M'$, it follows that $M' \vdash e_1 r e_2$. Using an application of (WEAK), then, $(M' + c) \vdash e_1 r e_2$.

We drop the remaining cases due to space restrictions. ■

We next define our notion of forward bisimulation and prove that analogy is a bisimulation.

Definition 5.6 *Call a binary relation \mathcal{R} on \mathcal{M}_W a bisimulation for $\mathcal{T}_F(W)$ when for every microcosms M_1 and M_2 of W such that $M_1 \mathcal{R} M_2$, the following hold:*

- $M_1 \xrightarrow{c} M'_1 \Rightarrow \exists M'_2 \triangleleft W. (M_2 \xrightarrow{c} M'_2) \wedge (M'_1 \mathcal{R} M'_2)$, and
- $M_2 \xrightarrow{c} M'_2 \Rightarrow \exists M'_1 \triangleleft W. (M_1 \xrightarrow{c} M'_1) \wedge (M'_1 \mathcal{R} M'_2)$.

Write \sim_F for the bisimilarity of $\mathcal{T}_F(W)$, i.e., the largest bisimulation for $\mathcal{T}_F(W)$.

Theorem 5.7 For every W , the relation \approx is a bisimulation for $\mathcal{T}_F(W)$.

Proof. Let $M, M' \triangleleft W$ and $M \approx M'$. Suppose that $M \xrightarrow{c} (M+c)$ and $\Pi = (M+c) \vdash e_1 r e_2$. When $c \notin \Pi$, by Lemma 5.4, $(M+c) \vdash e_1 r e_2$. When $c \in \Pi$, by Lemma 5.5, $(M+c) \vdash e_1 r e_2$. The result follows by symmetry. ■

Now that we are armed with Theorem 5.7, it is easy to prove Theorem 5.8. We would like to draw the reader's attention to the small length of the proof and the simple technique used for it. Such a comfort is a consequence of bisimulation being such a strong concept.

For a given n , write \bar{c} for c_1, c_2, \dots, c_n and $n = |\bar{c}|$. Extend \rightarrow , accordingly, to $\xrightarrow{\bar{c}}$ where $\xrightarrow{\bar{c}}$ abbreviates $\xrightarrow{c_1} \circ \xrightarrow{c_2} \circ \dots \circ \xrightarrow{c_n}$. Furthermore, write $\bar{c}' = p(\bar{c})$ when \bar{c}' is a permutation of \bar{c} .

Theorem 5.8 Suppose that $M_0 \approx M'_0$. Suppose also that $M_0 \xrightarrow{\bar{c}} M$ and $M'_0 \xrightarrow{\bar{c}'} M'$, where $\bar{c}' = p(\bar{c})$. Then, $M \approx M'$.

Proof. We proceed by strong induction on n , where $n = |\bar{c}|$:

- $n = 1$. By Theorem 5.7.
- $n = k$. Suppose that the theorem is correct for every $n < k$. The case when $\bar{c} = \bar{c}'$ is immediate.

Otherwise, let k_0 be the first position where \bar{c} and \bar{c}' disagree. That is, $M_0 \xrightarrow{\bar{c}_l} M_{k_0-1} \xrightarrow{c_{k_0}} M_{k_0} \xrightarrow{\bar{c}_r} M$ and $M'_0 \xrightarrow{\bar{c}'_l} M'_{k_0-1} \xrightarrow{c'_{k_0}} M'_{k_0} \xrightarrow{\bar{c}'_r} M'$ such that $\bar{c}_l = \bar{c}'_l$, $c_{k_0} \neq c'_{k_0}$, and $\bar{c}'_r = p(\bar{c}_r)$. Then, $M_{k_0} \approx M'_{k_0}$ is immediate from Theorem 5.7. And, given that $|c_{k_0}\bar{c}_r| = |c'_{k_0}\bar{c}'_r| < k$, by the inductive hypothesis, $M \approx M'$.

The result follows. ■

Theorem 5.9 For W , the relation \approx is the bisimilarity of $\mathcal{T}_F(W)$, i.e., $\approx = \sim_F$.

Proof. Given that \sim_F is the largest bisimulation of $\mathcal{T}_F(W)$ (c.f., Definition 5.6), it suffices that we show $\sim_F \subseteq \approx$. To that end, suppose that $M \sim_F M'$; we will show that $M \approx M'$. Choose an arbitrary pair of events e_1 and e_2 such that $\Pi = M \vdash e_1 r e_2$ and $\Pi' = M' \vdash e_1 r' e_2$. The proof is by parallel induction on Π and Π' and proceeds by case distinction on $\text{lr}(\Pi)$ and $\text{lr}(\Pi')$. The goal is to show that, in all the possible cases, $r = r'$. We only show one case here:

- $M \vdash e_1 r e_2$ and $M' \vdash e_1 r' e_2$ but $\{r, r'\} \cap \{?\} = \emptyset$. Note first that, in this case, by Corollary 3.8, regardless of M and M' being bisimilar, either $r \sqsubseteq r'$ or $r' \sqsubseteq r$. We now show that, in the case of bisimilarity, $r = r'$. Let us assume that $r \sqsubseteq r'$; let us also assume that $r = <$ and $r' = <>$; the proof is similar otherwise.

When $M \vdash e_1 < e_2$ but $M' \vdash e_1 <> e_2$, by Lemma 3.2, $M \mid^\circ e_1 < e_2$ and $M' \mid^\circ e_1 <> e_2$, respectively. Hence, for an event e_3 such that $W \vdash e_2 < e_3$, it follows using an application of (IN-TR) that $M \mid^\circ e_1 < e_3$. Using an application of (IN-OK), then, $M \vdash e_1 < e_3$. This is whilst, with the given information, $M' \mid^\circ e_1 _ e_3$ is not derivable. Thus, $(M' + e_1 < e_3)$ is defined, but, $(M' + e_1 < e_3)$ is not. Let $c = e_2 < e_3$. Then, $M' \xrightarrow{c} (M' + c)$ but $M \xrightarrow{c} _$ is not implied, which contradicts $M \sim_F M'$. (See Definition 5.6.)

We omit the remaining cases due to space restrictions. ■

A shortcoming of \sim_F is that it only studies microcosm evolution via addition. Whereas microcosms can well evolve via update too. We leave the study of \sim_F in presence of updates (as well as additions) to future work. The same applies to \sim_B , which we will consider next.

6 Backward Bisimilarity

Only limited resources are available to devices, especially the edge devices. Emptying the disk or memory of such a device is routine then. To that end, usually, one removes the outdated data to come to a new manageable state. This section deals with when (causal) information is to be removed from devices, say due to resource limitation or outdatedness. That too can be seen as an evolution for a microcosm, albeit *backward* (Definition 6.1). We show that microcosm analogy (Definition 5.1) gives rise to a bisimilarity for backward evolution as well (Theorem 6.5). Besides, this section presents the backward counterpart of Theorem 5.8 that proves: The order of removal of causal information from bisimilar devices does not matter in that they will again be bisimilar once they are both done with the set of correspondences (Theorem 6.4).

Definition 6.1 Define $\overset{c}{\leftarrow}$ for the transition system $\mathcal{T}_B(W) = (\mathcal{M}_W, \mathcal{C}_W^*, \overset{c}{\leftarrow})$ such that $M \overset{c}{\leftarrow} M'$ when $M = (M' + c)$ for some $c \in \mathcal{C}_W^*$. Call $\mathcal{T}_B(W)$ the backward transition system of W .

The notation $M \overset{c}{\leftarrow} M'$ is indeed intended to be read from left to right to denote getting from M to M' by the removal of c .

Definition 6.2 Call a binary relation \mathcal{R} on \mathcal{M}_W a bisimulation for $\mathcal{T}_B(W)$ when for every microcosms M_1 and M_2 of W such that $M_1 \mathcal{R} M_2$, the following hold:

- $M_1 \overset{c}{\leftarrow} M'_1 \Rightarrow \exists M'_2 \triangleleft W. (M_2 \overset{c}{\leftarrow} M'_2) \wedge (M'_1 \mathcal{R} M'_2)$, and
- $M_2 \overset{c}{\leftarrow} M'_2 \Rightarrow \exists M'_1 \triangleleft W. (M_1 \overset{c}{\leftarrow} M'_1) \wedge (M'_1 \mathcal{R} M'_2)$.

Write \sim_B for the bisimilarity of $\mathcal{T}_B(W)$, i.e., the largest bisimulation for $\mathcal{T}_B(W)$.

Theorem 6.3 For every W , the relation \approx is a bisimulation for $\mathcal{T}_B(W)$.

We extend $\overset{c}{\leftarrow}$, like $\overset{c}{\rightarrow}$ to $\overset{\bar{c}}{\leftarrow}$ where $\overset{\bar{c}}{\leftarrow}$ abbreviates $\overset{c_1}{\leftarrow} \circ \overset{c_2}{\leftarrow} \circ \dots \circ \overset{c_n}{\leftarrow}$. In words, the following theorem states that the order of removal is irrelevant so long as the same set of correspondences are removed from analogous microcosms.

Theorem 6.4 Suppose that $M_0 \approx M'_0$. Suppose also that $M_0 \overset{\bar{c}}{\leftarrow} M$ and $M'_0 \overset{\bar{c}'}{\leftarrow} M'$, where $\bar{c}' = p(\bar{c})$. Then, $M \approx M'$.

Proof. Similar to Theorem 5.8. ■

Theorem 6.5 For W , the relation \approx is the bisimilarity of $\mathcal{T}_B(W)$, i.e., $\approx = \sim_B$.

7 Related Work

The partial knowledge of a microcosm w.r.t. its enclosing world of events resembles the classical “knowledge vs common knowledge” model [15, 11]. The latter works, however, take an algorithmic approach. Whereas our work is proof-theoretic. Ben-Zvi and Moses [5, 4] take the same approach to coin the *Syn-causality* as an extension to happens-before for synchronised computations. Gonczarowski and Moses [13] too generalise the classic model to characterise the interactive epistemic state when temporal constraints must be met. The final work in this thread [1] extends the classic model for reasoning about trust in distributed settings.

Burckhardt [6] takes a novel approach to define causal consistency not just in terms of happens-before, but also w.r.t. arbitration order and visibility order. The gain is a more precise definition of

how causality is used to ensure consistency. In addition to being model theoretic, unlike our work, his approach is not based on explicit causality [3].

One particular motivation for confining the universal knowledge of a world of events to microcosms is scalability. Systems that reduce the overhead of maintaining scalable causal consistency in wide-area replicated key-value stores include Orbe [9], COPS [17], Eiger [18], and ChainReaction [2]. COPS, in particular, defines *causal+ consistency*, which extends causal consistency with convergent conflict handling. This ensures that replicas that see concurrent updates will be updated in a consistent fashion. The systems mentioned above can incur significant overhead (in computation, storage, network load, and latency) to maintain causal consistency in scalable fashion. Du et. al [10] explain the performance overhead of causal consistency vs. eventual consistency. They introduce a protocol to reduce this overhead at the cost of degrading the quality-of-service (offered to the client) by significantly increasing data staleness.

8 Conclusion and Future Work

To the best of our knowledge, this is the first proof-theoretic modelling of causality in distributed systems, with special emphasis on partiality of causal knowledge. In our model, a device has strictly less causal information than a holistic causality store (Lemma 3.7). We offer rules for deducing causal information both when a device is online and offline (Definitions 3.1 and 4.1). We prove properties of our deductions, which are both theoretically attractive and practically valuable (Theorems 3.4, 3.5, Corollary 3.8, and Lemmata 4.2 and 4.3). We refute a causality folklore using a mechanical proof (Lemma 4.4). We define two notions of bisimilarity (Definitions 5.6 and 6.2) to prove that the order of addition or removal of causal data is irrelevant for bisimilar devices (Theorems 5.8 and 6.4, respectively).

There are two immediate improvements to our model that form interesting future work. The first is the study of how to retain (M4) whilst still not disallowing arrival of new information (like Fig. 2). The second is getting forward bisimilarity (and, therefore, backward bisimilarity) to also consider evolution from one microcosm to another by updates (as well as additions).

Our modelling does not take into consideration that information about concurrent events might arrive not at the same time. That lag makes a device observe an internal ordering for concurrent events. The interplay between the concurrency and the internal order becomes more interesting when relaying the concurrency to the next device in the vicinity. Studying that interplay is future work. We anticipate that a new set of proof systems will be required, their status w.r.t. the ones in this paper also requires dedicated study. Another related future work is to take arbitration and visibility into account.

The ability to reason about partial causal information suggests positive interaction with *causal+ consistency*: replicas that are actually causal but for which the causality is not known yet will remain consistently updated as the known causality increases (i.e., updates do not have to be redone as knowledge increases). This is an important property of *causal+ consistency* that can be a useful model to have together with the deduction systems introduced in this paper. Future work will reveal how the ability to deduce causality can increase the efficiency of COPS (and its counterparts) by reducing the overhead.

Acknowledgements This work was partially funded by the SyncFree project in the European Seventh Framework Programme under Grant Agreement 609551 and by the Erasmus Mundus Joint Doctorate Programme under Grant Agreement 2012-0030. Our special thanks to the SyncFree peers for their prolific comments on the early versions of this work. We would like to also thank the anonymous referees for their constructive discussion over the ICE forum.

References

- [1] A. Abdul-Rahman (2005): *A Framework for Decentralised Trust Reasoning*. Ph.D. thesis, U. London.
- [2] S. Almeida, J. Leitão & L. E. T. Rodrigues (2013): *ChainReaction: A Causal+ Consistent Datastore Based on Chain Replication*. In Z. Hanzálek, H. Härtig, M. Castro & M. F. Kaashoek, editors: *8th EuroSys*, ACM, pp. 85–98, doi:10.1145/2465351.2465361.
- [3] P. Bailis et al. (2012): *The Potential Dangers of Causal Consistency and an Explicit Solution*. In M. J. Carey & S. Hand, editors: *3rd SOCC*, ACM, pp. 22–1–22–7, doi:10.1145/2391229.2391251.
- [4] I. Ben-Zvi (2010): *Causality, Knowledge and Coordination in Distributed Systems*. Ph.D. thesis, Technion.
- [5] I. Ben-Zvi & Y. Moses (2010): *Beyond Lamport’s Happened-Before: On the Role of Time Bounds in Synchronous Systems*. In N. A. Lynch & A. A. Shvartsman, editors: *24th DISC, LNCS 6343*, Springer, pp. 421–436, doi:10.1007/978-3-642-15763-9_42.
- [6] S. Burckhardt (2014): *Principles of Eventual Consistency*. *FTPL* 1(1-2), pp. 1–150, doi:10.1561/2500000011.
- [7] B. Charron-Bost (1991): *Concerning the Size of Logical Clocks in Distributed Systems*. *Inf. Proc. Lett.* 39(1), pp. 11–16, doi:10.1016/0020-0190(91)90055-M.
- [8] G. DeCandia et al. (2007): *Dynamo: Amazon’s Highly Available Key-Value Store*. In: *21st SOSP*, pp. 205–220, doi:10.1145/1294261.1294281.
- [9] J. Du et al. (2013): *Orbe: Scalable Causal Consistency using Dependency Matrices and Physical Clocks*. In G. M. Lohman, editor: *SOCC*, ACM, pp. 11:1–11:14, doi:10.1145/2523616.2523628.
- [10] J. Du et al. (2014): *Closing the Performance Gap between Causal Consistency and Eventual Consistency*. In: *1st PaPEC*, EPFL-CONF-198281, ACM.
- [11] R. Fagin et al. (2003): *Common Knowledge Revisited*. In V. F. Hendricks, K. F. Jørgensen & S. A. Pedersen, editors: *Knowledge Contributors, Synthese Library 322*, Springer Netherlands, pp. 87–104, doi:10.1007/978-94-007-1001-6_5.
- [12] S. Gilbert & N. Lynch (2002): *Brewer’s Conjecture and the Feasibility of Consistent Available Partition-Tolerant Web Services*. 33, pp. 51–59, doi:10.1145/564585.564601.
- [13] Y. A. Gonczarowski & Y. Moses (2013): *Timely Common Knowledge*. In B. C. Schipper, editor: *14th TARK*.
- [14] A. Gotsman et al. (2016): *’Cause I’m Strong Enough: Reasoning about Consistency Choices in Distributed Systems*. In R. Bodík & R. Majumdar, editors: *43rd POPL*, ACM, pp. 371–384, doi:10.1145/2837614.2837625.
- [15] J. Y. Halpern & Y. Moses (1990): *Knowledge and Common Knowledge in a Distributed Environment*. *JACM* 37(3), pp. 549–587, doi:10.1145/79147.79161.
- [16] L. Lamport (1978): *Time, Clocks, and the Ordering of Events in a Distributed System*. *Commun. ACM* 21(7), pp. 558–565, doi:10.1145/359545.359563.
- [17] W. Lloyd et al. (2011): *Don’t Settle for Eventual: Scalable Causal Consistency for Wide-Area Storage with COPS*. In: *23rd SOSP*, ACM, New York, NY, USA, pp. 401–416, doi:10.1145/2043556.2043593.
- [18] W. Lloyd et al. (2013): *Stronger Semantics for Low-Latency Geo-Replicated Storage*. In N. Feamster & J. C. Mogul, editors: *10th NSDI, USENIX*, pp. 313–328.
- [19] F. B. Schneider (1990): *Implementing Fault-Tolerant Services using the State Machine Approach: A Tutorial*. *ACM CSUR* 22(4), pp. 299–319, doi:10.1145/98163.98167.
- [20] R. Schwarz & F. Mattern (1994): *Detecting Causal Relationships in Distributed Computations: In Search of the Holy Grail*. *Dist. Comp.* 7(3), pp. 149–174, doi:10.1007/BF02277859.
- [21] P. Sérgio Almeida et al. (2014): *Scalable and Accurate Causality Tracking for Eventually Consistent Stores*. In: *14th IFIP DAIS*, pp. 67–81, doi:10.1007/978-3-662-43352-2_6.
- [22] C. K. Yap (1998): *Theory of Complexity Classes*. 1. <https://cs.nyu.edu/yap/book/complexity/>.