What's "Next"?

Walker White, Mirek Riedewald, Johannes Gehrke, Alan Demers
Department of Computer Science, Cornell University
{wmwhite,mirek,johannes,ademers}@cs.cornell.edu

ABSTRACT

Event processing systems have wide applications ranging from monitoring RSS feeds to managing events from RFID readers, and there exists much work on them in the literature. Many competing temporal models for event systems have been proposed, with no consensus on which approach is best. In this paper we determine the important properties for such temporal models. Our approach is to define a very general temporal model capable of representing time in all of the major event systems. We introduce axioms motivated by the time stamp ordering relation and the semantics of the successor operator, which is present in all event systems. Only two of our axioms are controversial; the remaining axioms are satisfied by all event systems.

We consider the temporal models obtained using our full set of axioms, and the models that result when one or the other of our controversial axioms is weakened. In one case we see that there is no acceptable temporal model. In the other two cases, we show that the resulting temporal model is effectively unique up to isomorphism, leaving us with only two different models. Finally, we argue that one of the two models is better than the other when both naturalness of semantics and efficiency of implementation are considered.

1. INTRODUCTION

Event processing systems are an important component of today's information system infrastructures, with widespread applications ranging from monitoring RSS feeds to handling events from RFID readers and stock tickers. Event processing systems are a rather mature research field, and there are many system designs to choose from [1, 4, 5, 7, 11, 12]. In general, an event processing system's input is a stream of *primitive* events, which are generated by external processes. For example, each new posting on a blog or news site could result in a new primitive event that is appended to an RSS feed for that site. Users of an event processing system register long-running queries, also called subscriptions, to detect patterns of interest in the primitive event stream in real-time. These patterns are referred to as *composite* events, because they are composed of several primitive events that together satisfy the query. Typical examples of composite events are series of RSS events on the same news topic or a monotonically increasing stock price sequence. Advanced event processing systems support queries not only over streams of primitive events, but also over composite events, enabling the discovery of more complex patterns based on simpler patterns.

A major difference among existing event systems is how they handle time. Every event system must have a time stamp model, so that it can handle (partially) ordered events. Furthermore, if the event system produces composite events, it needs a rule for generating the time stamp of a composite event from the time stamps of its components. This is similar to the problem of performing a join in a temporal database between two time-stamped tables.

Time stamp models have been studied, but to date an important aspect of the temporal model has been overlooked: the notion of the *"next"* event or *successor*. In any system that treats composite events, the event definition language includes at least one *sequencing* (or immediate concatenation) operator, whose usage is denoted $E_1\,;E_2$ for events E_1 and E_2. In fact, this is the fundamental operator in any event system and it is virtually impossible to formulate a query for non-trivial patterns without using it. Recognizing this pattern requires a definition of immediate successor. There is a standard definition of immediate successor in any partially ordered set, and hence in any temporal model: t_1 is a successor of t_0 if $t_0 \leq t_1$ and there is no t_2 with $t_0 \leq t_2 \leq t_1$. Nevertheless, successors add a degree of freedom to the temporal model: Two temporal models with the same time stamp ordering relation can support different notions of successor, and at least one event system [12] has two different sequencing operators.

Most event systems have additional features in their language for composing an event with some later event (not just the next one). Some of them even allow for event consumption that removes events from the system [4]. However, sequencing is basic to all event systems, and any temporal model that does not correctly support sequencing cannot be used in an event system, regardless of its other features. Therefore, we are concerned with finding temporal models with the "right" definition of successor.

1.1 The Importance of Successor

To illustrate how the details of the successor definition of the temporal model affect the behavior of an event system, we present an example query involving sequencing, and consider its behavior under different temporal models.

Consider the following query over RSS feeds.

Query 1. After a new posting on Slashdot referring to a product announcement from Apple, notify me of the first blog to have added at least two new postings that link to the same Apple product announcement.

As before, we assume that individual blog postings are the primitive input events. The "two new postings" together form a composite event. Notice that we have two sequencing operations in this query: one to detect the two new postings at a blog site, and the other to find the overall sequence of the Slashdot posting followed by the two blog postings. It is a fairly sophisticated parameterized query [5], and not all existing event systems can express it. But our purpose here is simply to compare how this query would be interpreted using various sequencing rules.

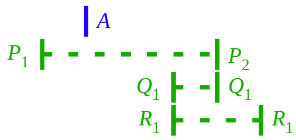

Figure 1: Overlapping Blog Posts

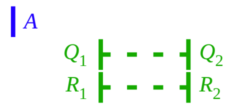

Figure 2: Simultaneous Blog Posts

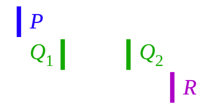

Figure 3: Associativity Example

Consider the sequence of events represented in Figure 1. In the figure, A represents a posting on Slashdot. Each P_i, Q_i and R_i represents a posting on another blog — **P**, **Q** and **R**, respectively — linking to the Apple announcement. The dashed horizontal lines connect the postings that make up the composite events. Intuitively, only the composite event from blog **Q** satisfies the query. The blog postings P_1P_2 are disqualified because the first posting P_1 precedes the Slashdot posting A. The blog **R** event R_1R_2 is disqualified because Q_1Q_2 clearly finished before it, and hence the R_1R_2 event cannot reasonably be considered "first".

Now consider the event stream of Figure 2. Intuitively, the composite events Q_1Q_2 and R_1R_2 both satisfy the query: each consists of a pair of postings linking to the Apple announcement, and they start and finish *simultaneously*. Thus, the query result should be a notification for both events. This requires a temporal model that allows an event to have multiple successors — both Q_1Q_2 and R_1R_2 must be successors of A.

We now consider how the temporal models of several existing event systems treat this query. A number of systems such as Snoop [4] or EPL [11] are based on *point time stamps*: Time stamp values come from a discrete, totally ordered domain. No such system can correctly answer this query: There is no way for a point time stamp to represent the fact that the first blog posting P_1 preceded blog posting A while the second posting P_2 followed it. Hence for point time stamps the P_1P_2 event will appear (incorrectly) to satisfy the query. A similar argument applies to the weak successors of Active Office [12]. In order to express this query correctly, the successor relation must prohibit overlap between time *intervals*.

Next consider the temporal models of SnoopIB [1] and ODE [7]. These models are interval-based, with a partial order on time stamp intervals that orders *non-overlapping* intervals in the natural way. They also use the standard successor as defined above. In this model both the Q_1Q_2 and R_1R_2 composite events from Figure 1 qualify as successors of the A event. In general, these models allow an event to have multiple successors — even an infinite set — and not all the successors are required to finish at the same time. Again, this temporal model does not accurately reflect the intuitive meaning of the query. In addition, the possibility for an event to have infinitely many successors of unbounded duration raises serious implementation issues, which we will discuss in Section 3.2.

Of the existing temporal models we have studied, the only one to capture the correct intuitive meaning of our example query on the event stream of Figure 1 is the strong successor of Active Office. But even this model fails on the stream of Figure 2. As argued above, the query result for this stream should be notification for both composite events Q_1Q_2 and R_1R_2. However, the strong successor rules of Active Office select only one of these composite events, using a tie-breaking scheme based on arbitrarily assigned, totally ordered unique identifiers.

Motivated by the above examples, we now informally describe another temporal model, the one used in the Cayuga system [5]. Cayuga uses interval time stamps like SnoopIB, but a different successor definition. Specifically, $t = [t_0, t_1]$ is a successor of $s = [s_0, s_1]$ if $t_0 > s_1$ and there is no event with time stamp $p = [p_0, p_1]$ such that $s_1 < p_0 < p_1 < t_1$. In other words, t is a successor of s if t follows s without overlap, and no p that follows s without overlap finishes before t. This definition deals correctly with our motivating example query in all cases, and also avoids infinite successor sets with their associated implementation difficulties.

1.1.1 *Implementation Considerations*

From an implementation point of view, left-associated query expressions are usually preferable over right-associated ones, because they are easier to implement, e.g., by a single finite automaton [5]. For example, Query 1 is a right-associated query; it has the form $E_0;(E_1;E_2)$. To process this query, we must first match the two blog postings $E_1;E_2$, before we match them to the Slashdot posting in E_0. The E_0 events arrive first, and thus this query cannot be processed as such by a single automaton.

Thus, a desirable property of an event system is that sequencing should be associative. Using associativity, the system could rewrite Query 1 as $(E_0;E_1);E_2$, allowing a substantially simpler implementation by using a single finite automaton [5]. Hence associativity can be viewed as an important enabler for query optimization in an event system implementation.

Unfortunately, as we show in this paper, associativity of sequencing has serious implications for the temporal model. Consider the event stream of Figure 3. Here P is an event matching some event expression E_P; Q_1 and Q_2 are events matching E_Q; and R is an event matching E_R. Using any of the temporal models discussed above, the expression $E_P;E_Q$ yields a single composite event PQ_1; and thus the left-associated expression $(E_P;E_Q);E_R$ yields only the event PQ_1R. However, expression $E_Q;E_R$ yields *two* composite events, Q_1R and Q_2R, and hence the right-associated expression $E_P;(E_Q;E_R)$ yields the events PQ_1R and PQ_2R using any of the temporal models except Active Office (which eliminates one of the composite events due to its tie-breaking rule). While Active Office handles this particular expression correctly, even it fails to be associative, as we show in Section 4.2.1.

We know of no existing system whose temporal model supports associative sequencing. As we show in Section 4.1, there is a very a good reason for this: Up to isomorphism, the *only* temporal model that supports associative sequencing is the *complete-history* model. This model, as its name suggests, requires a system to store the time stamps of *all* the primitive events that make up each composite event. With no upper bound on the size of a time stamp representation, complete-history is prohibitively expensive, and we are not aware of an implemented system using it.

1.2 Outline of Contributions

The technical content of this paper is a thorough study of temporal models for event systems through an axiomatic approach.

We start in Section 2 by giving a formal definition of a temporal model. This definition is capable of describing the temporal models of *all* event systems we are aware of, and captures the subtle distinction between models with identical time stamp orderings but different successor definitions. This gives us a uniform framework

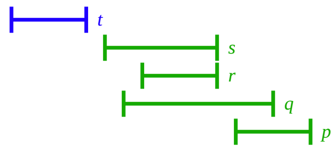

Figure 4: Intervals with the Standard Partial Order

in which to discuss existing design choices. In Section 3, we then present a set of axioms, algebraic properties that intuitively a temporal model is expected to have. Nearly all of our axioms are universally accepted and satisfied by all existing event systems; they are accepted in the temporal logic community as well [13]. Two of the axioms, which we call STRONG THICKENING and $\otimes$-DECOMPOSITION, are not universally accepted — that is, violated by at least one existing event system.

Given these axioms, we prove several results. First, in Section 4.1 we consider the *complete-history* model, a model in which we remember all the time stamps of primitive events that make up a composite event. We show in Section 4.1.2 that the complete-history model is the unique model, up to isomorphism, that satisfies all of our axioms. The complete-history model is not used by any event system as it is impractically expensive.

Then in Section 4.2, we discuss the results of eliminating or weakening the controversial axioms. Eliminating the STRONG THICKENING axiom does not result in any obvious benefits; there is no interval model in which we can limit the successor to having a single time stamp. We also consider weakening the other controversial axiom, $\otimes$-DECOMPOSITION. Given a few minor additional assumptions, the resulting model is Cayuga, unique up to isomorphism. Hence reasonable weakenings of our axiom system yield only two temporal models. Among these two models, the Cayuga model appears to be the best compromise with respect to both natural query semantics and implementation efficiency.

We end the paper with a discussion of related work (Section 5) and concluding remarks with a discussion of future work (Section 6).

2. TEMPORAL MODELS

In Section 1, we saw how the successor was defined differently in three event systems: SnoopIB, Active Office, and Cayuga. All three of these event systems use intervals as time stamps, and they have the same partial order on these intervals: $[s_0, s_1] < [t_0, t_1]$ if and only if $s_1 < t_0$. (We will refer to this partial order on intervals as the *canonical partial order* on intervals.) However, the three systems differ in how they choose a successor.

To see the difference between these three event systems, consider the intervals illustrated in Figure 4. The intervals s, r, q, and p all follow t in the canonical partial order. In SnoopIB, the immediate successors of t are s, r, and q; p is not an immediate successor as r (and also s) is between it and t. In Cayuga, the immediate successors of t are only s and r, as they have the earliest end time. Finally, Active Office (strong successor) chooses s as the unique immediate successor of t as it is the interval following t with the earliest end and start time (ordered first by end time, then by start time).

The example illustrates that despite using the same interval time stamps and the same partial order of intervals, the three systems behave very differently. In order to study these differences, we need to extend previous work on temporal models [13] by considering the successor operation explicitly in the model.

To be as general as possible, we define the *successor function* SUCC as a function that takes as input a time stamp t together with a set of time stamps $\mathcal{F}$ and produces the set $\text{SUCC}(t, \mathcal{F})$ of immediate successors of t in $\mathcal{F}$. The intuition for this model is that *candidate set* $\mathcal{F}$ represents the set of time stamps from which the immediate successor is chosen. Given an event expression $E_1 ; E_2$, an event system matches this expression by doing the following:

(a) Determine the set of candidate time stamps $\mathcal{F}$ for events matching E_2.

(b) For each event matching E_1 at time t, compose it with any event matching E_2 at time $\text{SUCC}(t, \mathcal{F})$.

The three event systems define $\text{SUCC}([s_0, s_1], \mathcal{F})$ as follows: SnoopIB:

$$\left\{[t_0,t_1]\in\mathcal{F} \,\middle|\, s_1 < t_0 \text{ and } \neg\exists[r_0,r_1]\in\mathcal{F} \text{ s.t. } s_1 < r_0 \le r_1 < t_0\right\}$$

Cayugua:

$$\left\{[t_0,t_1]\in\mathcal{F} \,\middle|\, s_1 < t_0 \text{ and } \neg\exists[r_0,r_1]\in\mathcal{F} \text{ s.t. } s_1 < r_0 \le r_1 < t_1\right\}$$

Active Office (strong successor):

$$\left\{[t_0,t_1]\in\mathcal{F} \,\middle|\, \begin{array}{l} s_1 < t_0 \wedge \neg\exists[r_0,r_1] \in \mathcal{F} \text{ s.t.} \\ s_1 < r_0 \wedge (r_1 < t_1 \vee (r_1 = t_1 \wedge r_0 < t_0)) \end{array}\right\}$$

The successor operation is not the only way that event models may differ. When we compute the result of a query like Query 1, we need to assign the composite event a new time stamp. In SnoopIB, Cayuga, and Active Office, the composite event gets the smallest interval containing the intervals of all events that make up the query result. For example, Query 1 is made up of three events. If these events happen at times 1 (Slashdot post), 2 (first blog post), and 4 (second blog post), then all three systems assign $[1, 4]$ as the result time stamp. However, ODE is different. It keeps a complete history of all of the time stamps in the component events. In the example, it would store $[1, 2, 4]$ as the result time stamp. However, notice that ODE still only uses the *boundaries* of this history when determining the immediate successor, treating the history like an interval. However, the fact remains that in ODE the time stamps $[1, 2, 4]$ and $[1, 3, 4]$ are different.

For the remainder of the paper, our approach will be more formal. Traditionally a *temporal model* is defined as $(T, \prec)$ where $\prec$ is a partial order on T [13]. To be able to study both immediate successor and event composition, we extend this definition of a temporal model to a quadruple $\mathbb{T} = (T, \prec, \text{SUCC}, \otimes)$. In this model, T is the set of all possible time stamps and $\prec$ is the partial order on these time stamps, as in the traditional model. In addition, the successor function $\text{SUCC} : T \times 2^T \rightarrow 2^T$ takes a time stamp t together with a set of candidates $\mathcal{F}$ and produces the set of immediate successors. Finally, the *composition operation* $\otimes$ takes the time stamps s and t of two events and produces the time stamp $s \otimes t$ for the corresponding composite event. Notice that SUCC and $\otimes$ are defined both for primitive as well as composite events. For convenience, we will identify $\mathbb{T}$ and T when the context is clear (e.g. a time stamp $t \in \mathbb{T}$).

While $\otimes$ behaves like a monoid operation, we do not always want it to be defined. For example, in an interval model like SnoopIB, we never want to compose two overlapping events. To avoid the use of partial operations, we introduce a special "undefined" time stamp $\perp$ to $\mathbb{T}$ such that for any t, $\mathcal{F}$, (a) $\perp \notin \text{SUCC}(t, \mathcal{F})$, (b) $\text{SUCC}(\perp, \mathcal{F}) = \emptyset$, and (c) $t \otimes \perp = \perp \otimes t = \perp$. We say that $s \otimes t$ *is defined* whenever $s \otimes t \neq \perp$.

2.1 Some Concrete Examples

We have already outlined how to express SnoopIB, Active Office, and Cayuga in our framework. As an illustrative example, we will give a complete implementation of ODE. In ODE, all time stamps are monotonically increasing finite sequences over the discrete linear order $\mathbb{Z}$. In other words, the time stamps are sequences $\sigma = \sigma(0)\sigma(1)\ldots\sigma(k-1)$ where $\sigma(i) < \sigma(i+1)$ for all $i < \ell(\sigma)-1$, with $\ell(\sigma) = k$ the *length* of the sequence. The partial order is defined as $\sigma \prec \tau$ exactly when $\sigma(\ell(\sigma)-1) < \tau(0)$, i.e., when the largest element of σ is less then the smallest element of τ. The successor operation is defined as

$$\textsc{succ}(\sigma, \mathcal{F}) = \{\, \tau \in \mathcal{F} \mid \sigma \prec \tau \text{ and } \neg\exists\rho \in \mathcal{F}, \sigma \prec \rho \prec \tau \,\}$$

Finally, for two events $\sigma < \tau$, the composition $\sigma \otimes \tau$ is the standard sequence composition (concatenation of sequences).

As the intermediate time stamps (i.e., not the interval boundaries) in ODE are not used in the definition of either the partial order or the successor function, they are not particularly useful with respect to the temporal model.

An interesting variation is the *complete-history model*. In this model, T, $\prec$ and $\otimes$ are exactly the same as in ODE. However, the successor function is different. We define a linear ordering on time stamp histories by letting $\sqsubseteq$ be the lexicographical ordering from the end of the sequences. In other words, $\sigma \sqsubseteq \tau$ if either

- $\sigma(\ell(\sigma)-i) < \tau(\ell(\tau)-i)$, and $\sigma(\ell(\sigma)-k) = \tau(\ell(\sigma)-k)$ for $k < i$, or
- $\ell(\sigma) < \ell(\tau)$ and $\sigma(\ell(\sigma)-i) = \tau(\ell(\tau)-i)$ for all $i < \ell(\sigma)$.

We use this linear order to break ties, and thus define

$$\textsc{succ}(\sigma, \mathcal{F}) = \{\, \tau \in \mathcal{F} \mid \sigma \prec \tau \text{ and } \neg\exists\rho \in \mathcal{F}, \sigma \prec \rho \sqsubseteq \tau \,\}$$

Notice that this model is also a generalization of time stamps in Active Office to complete histories, though it does not use data elements (identifiers) to break ties, only time stamp ordering.

3. AXIOMATIZING TEMPORAL MODELS

Our temporal model provides us with a very general framework for studying time in event systems. It can represent time stamps that are points, intervals, sets of points, sets of intervals, and so on. One of the reasons for this generality is that we have put no restrictions on the definitions of SUCC and $\otimes$. This means we can have aberrant behavior, e.g., a model in which two time stamps are successors of each other (i.e., $t_0 \in \textsc{succ}(t_1, \mathcal{F})$, $t_1 \in \textsc{succ}(t_0, \mathcal{F})$). Clearly we do not want such things to happen.

As in any algebraic model, we prevent such aberrant behavior by adding axioms that express properties of "reasonable" temporal models. Since adding axioms restricts the class of valid models, we want to be sure that we are not excluding perfectly acceptable models. Therefore, it is important that our axioms all be properly motivated.

We separate our axioms into two categories: *accepted axioms* and *desirable axioms*. Accepted axioms are non-controversial; they are satisfied by the temporal models in all of the major event systems. Desirable axioms, on the other hand, are each violated by at least one major event system. However, as we shall demonstrate, there are compelling reasons for wanting our temporal models to satisfy the desirable axioms.

3.1 Accepted Axioms

Many of the accepted axioms have already been implicitly mentioned in our discussion of temporal models. For the sake of completeness, in this section we will make all of these assumptions explicit. As we have several axioms, we organize them according to their defining feature: $\prec$, SUCC, or $\otimes$.

3.1.1 The $\prec$ Axioms

As in traditional temporal models, $\prec$ should be a partial order. The following two axioms capture this property.

AXIOM 1 (TRANSITIVITY). *If $t_0 \prec t_1$, $t_1 \prec t_2$, then $t_0 \prec t_2$.*

AXIOM 2 (IRREFLEXIVITY). *For any $t \in \mathbb{T}$, $t \not\prec t$.*

3.1.2 The SUCC *Axioms*

Another implicit assumption of our discussion has been that we always chose the successor time stamp from the candidate set $\mathcal{F}$. This assumption is expressed by the following axiom.

AXIOM 3 (CANDIDATE PRESENCE). *For all $t \in \mathbb{T}$ and $\mathcal{F} \subseteq \mathbb{T}$, $\textsc{succ}(t, \mathcal{F}) \subseteq \mathcal{F}$*

Additionally, the idea of successor is tightly-coupled with the partial order $\prec$. For example, if b is a successor of a, we generally assume that a "happens before" b. We capture this idea with the following axiom.

AXIOM 4 (RESPECTING ORDER). *For any $t, s \in \mathbb{T}$, $t \prec s$ if and only if there is some $\mathcal{F}$ such that $s \in \textsc{succ}(t, \mathcal{F})$*

Another issue with the successor operation is that we only want canonical successors to be successors. In all the major event systems, the elements of $\textsc{succ}(t, \mathcal{F})$ are *natural $\prec$-successors* of t. That is, $\textsc{succ}(t, \mathcal{F})$ contains only elements of $\mathcal{F}$ that follow t and have no $\prec$-intermediate elements. In fact the existing event systems differ only in how they choose from these $\prec$-successors; SnoopIB and ODE take them all, while Cayuga and Active Office are more selective and "break ties". More generally, we want to ensure the following intuitive behavior: removing any time stamps other than the successor from the candidate set should have no effect on the current successor. This is formalized as follows.

AXIOM 5 (THINNING). *Suppose $t_1 \in \textsc{succ}(t_0, \mathcal{F})$. Then for any $\mathcal{T} \subseteq \mathcal{F}$ with $t_1 \in \mathcal{T}$, $t_1 \in \textsc{succ}(t_0, \mathcal{T})$.*

This axiom also addresses the following issue. We know from Axiom 4 (RESPECTING ORDER) that $t_0 \prec t_1$ whenever there is *some* $\mathcal{F} \subseteq \mathbb{T}$ such that $t_1 \in \textsc{succ}(t_0, \mathcal{F})$. But this means we could have a temporal model that permits only singleton candidate sets (i.e. $\textsc{succ}(t, \mathcal{F}) = \emptyset$ if $|\mathcal{F}| > 1$). This would correspond to an event system that shuts down if it ever receives more than one future event. Clearly this is undesirable behavior. To prevent RESPECTING ORDER from degenerating as such, we need to be able to add and remove elements from the candidate sets in limited ways. THINNING addresses removal.

Adding new time stamps to a candidate set is much more subtle. Consider the intervals illustrated in Figure 5. Suppose we are trying to pick the successors of t, and start with the candidate set $\mathcal{F} = \{\, s \,\}$ (so trivially, s is the unique successor). If we extend $\mathcal{F}$ to the candidate set $\mathcal{F}' = \{\, r_1, s \,\}$, then the successor depends on our choice of event system. In EPL and the original Snoop, the time of an event is identified with the end of its interval, and so r_1 is the successor in this model. However, in all event systems with interval time stamps, the successor is still s, since the interval r_1 started before the end of t. Similarly, the effect of adding r_2 to $\mathcal{F} = \{\, s \,\}$ is also system dependent. In Active Office and Cayuga, the addition has no effect on the successor, as the interval ends later

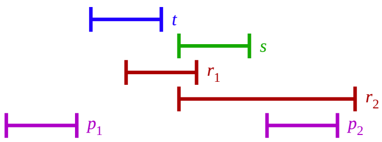

Figure 5: Adding Elements to a Candidate Set

than s. However, in SnoopIB, r_2 is also a successor, and hence this addition changes the contents of $\textsc{succ}(t, \mathcal{F})$.

Fortunately, all event systems agree that the addition of time stamps like p_1 or p_2 to $\mathcal{F}$ does not effect the value of $\textsc{succ}(t, \mathcal{F})$. These are time stamps that are either far in the past or far in the future, in that they do not overlap the time between t and any of its successors. Thus, in all event systems, we are permitted to add to a candidate set via the following axiom.

AXIOM 6 (THICKENING). *Let $\mathcal{A}$ be such that, for any $s \in \mathcal{A}$, either $s \prec t$ or $p \prec s$ for some $p \in \textsc{succ}(t, \mathcal{F})$. Then* $\textsc{succ}(t, \mathcal{F}) = \textsc{succ}(t, \mathcal{F} \cup \mathcal{A})$.

THICKENING is important for two reasons. First of all, in an event system the candidate set $\mathcal{F}$ is effectively infinite. It represents the time stamps of all the events that appear in the stream after t. An event system therefore never knows the full contents of $\mathcal{F}$ ahead of time; it only learns the values of these time stamps as they arrive. Hence, if we expect to have a real-time event processing system, the definition of successor cannot rely on such future events.

Additionally, as we mentioned before, we want to ensure that $\textsc{succ}(t, \mathcal{F})$ only chooses elements from $\mathcal{F}$ that are the $\prec$-successors of t. The following propositions demonstrate that these axioms are enough to guarantee this.

PROPOSITION 1. *For any $t_0, t_1 \in \mathbb{T}$, $t_0 \prec t_1$ if and only if $t_1 \in \textsc{succ}(t_0, \{ t_1 \})$.*

PROOF. By Axiom 4 (RESPECTING ORDER), $t_0 \prec t_1$ if and only there is some $\mathcal{F}$ such that $t_1 \in \textsc{succ}(t_0, \mathcal{F})$. By Axiom 5 (THINNING), we can choose $\mathcal{F} = \{ t_1 \}$. $\square$

PROPOSITION 2. *If $t_1 \in \textsc{succ}(t_0, \mathcal{F})$, then $t_0 \prec t_1$ and there is no $s \in \mathcal{F}$ with $t_0 \prec s \prec t_1$.*

PROOF. Suppose that there is such an s. As $s \in \mathcal{F}$, we get $t_1 \in \textsc{succ}(t_0, \{ s, t_1 \})$ from Axiom 5 (THINNING). As $t_0 \prec s$, $\{ s \} = \textsc{succ}(t_0, \{ s \})$ by Proposition 1. Hence $\{ s \} = \textsc{succ}(t_0, \{ s, t_1 \})$, by Axiom 6 (THICKENING), a contradiction. $\square$

A related issue is the problem of *blocking*. By Axiom 4 (RESPECTING ORDER), we know that r_1 in Figure 5 can never be a successor of t. However, there is nothing to prevent us from saying that, since r_1 ends before s, it "blocks" s from being the successor of t, and hence $\textsc{succ}(t, \{ s, r_1 \}) = \emptyset$. In this case we have an element r_1 that is not the successor, which must be removed from the set in order to make s a successor. Again, this is behavior not found in any of the existing event systems.

AXIOM 7 (NON-BLOCKING). *If $\textsc{succ}(t, \mathcal{F}) \cap \mathcal{A} = \emptyset$, then* $\textsc{succ}(t, \mathcal{F}) = \textsc{succ}(t, \mathcal{F} \setminus \mathcal{A})$.

3.1.3 The $\otimes$ Axioms

The $\otimes$ operator is used to combine time stamps from sequenced events. Hence our first axiom is concerned with when sequencing is defined. In particular, $t_0 \otimes t_1$ should only be defined if the t_i are the time stamps to two events that can be sequenced.

AXIOM 8 (CONSERVATIVE COMPOSITION). *$t_0 \otimes t_1$ is defined if and only if $t_0 \prec t_1$.*

Because event systems must process events in real-time, event sequencing should happen in a "timely" manner. In other words, the sequenced event should have a time stamp that allows us to add it to the output stream immediately. For example, suppose we compose two events with time stamps $t_0 = [0, 1]$ and $t_1 = [2, 3]$. We should not allow $t_0 \otimes t_1 = [0, 5]$ as the interval $[4, 4]$ follows t_1, but not $t_0 \otimes t_1$; hence we could not add $t_0 \otimes t_1$ to the stream until we are sure that all events with time $[4, 4]$ have passed. This constraint is implemented in all event systems by ensuring that $t_0 \otimes t_1$ and t_1 always share the same successors.

AXIOM 9 ($\otimes$-ELIMINATION). *Suppose $t_0 \prec t_1$. Then $t_2 \in \textsc{succ}(t_0 \otimes t_1, \mathcal{F})$ if and only if $t_2 \in \textsc{succ}(t_1, \mathcal{F})$.*

3.1.4 The $\mathbb{T}^\circ$ Axioms

The operation $\otimes$ is used to construct time stamps created by the sequencing operation. Intuitively all time stamps should ultimately be derived via $\otimes$ from some universe of "base" time stamps, e.g., the universe of clock ticks that define some event occurrence. We will refer to this set of base time stamps as $\mathbb{T}^\circ$. For example, in all event systems discussed in this paper, time stamps are defined over a universe that is isomorphic to $\mathbb{Z}$ (global clock, discrete time).

AXIOM 10 (PRIMITIVE REPRESENTATION). *There is a set $\mathbb{T}^\circ \subseteq \mathbb{T}$ such that*

- *for any $s \in \mathbb{T}^\circ$, there is no $t_0, t_1 \in \mathbb{T}$ with $s = t_0 \otimes t_1$.*
- *for any $t \in \mathbb{T}$, there are $s_i \in \mathbb{T}^\circ$ such that $t = s_0 \otimes \cdots \otimes s_n$.*

PROPOSITION 3. *The set $\mathbb{T}^\circ$ in* PRIMITIVE REPRESENTATION *is unique.*

PROOF. Let $\mathbb{T}_1$ and $\mathbb{T}_2$ both satisfy the properties of $\mathbb{T}^\circ$ in Axiom 10 (PRIMITIVE REPRESENTATION), and suppose $\mathbb{T}_1 \neq \mathbb{T}_2$. Without loss of generality, there is some $s \in \mathbb{T}_1$ with $s \notin \mathbb{T}_2$. As $s \notin \mathbb{T}_2$, there is some $p_0, \ldots, p_n \in \mathbb{T}_2$, $n \geq 1$ such that $s = p_0 \otimes (p_1 \otimes \cdots \otimes p_n)$. However, as $s \in \mathbb{T}_1$, this is a contradiction. $\square$

In essence, PRIMITIVE REPRESENTATION asserts that $\mathbb{T}$ is a free monoid with respect to $\otimes$ over $\mathbb{T}^\circ$. Note that this axiom only says that base time stamps exist, and does not require them to be points, intervals, or anything in particular. Furthermore, the decomposition in PRIMITIVE REPRESENTATION need not be unique. For example, in Active Office, $[1, 3] = [1, 1] \otimes [3, 3] = [1, 1] \otimes [2, 2] \otimes [3, 3]$.

However, all existing event systems have a global clock and all time stamps are defined in terms of values of this clock. Thus there is an implicit linear order on the base time stamps. Notice that this does not imply that *event* time stamps, including those of primitive events, are linearly ordered. For example, even though the natural numbers are linearly ordered, intervals of natural numbers can overlap and hence are only partially ordered (as pointed out earlier). The underlying assumption of a global clock is formalized by asserting that $\mathbb{T}^\circ$ is isomorphic to the linear order $\mathbb{Z}$.

AXIOM 11 (LINEARITY). *Let $\mathbb{T}^\circ$ be the unique set identified in* PRIMITIVE REPRESENTATION. *The ordering $\prec$ is an infinite discrete linear ordering on $\mathbb{T}^\circ$.*

From this axiom it may appear that we cannot handle real-valued time stamps. However, we can remove the discrete requirement from LINEARITY provided that we stipulate that all candidate sets are well-founded. If we had a non-well-founded candidate set $\mathcal{F}$ with an infinite descending sequence converging to t, then $\textsc{succ}(t, \mathcal{F})$ would not be well-defined, even though there are elements in $\mathcal{F}$ after t. As $\mathcal{F}$ corresponds to a set of time stamps for incoming events, well-foundedness is a realistic assumption. Furthermore, as all event expressions are finite, there is no distinguishable difference between requiring that $\mathbb{T}^\circ$ be discrete and requiring that all $\mathcal{F}$ be well-founded. Therefore, for simplicity, we keep the discreteness assumption.

3.2 Desirable Axioms

All of the axioms in the previous section are satisfied by the existing event systems. However, there are several axioms that we would like our models to satisfy for implementation reasons. In this section we will introduce these axioms.

3.2.1 The "Time-Out" Axiom

In Section 1, we saw an important problem that occurs in the SnoopIB system. In SnoopIB, overlapping pairs of events in Query 1 can result in an infinite number of matches for each Slashdot posting. All we need is for each blog in existence to post one link immediately, and then some link in the future. The second link can be posted in an hour, a day, or even years from now; as all these event pairs overlap, they are all successors to the Slashdot posting.

For a more formal illustration of this problem, in SnoopIB,

$$\textsc{succ}([0,0], \{\, [1,x] \mid 1 \leq x \,\}) = \{\, [1,x] \mid 1 \leq x \,\} \qquad (1)$$

Hence this definition of successor is very difficult to implement in an event system. Even though the time stamps may be partially ordered, the events necessarily arrive real-time in a linear fashion. In models with interval time stamps, they typically arrive to the stream at the time corresponding to the end of the interval. (This is the time when the event "happens".) Hence for the candidate set $\mathcal{F} = \{\, [1,1], [2,2], [1,3] \,\}$, $[2,2]$ will arrive before an event with time stamp $[1,3]$, even though $[1,3]$ is a successor time stamp to $[0,0]$ and $[2,2]$ is not.

This is particularly troubling as most high-performance event systems [5, 7, 12] use automata to process events. Suppose we have an automaton to recognize sequenced events of the form $E_1 \mathbin{;} E_2$, and suppose our automaton recognizes a match to E_1 at time $[0,0]$. Assume the next matching events for E_2 have times $[1,5]$, $[6,6]$, $[2,7]$, and $[7,8]$. After seeing the first match at time $[1,5]$, the automaton needs to remember the value 5 so that it can eliminate any event whose start time is afterwards, such as $[6,6]$ or $[7,8]$. Furthermore, it must remember this value for *each* new match to E_1.

One solution to this problem, adopted by Cayuga, is to spawn a new instance of the automaton processing $E_1 \mathbin{;} E_2$ for each match to E_1. However, in this case (1) demonstrates that an instance can never be garbage-collected in SnoopIB. There could always be an event with interval $[1,x]$ arriving at some future time x. As a result, memory usage grows without bound. Notice that any implementation, not only automaton-based ones, suffer from this issue. Therefore, we need an axiom that limits the effect that events with arbitrary long duration can have on the system.

AXIOM 12 (STRONG THICKENING). *Let* $t \in \textsc{succ}(s, A)$. *For any* $u, v \in \mathbb{T}$, *if* $t \prec v$, *then* $\textsc{succ}(s, A) = \textsc{succ}(s, A \cup \{\, u \otimes v \,\})$.

In essence, STRONG THICKENING is a "time-out" axiom. It guarantees that once we see a successor, we can ignore any events that happen afterwards.

3.2.2 Associativity

In Section 1.1.1 we saw that it would be advantageous for us to associate event sequencing. So, our next desirable axiom is one that guarantees associativity. Naively, it would seem to be enough for us to require that $\otimes$ is associative.

AXIOM 13 ($\otimes$-ASSOCIATIVITY). $(t_0 \otimes t_1) \otimes t_2 = t_0 \otimes (t_1 \otimes t_2)$, *for all* $t_i \in \mathbb{T}$.

However, this axiom is satisfied by both SnoopIB and Cayuga, which we have already seen are not associative. In fact, the only systems that violate this axiom are the point models of Snoop and EPL. Recall that we denote $t_0 \otimes t_1 = \perp$ if $t_0 \otimes t_1$ is undefined. So $\otimes$-ASSOCIATIVITY implicitly guarantees that $(t_0 \otimes t_1) \otimes t_2$ is defined exactly when $t_0 \otimes (t_1 \otimes t_2)$ is. In Snoop and EPL, the time stamp $2 \otimes (1 \otimes 3) = 3$ is defined, but $(2 \otimes 1) \otimes 3$ is not. In fact, this issue is the reason for the observation from [2, 6] that the two sequencings

$$E_1 \mathbin{;} (E_2 \mathbin{;} E_3) \quad \text{and} \quad E_2 \mathbin{;} (E_1 \mathbin{;} E_3) \qquad (2)$$

are equivalent for event systems with point time stamps. So while $\otimes$-ASSOCIATIVITY does not give us sequencing associativity, it is important in that it prevents us from sequencing events that should not be sequenced. In fact, we can express this observation as the following proposition.

PROPOSITION 4. $(t_0 \otimes t_1) \otimes t_2$ *is defined if and only if* $t_0 \prec t_1 \prec t_2$.

PROOF. Suppose $(t_0 \otimes t_1) \otimes t_2$ is defined. Then $t_0 \prec t_1$ by Axiom 8 (CONSERVATIVE COMPOSITION). Furthermore, by Axiom 13 ($\otimes$-ASSOCIATIVITY), we know that $t_0 \otimes (t_1 \otimes t_2)$ is defined and hence $t_1 \prec t_2$.

Now suppose $t_0 \prec t_1 \prec t_2$. By CONSERVATIVE COMPOSITION, $t_0 \otimes t_1$ is defined, and hence $t_0 \otimes t_1 \prec t_2$ by Axiom 9 ($\otimes$-ELIMINATION). Therefore, $(t_0 \otimes t_1) \otimes t_2$ is defined. $\square$

In order to find the correct axiom for associativity, we first need to formally understand what it means for sequencing to be associative. An event system processes expressions on a stream S of events. Events in a data stream consist of both data fields (which define the type of the event) and a time stamp. We typically denote these elements $\langle a, t \rangle \in D \times \mathbb{T}$ where $a \in D$ is the data and $t \in \mathbb{T}$ is the time stamp. As it is not relevant to our discussion, we make no stipulation on the nature of the data domain D. In traditional event systems D is the finite set of all event symbols, while in parametrized event systems such as Cayuga, D can be an infinite set of data tuples.

Given an event expression E, an event system returns $[\![E]\!]_S$, the set of all events in S that match E. For the sequencing operator, this set is defined as

$$[\![E_1 \mathbin{;} E_2]\!]_S = \left\{ \langle a_1 \oplus a_2, t_1 \otimes t_2 \rangle \;\middle|\; \begin{array}{l} \langle a_1, t_1 \rangle \in [\![E_1]\!]_S, \langle a_2, t_2 \rangle \in [\![E_2]\!]_S, \\ t_2 \in \textsc{succ}(t_1, \{s \mid \langle b, s \rangle \in [\![E_2]\!]_S\}) \end{array} \right\} \qquad (3)$$

Note that the data domain of the complex event $E_1 \mathbin{;} E_2$ is the set $\{a \oplus b | a \in D_1, b \in D_2\}$, where $a \oplus b$ is some data composition of data values a and b. We give no semantics for this operation as it will not be relevant to the discussion; in practice it is usually tuple concatenation.

An important part of this definition is that the time stamp composition ignores the data elements of the events. Otherwise, we cannot express associativity solely in terms of our temporal models. Some event systems do use data elements in their sequencing definition. For example, Active Office uses the element ID of an event to break ties when determining the successor. In those systems, we assume that the relevant ID information is added as part of the time stamp, hence sequencing is determined only from the time stamp. As data can be real-valued, this assumption does require our models to support real-valued time stamps. However, as we discussed in Section 3.1.4, this is not a problem.

From the definition of sequencing in (3), associativity requires that for any event expressions E_0, E_1, E_2 and stream S,

$$\begin{aligned} &[\![(E_0\,;E_1)\,;E_2]\!]_S \\ &= \left\{\left\langle \begin{matrix}(a_0\oplus a_1)\oplus a_2,\\ (t_0\otimes t_1)\otimes t_2\end{matrix}\right\rangle \,\middle|\, \begin{matrix}\langle a_i,t_i\rangle\in[\![E_i]\!]_S, t_1\in\text{SUCC}(t_0,\mathcal{F}_{E_1}),\\ t_2\in\text{SUCC}(t_0\otimes t_1,\mathcal{F}_{E_2})\end{matrix}\right\} \\ &= \left\{\left\langle \begin{matrix}a_0\oplus(a_1\oplus a_2),\\ t_0\otimes(t_1\otimes t_2)\end{matrix}\right\rangle \,\middle|\, \begin{matrix}\langle a_i,t_i\rangle\in[\![E_i]\!]_S, t_2\in\text{SUCC}(t_1,\mathcal{F}_{E_2}),\\ t_1\otimes t_2\in\text{SUCC}(t_0,\mathcal{F}_{E_1\,;E_2})\end{matrix}\right\} \\ &= [\![E_0\,;(E_1\,;E_2)]\!]_S \end{aligned} \tag{4}$$

where $\mathcal{F}_E = \{t\,|\,\langle a,t\rangle \in [\![E]\!]_S\}$. Note that this equation entails a relationship between $\mathcal{F}_{E_i}$ and $\mathcal{F}_{E_1\,;E_2}$. For candidate sets $\mathcal{F}_0$, $\mathcal{F}_1$, we define

$$\mathcal{F}_0\,;\mathcal{F}_1 = \{\, t_0\otimes t_1 \mid t_0\in\mathcal{F}_0, t_1\in\text{SUCC}(t_0,\mathcal{F}_1)\,\} \tag{5}$$

It should be clear then that (4) implies the following axiom.

AXIOM 14 ($\otimes$-DECOMPOSITION). *Suppose $t_0, t_1, t_2 \in \mathbb{T}$, with $t_1 \prec t_2$, and $\mathcal{F}_1, \mathcal{F}_2 \subseteq \mathbb{T}$. Also suppose that $t_2 \in \text{SUCC}(t_1, \mathcal{F}_2)$. Then $t_1 \in \text{SUCC}(t_0, \mathcal{F}_1)$ if and only if $t_1 \otimes t_2 \in \text{SUCC}(t_0, \mathcal{F}_1\,;\mathcal{F}_2)$.*

It is apparent from (4) that both $\otimes$-ASSOCIATIVITY and $\otimes$-DECOMPOSITION are necessary for associativity. The following proposition establishes that they are sufficient as well.

PROPOSITION 5. *Suppose $\mathbb{T}$ is a temporal model satisfying Axiom 9 ($\otimes$-ELIMINATION), Axiom 13 ($\otimes$-ASSOCIATIVITY), and Axiom 14 ($\otimes$-DECOMPOSITION). Let E_1, E_2, E_3 be event expressions, and suppose $\oplus$ is associative over the data elements of the event stream S. Then $[\![(E_0\,;E_1)\,;E_2]\!]_S = [\![E_0\,;(E_1\,;E_2)]\!]_S$.*

PROOF. Suppose that

$$\langle (a_0\oplus a_1)\oplus a_2, (t_0\otimes t_1)\otimes t_2\rangle \in [\![(E_0\,;E_1)\,;E_2]\!]_S$$

with $\langle a_i, t_i\rangle \in [\![E_i]\!]_S$. As $\oplus$ is associative, and $\otimes$ is associative by Axiom 13 ($\otimes$-ASSOCIATIVITY)

$$\begin{aligned}&\langle (a_0\oplus a_1)\oplus a_2, (t_0\otimes t_1)\otimes t_2\rangle \\ &\quad= \langle a_0\oplus(a_1\oplus a_2), t_0\otimes(t_1\otimes t_2)\rangle\end{aligned} \tag{6}$$

Furthermore, $t_1 \in \text{SUCC}(t_0, \mathcal{F}_{E_1})$ and $t_2 \in \text{SUCC}(t_0\otimes t_1, \mathcal{F}_{E_2})$. By Axiom 9 ($\otimes$-ELIMINATION), we have $t_2 \in \text{SUCC}(t_1, \mathcal{F}_{E_2})$. Hence $t_1\otimes t_2 \in \text{SUCC}(t_0, \mathcal{F}_{E_1\,;E_2})$ by Axiom 14 ($\otimes$-DECOMPOSITION), and thus

$$\langle (a_0\oplus a_1)\oplus a_2, (t_0\otimes t_1)\otimes t_2\rangle \in [\![E_0\,;(E_1\,;E_2)]\!]_S$$

Now suppose

$$\langle a_0\oplus(a_1\oplus a_2), t_0\otimes(t_1\otimes t_2)\rangle \in [\![E_0\,;(E_1\,;E_2)]\!]_S$$

Again, $\otimes$-ASSOCIATIVITY gives us (6). Furthermore, $t_2 \in \text{SUCC}(t_1, \mathcal{F}_{E_2})$ and $t_1\otimes t_2 \in \text{SUCC}(t_0, \mathcal{F}_{E_1\,;E_2})$. So $t_1 \in \text{SUCC}(t_0, \mathcal{F}_{E_1})$ by $\otimes$-DECOMPOSITION. □

4. IDENTIFYING TEMPORAL MODELS

Now that we have stated our axioms, we would like to find the "best" model that satisfies all of these axioms. Note that none of the definitions of successor in Section 1 satisfy all of them. Cayuga and Active Office violate $\otimes$-DECOMPOSITION, and hence do not support associativity. SnoopIB and ODE also violate this axiom, and in addition violate STRONG THICKENING. Systems with point time stamps even violate Axiom 13 ($\otimes$-ASSOCIATIVITY). Hence to satisfy all of the axioms, we need to find a new temporal model for event systems.

In this section we characterize the models that satisfy all of our axioms up to isomorphism. We also identify the trade-offs that the models in Section 1 make by violating one or more of the desirable axioms.

4.1 Satisfying All Axioms

There is at least one model that satisfies all of the axioms. That is the complete-history model from Section 2.1; we leave verification of this fact as an exercise for the reader. Unfortunately, this particular model is impractical because of its memory requirements. In any event system, each base time stamp (i.e., an element of $\mathbb{T}^\circ$) requires a memory word. A complete history of time stamps for a composite event would require as many words as there are primitive events that form the composite event. This is particularly bad for queries in which the history can grow without bound. In addition to regular sequencing, all of the major event systems have an iterated sequencing operator, similar to Kleene-*. This operator is illustrated by the following blog query.

Query 2. Send me a sequence of links to blog postings, in which the first posting is a reference to a file on a sensitive site, and each later posting has a link to the previous.

This sequence can be composed of any number of blog postings. In the complete-history model, we have to store and remember the time stamps for all of the postings in the sequence.

To get a model that uses bounded memory for time stamps, we need to compress the time stamp representation. In other words, we want some fixed n such that every $t \in \mathbb{T}$ can be written $t = p_0\otimes\cdots\otimes p_n$ for $p_i \in \mathbb{T}^\circ$. Unfortunately, as the following theorem demonstrates, this is impossible.

THEOREM 1. *Assume $\mathbb{T}$ is a temporal model satisfying Axioms 1-14. For each $t \in \mathbb{T}$, there is a unique sequence $p_0, \ldots, p_n \in \mathbb{T}^\circ$ with $t = p_0\otimes\cdots\otimes p_n$, where n depends upon t.*

This theorem demonstrates that any temporal model that satisfies all of the axioms must keep a complete history of the time stamps. Intuitively this is the case because any time stamp in the history can be used to determine its order with respect to another history. From this theorem, we can prove an even stronger result, namely that complete-history is the *only* model of the axioms, up to isomorphism.

THEOREM 2. *Let $\mathbb{T}$ be a temporal model satisfying Axioms 1-14. Let $\mathbb{S}$ be the complete-history model. If we identify $\mathbb{T}^\circ$ with $\mathbb{Z}$, the mapping $t_0\otimes\cdots\otimes t_n \mapsto \sigma$ where $\sigma(i) = t_i$ is an isomorphism.*

The remainder of this section is the proof of these two theorems.

4.1.1 *Proof of Theorem 1*

To prove Theorem 1, we will assume from here on that $\mathbb{T}$ is a temporal model satisfying all of the axioms (Axioms 1-14). Before we prove Theorem 1, we first need a way of distinguishing time stamps. To do this, we introduce two equivalence relations.

Definition 1. For any $t_0, t_1 \in \mathbb{T}$, we say that t_0, t_1 have the *same end time* (denoted $t_0 \sim_E t_1$) when, for any $s \in \mathbb{T}$, $t_0 \prec s$ if and only if $t_1 \prec s$. Similarly, t_0, t_1 have the *same start time* (denoted $t_0 \sim_S t_1$) when, for any $s \in \mathbb{T}$, $s \prec t_0$ if and only if $s \prec t_1$.

Intuitively, these relations give us an abstract way to identify the start and end time of a time stamp without having to assume our time stamps are actually intervals. The following propositions below guarantee that every time stamp t has a unique start time $t \sim_S s_0 \in \mathbb{T}^\circ$, and a unique end time $t \sim_E s_1 \in \mathbb{T}^\circ$. Thus we can unambiguously speak of a time stamp "interval" in an abstract sense.

PROPOSITION 6. *Suppose $t_0 \prec t_1$. Then $t_0 \otimes t_1 \sim_E t_1$ and $t_0 \otimes t_1 \sim_S t_0$.*

PROOF. $t_0 \otimes t_1 \sim_E t_1$ is immediate from Axiom 9 ($\otimes$-ELIMINATION), so we need only prove $t_0 \otimes t_1 \sim_S t_0$. First suppose $s \prec t_0 \otimes t_1$. By Axiom 8 (CONSERVATIVE COMPOSITION), $s \otimes (t_0 \otimes t_1)$ is defined. Thus $s \prec t_0$ by Axiom 13 ($\otimes$-ASSOCIATIVITY) and Proposition 4. Now suppose $s \prec t_0$. By Proposition 4, $(s \otimes t_0) \otimes t_1$ is defined. So $s \otimes (t_0 \otimes t_1)$ is defined by $\otimes$-ASSOCIATIVITY, and hence $s \prec t_0 \otimes t_1$. $\square$

PROPOSITION 7. *Suppose $p_0, p_1 \in \mathbb{T}^\circ$. Then $p_0 = p_1$ if and only if $p_0 \sim_E p_1$. Similarly, $p_0 = p_1$ if and only if $p_0 \sim_S p_1$*

PROOF. If $t_0 = t_1$ then $t_0 \sim_E t_1$ is clear. Suppose then that $t_0 \sim_E t_1$ but $t_0 \neq t_1$. By Axiom 11 (LINEARITY) we can assume $t_0 \prec t_1$ without loss of generality. But as $t_0 \sim_E t_1, t_1 \prec t_1$, which contradicts Axiom 2 (IRREFLEXIVITY).

The proof for $\sim_S$ is analogous. $\square$

To prove Theorem 1, we will need to induct over the length of a decomposition $t = p_0 \otimes \cdots \otimes p_n$ of t. We can reduce a time stamp to one with smaller decomposition length by using Axiom 14 ($\otimes$-DECOMPOSITION). However, in order to make use of this axiom, we need to understand what happens when we apply SUCC twice. Proposition 8 tells us that all of the successors have the same end time. Hence by Axiom 9 ($\otimes$-ELIMINATION), if t_0, t_1 are both successors of s from the same candidate set, $\text{SUCC}(t_0, \mathcal{F}) = \text{SUCC}(t_1, \mathcal{F})$.

PROPOSITION 8. *Suppose $t_0, t_1 \in \text{SUCC}(s, \mathcal{F})$. Then $t_0 \sim_E t_1$.*

PROOF. Suppose that $t_0, t_1 \in \text{SUCC}(s, \mathcal{F})$ with $t_0 \not\sim_E t_1$. As elements of $\mathbb{T}$ are built up from $\mathbb{T}^\circ$, we know from Proposition 6 that there are $p_0, p_1 \in \mathbb{T}^\circ$ with $p_i \sim_E t_i$. As $\sim_E$ is an equivalence relation, $p_0 \not\sim_E p_1$. So from Axiom 13 ($\otimes$-ASSOCIATIVITY) and Axiom 11 (LINEARITY), we can assume without loss of generality that $p_0 \prec p_1$. Thus as $p_0 \sim_D t_0$, we have $p_0 \prec t_1$. We now consider two cases.

First, suppose $t_1 \in \mathbb{T}^\circ$. In this case $t_1 = p_1$ and so $t_0 \prec t_1$. As $t_0, t_1 \in \text{SUCC}(s, \mathcal{F})$, $\{ t_0, t_1 \} = \text{SUCC}(s, \{ t_0, t_1 \})$ by Axiom 3 (CANDIDATE PRESENCE) and Axiom 5 (THINNING). Then, again by these two axioms, $\{ t_0 \} = \text{SUCC}(s, \{ t_0 \})$. But $t_0 \prec t_1$, and this so this contradicts Axiom 6 (THICKENING).

Now suppose $t_1 \notin \mathbb{T}^\circ$. We write $t_1 = v_0 \otimes \cdots \otimes v_m \otimes p_1$ where $v_i \in \mathbb{T}^\circ$. Again by CANDIDATE PRESENCE and THINNING we have that $\{ t_0, t_1 \} = \text{SUCC}(s, \{ t_0, t_1 \})$ and $\{ t_0 \} = \text{SUCC}(s, \{ t_0 \})$. But as $t_0 \prec p_1$, Axiom 12 (STRONG THICKENING) gives us $\{ t_0 \} = \text{SUCC}(s, \{ t_0, t_1 \})$, a contradiction. $\square$

COROLLARY 1. *Suppose $t_0, t_1 \in \text{SUCC}(s, \mathcal{F})$ with $t_0, t_1 \in \mathbb{T}^\circ$. Then $t_0 = t_1$.*

PROOF. Apply Proposition 7 to Proposition 8. $\square$

While these propositions appear fairly technical, they are enough to prove that every time stamp has at most one successor. This is such a powerful result that we state it as a theorem in its own right.

THEOREM 3. *Let $s, t_0, t_1 \in \mathbb{T}$. For all $\mathcal{F} \subseteq \mathbb{T}$, if $t_0, t_1 \in \text{SUCC}(s, \mathcal{F})$, then $t_0 = t_1$.*

PROOF. Using Axiom 10 (PRIMITIVE REPRESENTATION), suppose $t_0 = u_0 \otimes \cdots \otimes u_n$, $t_1 = v_0 \otimes \cdots \otimes v_m$, where $u_i, v_j \in \mathbb{T}^\circ$. We proceed by induction on n and m. The case for $n, m = 1$ is covered by Corollary 1. Suppose we know it is true for any $n, m \leq k$, and take some t_0, t_1 with $n, m \leq k + 1$. Without loss of generality, $m = k + 1$.

First we consider the case for $n > 1$. By Propositions 7 and 8, we have that $u_n = v_m$. Let $q = u_n$, and let $p_0 = u_0 \otimes \cdots \otimes u_{n-1}$ and $p_1 = v_0 \otimes \cdots \otimes v_{m-1}$. Hence $t_0 = p_0 \otimes q$, $t_1 = p_1 \otimes q$. By Axiom 5 (THINNING), $p_0 \otimes q, p_1 \otimes q \in \text{SUCC}(s, \{ p_0 \otimes q, p_1 \otimes q \})$. Hence $p_0, p_1 \in \text{SUCC}(s, \{ p_0, p_1 \})$ and $q \in \text{SUCC}(p_i, \{ q \})$ by Axiom 14 ($\otimes$-DECOMPOSITION). Then $p_0 = p_1$ by our induction hypothesis, and so we are done.

Now suppose $t_0 \in \mathbb{T}^\circ$. By Propositions 8 and 7, $v_m = t_0$. Now let $r \prec s$. Then $s \in \text{SUCC}(r, \{ s \})$. As $\text{SUCC}(s, \{ t_0, t_1 \}) = \{ t_0, t_1 \}$, $\otimes$-DECOMPOSITION gives us

$$\begin{aligned}\{ s \otimes t_0, s \otimes v_0 \otimes \cdots \otimes v_{m-1} \otimes t_0 \} \\ = \text{SUCC}(r, \{ s \otimes t_0, s \otimes v_0 \otimes \cdots \otimes t_0 \})\end{aligned}$$

Again by $\otimes$-DECOMPOSITION, we have that $s, s \otimes \cdots \otimes v_{m-1} \in \text{SUCC}(r, \{ s, s \otimes \cdots \otimes v_{m-1} \})$. However, $s \prec v_0 \prec v_{m-1}$, and so this case contradicts Proposition 8. $\square$

To prove Theorem 1, we need one more result. Proposition 9 establishes that a single usage of $\otimes$ cannot collapse two different time stamps into a single time stamp.

PROPOSITION 9. *Suppose $t_0, t_1, s \in \mathbb{T}$ with $t_0, t_1 \prec s$ and $t_0 \neq t_1$. Then $t_0 \otimes s \neq t_1 \otimes s$*

PROOF. By Proposition 1, $s \in \text{SUCC}(t_i, \{ s \})$ for each i. Suppose for a contradiction that $t_0 \otimes s = t_1 \otimes s$. Let $r \in \mathbb{T}^\circ$ be such that $r \prec t_0$. By Proposition 6, $r \prec t_0 \otimes s = t_1 \otimes s$. Hence

$$t_0 \otimes s = t_1 \otimes s \in \text{SUCC}(r, \{ t_0 \otimes s, t_1 \otimes s \})$$

By Axiom 14 ($\otimes$-DECOMPOSITION), $t_0, t_1 \in \text{SUCC}(r, \{ t_0, t_1 \})$. But this contradicts Theorem 3. $\square$

PROPOSITION 10. *Suppose $t_0, t_1 \in \mathbb{T}$ with $s_0, s_1 \in \mathbb{T}^\circ$ such that $t_0 \sim_E s_0$ and $t_1 \sim_S s_1$. Then $t_0 \prec t_1$ if and only if $s_0 \prec s_1$. In other words, t_1 follows t_0 if and only if the start time of t_1 follows the end time of t_0.*

PROOF. Suppose $t_0 \sim_E s_0$ and $t_1 \sim_S s_1$. First suppose that $t_0 \prec t_1$. As $t_0 \sim_E s_0$, $s_0 \prec t_1$ by the definition of $\sim_E$. Similarly, $s_0 \prec s_1$ as $t_1 \sim_S s_1$. The proof for when $s_0 \prec s_1$ is analogous. $\square$

PROOF OF THEOREM 1. Let $t_0 = u_0 \otimes \cdots \otimes u_n$, $t_1 = v_0 \otimes \cdots \otimes v_m$. Also suppose that $n \neq m$, or $n = m$ and $u_i \neq v_i$ for some $i \leq n$. We need to show that $t_0 \neq t_1$.We proceed by induction on n and m. The case for $n = m = 1$ is obvious. So suppose we know that the $t_0 \neq t_1$ for $n, m \leq k$. Let $m = k + 1$ and $n \leq m$.

First we consider the case where $n > 1$. Suppose for a contradiction that $t_0 = t_1$. Then by Proposition 6 and 7, $u_n = v_m = q$. Let $p_0 = u_0 \otimes \cdots \otimes u_{n-1}$ and $p_1 = v_0 \otimes \cdots \otimes v_{m-1}$. By Proposition 9, $p_0 = p_1$, contradicting our induction hypothesis.

Now suppose $n = 1$. Then $n \neq m$ and so $u_0 \neq v_0$. However, $t_0 \sim_S t_1$ and so this contradicts Proposition 7. $\square$

4.1.2 *Proof of Theorem 2*

Theorem 3 proves that there is at most one successor at any time. While this is true in complete-history, this is not enough to establish complete-history as the unique temporal model. We need to prove that this unique successor is structurally identical to the one in complete-history. In particular, we need to know that the partial order $\prec$ behaves just like the interval partial order $<$. This fact follows from the next proposition.

PROPOSITION 10. *Suppose $t_0, t_1 \in \mathbb{T}$ with $s_0, s_1 \in \mathbb{T}^\circ$ such that $t_0 \sim_E s_0$ and $t_1 \sim_S s_1$. Then $t_0 \prec t_1$ if and only if $s_0 \prec s_1$. In other words, t_1 follows t_0 if and only if the start time of t_1 follows the end time of t_0.*

PROOF. Suppose $t_0 \sim_E s_0$ and $t_1 \sim_S s_1$. First suppose that $t_0 \prec t_1$. As $t_0 \sim_E s_0$, $s_0 \prec t_1$ by the definition of $\sim_E$. Similarly, $s_0 \prec s_1$ as $t_1 \sim_S s_1$. The proof for when $s_0 \prec s_1$ is analogous. $\square$

Thus the only difference between any two models that satisfy all the axioms could lie in how they break ties between overlapping "intervals". We therefore need to establish that any two such models must break ties in the same way. Complete-history uses the linear order $\sqsubseteq$ to break ties. Because of Theorem 1, we can extend the definition of $\sqsubseteq$ to arbitrary temporal models in the usual way, identifying $t_0 \otimes \cdots \otimes t_n$ with σ as specified in Theorem 2. As the following proposition demonstrates, for small candidacy sets, $\sqsubseteq$ is our only option to choose a successor.

PROPOSITION 11. *Suppose $\textsc{succ}(s, \{ t_0, t_1 \}) = \{ t_0 \}$ with $t_i \in \mathbb{T}$. If $s \prec t_1$, then $t_0 \sqsubseteq t_1$.*

PROOF. Suppose that $\textsc{succ}(s, \{ t_0, t_1 \}) = \{ t_0 \}$ with $t_i \in \mathbb{T}$ such that $s \prec t_1$. Furthermore, suppose for a contradiction that $t_0 \not\sqsubseteq t_1$. As $\sqsubset$ is a linear order, $t_1 \sqsubset t_0$. We break our proof up into two cases: the case when $t_0 \not\sim_E t_1$ and the case when $t_0 \sim_E t_1$.

First consider the case $t_0 \not\sim_E t_1$. By Axiom 10 (PRIMITIVE REPRESENTATION) and Proposition 6, there are $p_0, p_1 \in \mathbb{T}^\circ$ with $p_i \sim_E t_i$. Then $p_0 \not\sim_E p_1$, and so from Axiom 13 ($\otimes$-ASSOCIATIVITY) and Axiom 11 (LINEARITY), either $p_0 \prec p_1$ or $p_1 \prec p_0$. As $t_1 \sqsubset t_0$, it is clear from Proposition 6 and the definition of $\sqsubseteq$ that $p_1 \prec p_0$. As $p_1 \sim_E t_1$, we have that $t_1 \prec p_0$. We now split into two further subcases.

First assume $t_0 \in \mathbb{T}^\circ$. In that case $t_0 = p_0$, and so $t_1 \prec t_0$. Thus we have $\textsc{succ}(s, \{ t_0, t_1 \}) = \{ t_1 \}$ by the arguments in the proof of Proposition 8.

Now assume $t_0 \notin \mathbb{T}^\circ$. By PRIMITIVE REPRESENTATION, we can write $t_0 = v_0 \otimes \cdots \otimes v_m \otimes p_0$. By Axiom 3 (CANDIDATE PRESENCE) and Axiom 5 (THINNING), we have that $\{ t_0, t_1 \} = \textsc{succ}(s, \{ t_0, t_1 \})$ and $\{ t_1 \} = \textsc{succ}(s, \{ t_1 \})$. But as $t_1 \prec p_0$, Axiom 12 (STRONG THICKENING) gives $\{ t_1 \} = \textsc{succ}(s, \{ t_0, t_1 \})$, a contradiction. So our proposition holds in the case $t_0 \not\sim_E t_1$.

Now we consider the case $t_0 \sim_E t_1$. We decompose $t_0 = v_0 \otimes \cdots \otimes v_m$, $t_1 = u_0 \otimes \cdots \otimes u_n$. As $t_0 \sim_E t_1$, $v_m = u_n$ by Proposition 6. Again we have two possibilities for our subcases.

The first possibility is that there is some $k > 0$ such that $v_{m-k} < u_{n-k}$ and $v_{m-i} < u_{n-i}$ for $i < k$. In that case $m, n > 0$, so we let $p_0 = v_0 \otimes \cdots \otimes v_{m-1}, p_1 = u_0 \otimes \cdots \otimes u_{n-1}$, and $q = v_m = u_n$. So $t_0 = p_0 \otimes q$ and $t_1 = p_1 \otimes q$. As $\textsc{succ}(s, \{ t_0, t_1 \}) = \{ t_0 \}$, we have $\{ p_0 \} = \textsc{succ}(s, \{ p_0, p_1 \})$ by Axiom 14 ($\otimes$-DECOMPOSITION). As $s \prec t_1$, Axiom 5 (THINNING) gives $\{ t_1 \} = \textsc{succ}(s, \{ t_1 \})$, and thus $s \prec p_1$ by $\otimes$-DECOMPOSITION. Therefore, $p_0 \sqsubseteq p_1$ by our induction hypothesis, and hence $t_0 \sqsubseteq t_1$.

The second possibility is that $m > n$ and $v_{m-i} = u_{n-i}$ for all $i \leq n$. This time we let $p = v_0 \otimes \cdots \otimes v_{m-n-1}$ and $q = v_{m-n} \otimes \cdots \otimes v_m$, and so $t_0 = p \otimes q$, $t_1 = q$. By LINEARITY, pick $r \prec s$. Then $s \in \textsc{succ}(r, \{ s \})$ and so $\otimes$-DECOMPOSITION gives

$$\{ s \otimes p \otimes q \} = \textsc{succ}(r, \{ s \otimes p \otimes q, s \otimes q \})$$

Then again by $\otimes$-DECOMPOSITION

$$\{ s \otimes p \} = \textsc{succ}(r, \{ s \otimes p, s \})$$

As $s \prec p$, $s \otimes p \not\subseteq p$ and $p \not\sim_D s$, which contradicts our proof of the case $t_0 \not\sim_E t_1$. $\square$

Theorem 3 guarantees that there is at most one successor, and Proposition 11 suggests that when we have a successor, we always use $\sqsubseteq$ to determine which one it is. Therefore, to prove Theorem 2, we only need to guarantee that, when there is some $s \in \mathcal{F}$ with $t \prec s$, then there is *at least* one element in $\textsc{succ}(t, \mathcal{F})$. Fortunately, this follows from Axiom 7 (NON-BLOCKING).

PROOF OF THEOREM 2. By Theorem 1, the mapping $t_0 \otimes \cdots \otimes t_n \mapsto \sigma$ is well-defined; it is clearly a bijection. We need to show that this mapping preserves the successor operation. We already know from Proposition 10 that $\prec$ and the interval order are the same. So we need only show that we break ties properly on all candidate sets.

Suppose $t_0 \in \textsc{succ}(s, \mathcal{F})$ with $s \prec t_1 \in \mathcal{F}$. Applying NON-BLOCKING to Proposition 11, we see that $t_0 \sqsubset t_1$ whenever $t_0 \neq t_1$. Thus $\sqsubseteq$ is the only way to break ties over arbitrary candidate sets. The only thing left to show is that $\textsc{succ}(s, \mathcal{F}) \neq \emptyset$ whenever $t \in \mathcal{F}$ with $s \prec t$. Suppose for a contradiction that $t \in \mathcal{F}$ with $s \prec t$, but $\textsc{succ}(s, \mathcal{F}) = \emptyset$. Then by NON-BLOCKING, $\textsc{succ}(s, \{ t \}) = \emptyset$. But this contradicts Proposition 1. $\square$

4.2 Relaxing the Desirable Axioms

The moral of Section 4.1 is that there is no efficient temporal model satisfying all of the axioms. Thus if we want an efficient temporal model, we need to relax our demands. This means that our primary goal now is to identify the least number of axioms that we need to relax in order to get an efficient temporal model. In order to answer this question, we first have to identify what we mean by an efficient temporal model.

Definition 2. An *interval model* is a model $\mathbb{T}$ in which

$$t_0 \otimes t_1 \otimes t_2 = t_0 \otimes t_2 \text{ for any } t_0, t_1, t_2 \in \mathbb{T} \tag{7}$$

An interval model allows us the most compact representation, as we only need to remember two primitive time stamps for each element of $\mathbb{T}^\circ$ (see Proposition 12 below). While this may seem like a fairly extreme restriction, our results in this section generalize for any model with bounded representation (i.e., there is some fixed n such that for each t, $t = p_0 \otimes \cdots \otimes p_n$ for some $p_i \in \mathbb{T}^\circ$). Thus we consider only interval models in order to simplify our analysis.

Because only the axioms in Section 3.2 are controversial, our analysis will only consider interval models that satisfy all of the accepted axioms (Axioms 1-11). Furthermore, of all the axioms in Section 3.2, we do not want to drop Axiom 13 ($\otimes$-ASSOCIATIVITY). That axiom is necessary to prevent the pathological behavior equating the two expressions in (2), which is clearly undesirable. Therefore, in this section we will determine what types of interval models we get if we relax either Axiom 12 (STRONG THICKENING) or Axiom 14 ($\otimes$-DECOMPOSITION).

As many of the of the propositions in Section 4.1 did not require the use of axioms in Section 3.2, we can still say a lot about these models. In particular, Propositions 6 and 7 require neither STRONG THICKENING nor $\otimes$-DECOMPOSITION, and so we have the following result, which we will use throughout the section.

PROPOSITION 12. *Let $\mathbb{T}$ be any interval model satisfying the accepted axioms and let $t \in \mathbb{T} \setminus \mathbb{T}^\circ$. There are unique $t_0, t_1 \in \mathbb{T}^\circ$ such that $t_0 \sim_S t$, $t_1 \sim_E t$. Furthermore, $t = t_0 \otimes t_1$.*

PROOF. By Axiom 10 (PRIMITIVE REPRESENTATION), we have that $t = v_0 \otimes \cdots \otimes v_n$ with $v_i \in \mathbb{T}^\circ$. By Proposition 6, $v_0 \sim_S t$ and $v_n \sim_E t$. Also as $\mathbb{T}$ is an interval model, $t = v_0 \otimes v_n$. Let $t_0 = v_0$, $t_1 = v_n$. We need only show that they are unique.

Suppose $t = s_0 \otimes s_1$ with $s_i \in \mathbb{T}^\circ$. By Proposition 6, $s_0 \sim_S t$ and hence $s_0 \sim_S t_0$. Thus $s_0 = t_0$ by Proposition 7. A similar argument shows that $s_1 = t_1$. $\square$

4.2.1 *Relaxing* STRONG THICKENING

STRONG THICKENING was an important part of the proof of Theorem 1, which prevents any model of the axioms from being an interval model. Therefore, we might suspect that we can get an associative interval model by relaxing this axiom. However, as we saw in Section 1.1.1, none of the existing interval models are associative. Furthermore, as the following theorem shows, there is no way to get an associative interval model of the accepted axioms.

THEOREM 4. *There is no interval model of the accepted axioms that is also associative.*

This theorem is true because any associative model must satisfy both Axiom 13 ($\otimes$-ASSOCIATIVITY) and Axiom 14 ($\otimes$-DECOMPOSITION). And any interval model of these two axioms can never have more than one successor. Suppose event E_1 has time stamp $[0, 0]$ and there are two instances of $(E_2\,;E_3)$ with time stamps $[1, 3]$ and $[2, 3]$, respectively. We cannot tell from the time stamp $[1, 3]$ whether E_2 had time stamp $[1, 1]$ or $[1, 2]$. So if we choose both the event at $[1, 3]$ and the one at $[2, 3]$ as the next occurrence of $(E_2\,;E_3)$, and the two E_2 events have time stamps $[1, 1]$ and $[2, 2]$, respectively, then we must choose both of them as the next E_2 event after E_1. However, this violates Axiom 6 (THICKENING), which is an accepted axiom.

As the following proposition shows, we can never limit the successor in an associative interval model to a single time stamp.

PROPOSITION 13. *Let $\mathbb{T}$ be any interval model of the accepted axioms which is associative. Let $t \in \mathbb{T}$ and let $\mathcal{F} \subseteq \mathbb{T}$ be such that $s_1 \sim_E s_2$ and $t \prec s_1$ for all $s_1, s_2 \in \mathcal{F}$. Then $\textsc{succ}(t, \mathcal{F}) = \mathcal{F}$.*

PROOF. By Proposition 12, there is some $t_0 \in \mathbb{T}^\circ$ with $t_0 \sim_S t$. Similarly, for each $s \in \mathcal{F}$ there is some $s_1 \in \mathbb{T}^\circ$ such that $s_1 \sim_E s$. Furthermore, as $p_1 \sim_E p_2$ for each $p_i \in \mathcal{F}$, by Proposition 7, there is a unique s_1 that works for all elements of $\mathcal{F}$. Now take any $p \in \mathcal{F}$. As $t \prec p$, $t \otimes p$ is defined. By Proposition 6, $t_0 \sim_S t \sim_S t \otimes p$ and $s_1 \sim_E p \sim_E t \otimes p$. Thus $t \otimes p = t_0 \otimes s_1$ for all $p \in \mathcal{F}$, and hence $\{\, t_0 \otimes s_1 \,\} = \{\, t \,\}\,; \mathcal{F}$.

By Axiom 11 (LINEARITY), there is some $r \in \mathbb{T}^\circ$ such that $r \prec t_0$. Hence $r \prec t_0 \otimes s_1$ by Proposition 6. Thus $\textsc{succ}(r, \{\, t_0 \otimes s_1 \,\}) = \{\, t_0 \otimes s_1 \,\}$ by Proposition 1. As $\{\, t_0 \otimes s_1 \,\} = \{\, t \,\}\,; \mathcal{F}$, $\textsc{succ}(t, \mathcal{F}) = \mathcal{F}$ by Axiom 14 ($\otimes$-DECOMPOSITION). $\square$

PROOF OF THEOREM 4. By Axiom 11 (LINEARITY), let $t_i \in \mathbb{T}^\circ$ with $0 \le i \le 3$ and $t_0 \prec t_1 \prec t_2 \prec t_3$. By Proposition 13,

$$\textsc{succ}(t_0, \{\, t_1 \otimes t_3, t_2 \otimes t_3 \,\}) = \{\, t_1 \otimes t_3, t_2 \otimes t_3 \,\}$$

Then by Axiom 14 ($\otimes$-DECOMPOSITION), $\textsc{succ}(t_0, \{\, t_1, t_2 \,\}) = \{\, t_1, t_2 \,\}$. However, $\textsc{succ}(t_0, \{\, t_1 \,\}) = \{\, t_1 \,\}$ and $t_1 \prec t_2$, which violates Axiom 6 (THICKENING). $\square$

As a result of this theorem, there is no obvious benefit for relaxing STRONG THICKENING.

4.2.2 *Relaxing* $\otimes$-DECOMPOSITION

Even though there is no hope for an associative interval model, we may still be able to construct an interval model that *approximates* associativity. All of the interval models in Section 1 satisfy $\otimes$-ASSOCIATIVITY. The only problem is how we treat the candidate sets of composite events. For full associativity, we require that $[\![(E_0\,;E_1)\,;E_2]\!]_S = [\![E_0\,;(E_1\,;E_2)]\!]_S$. Suppose instead that we have a model in which $[\![(E_0\,;E_1)\,;E_2]\!]_S \supseteq [\![E_0\,;(E_1\,;E_2)]\!]_S$. In such a model we could rewrite the expression $E_0\,;(E_1\,;E_2)$ as a left-associated expression, and eliminate the false positives in post-processing. Unfortunately, even this is impossible in an interval model.

THEOREM 5. *Let $\mathbb{T}$ be an interval model satisfying all axioms but $\otimes$-DECOMPOSITION. Then there are expressions E_i and a stream S such that $[\![(E_0\,;E_1)\,;E_2]\!]_S \not\supseteq [\![E_1\,;(E_1\,;E_2)]\!]_S$.*

PROOF SKETCH. Suppose for a contradiction that $[\![(E_0\,;E_1)\,;E_2]\!]_S \supseteq [\![E_0\,;(E_1\,;E_2)]\!]_S$ for any E_i and S. This means that we get the reverse direction of $\otimes$-DECOMPOSITION. In other words, for any t_0, t_1, t_2, and $\mathcal{F}_1, \mathcal{F}_2$,

$$\begin{aligned} t_1 \otimes t_2 \in \textsc{succ}(t_0, \mathcal{F}_1\,; \mathcal{F}_2),\ t_2 \in \textsc{succ}(t_1, \mathcal{F}_2) \\ \Rightarrow t_1 \in \textsc{succ}(t_0, \mathcal{F}_1) \end{aligned} \tag{8}$$

From (8), we can reproduce enough of the proof of Proposition 11 to show that we have to break ties by start time. We then use the forgetfulness of (7) to show that the decomposition over $\mathbb{T}^\circ$ is not unique, and hence our tie breaking is not well-defined. $\square$

It is also possible to approximate associativity when $[\![(E_0\,;E_1)\,;E_2]\!]_S \subseteq [\![E_0\,;(E_1\,;E_2)]\!]_S$. This approximation guarantees that we will never produce false positives if we rewrite $E_0\,;(E_1\,;E_2)$ as a left-associated expression. Satisfying this property requires the forward direction of $\otimes$-DECOMPOSITION, namely

$$\begin{aligned} t_1 \in \textsc{succ}(t_0, \mathcal{F}_1),\ t_2 \in \textsc{succ}(t_1, \mathcal{F}_2) \\ \Rightarrow t_1 \otimes t_2 \in \textsc{succ}(t_0, \mathcal{F}_1\,; \mathcal{F}_2) \end{aligned} \tag{9}$$

It is easy to verify that Cayuga has this property. Furthermore, any interval model with this property must accept almost all $\prec$-successor time stamps with the minimal end time, and thus is at most a minor variation of Cayuga.

PROPOSITION 14. *Let $\mathbb{T}$ be an interval model satisfying* (9) *and all axioms but $\otimes$-DECOMPOSITION. Let $\mathcal{F}$ be a candidate set in which $t = r_0 \otimes r_1 \otimes r_2$ for every $t \in \mathcal{F}$. Then*

$$\textsc{succ}(s, \mathcal{F}) = \{t | s \prec t \in \mathcal{F} \textit{ and end time } t \sim_E q \in \mathbb{T}^\circ \textit{ is least}\}$$

PROOF. The key idea of this proof is to use the interior element r_1 of each time stamp, together with Axiom 12 (STRONG THICKENING), to show that no time stamp can block another with the same end time. Given Proposition 10 and several other axioms, we can assume without loss of generality that all events in $\mathcal{F}$ have the same end time and follow s. We let $q \in \mathbb{T}^\circ$ be the unique end time of all these elements. Then by (7), every element of $\mathcal{F}$ can be expressed as $p_i \otimes q$ with $p_i \in \mathbb{T}^\circ$.

Let $\mathcal{F}' = \{\, p_i \mid p_i \otimes q \in \mathcal{F}, p_i \in \mathbb{T}^\circ \,\}$ be the set of start times of all these time stamps. As $\mathcal{F}'$ is linearly ordered, pick $p_0 \otimes q \in \mathcal{F}$ such that $p_0 \in \mathcal{F}'$ is least. By Axiom 6 (THICKENING), $p_0 \in \textsc{succ}(s, \mathcal{F}')$, and thus $p_0 \otimes q \in \textsc{succ}(s, \mathcal{F})$ by (9).

As every element in $\mathcal{F}$ has form $r_0 \otimes r_1 \otimes q$, there is some $r \in \mathbb{T}^\circ$ such that $p_i \prec r \prec q$ for all $p_i \in \mathcal{F}'$. Let r be the greatest such primitive time stamp. Take any $t = p_i \otimes q \in \mathcal{F}$, and define $\mathcal{F}_t = \{ p_j \otimes r \mid i \neq j \} \cup \{ p_i \}$. Note that $\mathcal{F}_t ; \{ q \} = \mathcal{F}$. By STRONG THICKENING, $p_i \in \textsc{succ}(s, \mathcal{F}_t)$, and so $t \in \textsc{succ}(s, \mathcal{F})$ by (9). $\square$

It is possible for an interval model satisfying (9) to have arbitrary behavior on time stamps of very short duration (i.e., the composition of one or two base time stamps), as they are too short for associativity to apply. However, in addition to (9), Cayuga also has a very weak form of associativity that applies when it is sequencing a stream of events with itself (i.e., an expression of the form $E_1 ;(E_2 ; E_2)$). In Cayuga, if there are no overlapping E_2 events in S, then $[\![E_1 ;(E_2 ; E_2)]\!]_S = [\![(E_1 ; E_2) ; E_2]\!]_S$. This property follows from a weaker version of $\otimes$-DECOMPOSITION, namely

$$\textsc{succ}(t, \{ p_i \otimes s \}_{i \in I}) = \textsc{succ}(t, \{ p_i \}_{i \in I}) ; \{ s \} \qquad (10)$$

This property, in addition to (9), uniquely characterizes Cayuga up to isomorphism, suggesting that this temporal model is the closest we can get to an associative interval model.

THEOREM 6. *Let $\mathbb{T}$ be an interval model satisfying* (9), (10) *and all axioms but* $\otimes$-DECOMPOSITION. *Let $\mathbb{S}$ be the Cayuga temporal model. If we identify $\mathbb{T}^\circ$ with $\mathbb{Z}$, the mapping $t_0 \otimes t_1 \mapsto [t_0, t_1] \in \mathbb{S}$, where $t_i \in \mathbb{T}^\circ$, is an isomorphism.*

PROOF. Proposition 12 guarantees that the mapping is a well-defined bijection. The proof of Proposition 8 does not require $\otimes$-DECOMPOSITION. Hence by Axiom 7 (NON-BLOCKING) and Axiom 12 (STRONG THICKENING), it is sufficient to show that $\textsc{succ}(t, \mathcal{F}) = \mathcal{F}$ for any $\mathcal{F}$ such that $p_1 \sim_E p_2$ and $t \prec p_1$ for all $p_i \in \mathcal{F}$.

Take any such t, $\mathcal{F}$. By Propositions 7 and 12, there is a unique s such that $p \sim_E s$ for all $p \in \mathcal{F}$. By Axiom 11 (LINEARITY), there is some $s \prec u \prec v$. By (10),

$$\textsc{succ}(t, \{ p_i \otimes (u \otimes v) \}_{p_i \in \mathcal{F}}) = \textsc{succ}(t, \mathcal{F}) ; \{ u \otimes v \}$$

Hence by Proposition 14,

$$\textsc{succ}(t, \{ p_i \otimes (u \otimes v) \}_{p_i \in \mathcal{F}}) = \{ p_i \otimes (u \otimes v) \}_{p_i \in \mathcal{F}}$$

By Proposition 12, each element of $\mathcal{F}$ has a unique start time. Thus $\textsc{succ}(t, \mathcal{F}) = \mathcal{F}$ $\square$

Proposition 14 demonstrates that Active Office does not approximate associativity in either direction. However, SnoopIB and ODE also satisfy (9) and (10), and thus approximate associativity equally well as Cayuga. Still, they do not satisfy STRONG THICKENING, and these results show that there is no gain from eliminating STRONG THICKENING. Thus there is no apparent advantage to adopting the temporal models of SnoopIB or ODE over Cayuga.

5. RELATED WORK

Initial implementations of event composition systems, such as Snoop [4] and EPL [11], used a linear temporal model based on detection times. Results from the Knowledge Representation community [2, 6] demonstrated that this temporal model did not correctly implement the semantics of sequencing in right-associated queries. Other attempts at event systems [1, 5, 7, 12] all use interval or history models. However, there has been no research into which definition of successor is most appropriate.

The work on EPL [11] is particularly notable as it provides a formal semantics for event languages. However, even though the language is well-defined, it still exhibits unusual behavior like equating the queries in (2). Instead of presenting yet another formal semantics, our work in this paper has been to determine criteria for evaluating and comparing alternate semantics.

The theory of temporal logic has covered many aspects of temporal models; an excellent survey can be found in van Benthem [13]. Bohlen et al [3] have examined the difference between point and interval models for time in database systems. Our temporal model is a general framework that includes all of these types of models, and many of our axioms in Section 3.1 were motivated by work in this area. To our knowledge, our paper is the first formulation of a temporal model that examines the definition of a successor operation different from the usual one defined by the partial order on time.

Kraemer and Seeger [8] have examined the difficulty of implementing a window join operation on streaming data with interval time stamps. However, their analysis only looks at implementing a specific temporal model, and is not an attempt to characterize all possible implementations, such as we have done in this paper.

6. CONCLUSIONS AND FUTURE WORK

While our approach has been motivated by practical implementation concerns, we have attempted to give a formal and rigorous analysis of the different ways in which we can define a sequencing operator in event composition systems. Admitting that two of the axioms in Section 3.2 are controversial, we have identified two canonical temporal models. One of the two models — complete-history — has serious implementation issues; in complete-history the time stamps are unbounded. The time stamp model that we introduce in Cayuga appears to be the best trade-off between ease of implementation and support of sequencing associativity and right-associated queries.

There are two axioms in Section 3.1 which, while accepted by all event composition systems, are controversial in the temporal logic community. In particular, while Axiom 11 (LINEARITY) is appropriate for synchronous event systems, it is not applicable to distributed event systems as initially studied by Lamport [9] and later by Liebig et al [10]. Future work is needed to determine the effect of removing this axiom from our framework.

An even more interesting solution to the synchronous assumption would be to remove both LINEARITY and Axiom 10 (PRIMITIVE REPRESENTATION). While the base time stamps are fundamental to our arguments, we can artificially construct them as equivalence classes over the relations $\sim_S$ and $\sim_E$. Further research is needed to determine what temporal models arise when we extend an existing model with these equivalence classes as time stamps.

7. REFERENCES

[1] R. Adaikkalavan and S. Chakravarthy. SnoopIB: Interval-based event specification and detection for active databases. In *Proc. ADBIS*, pages 190–204, 2003.

[2] J. Allen and G. Ferguson. Actions and events in interval temporal logic. Technical Report 521, University of Rochester, jul 1994.

[3] M. H. Bohlen, R. Busatto, and C. S. Jensen. Point-versus interval-based temporal data models. In *Proc. ICDE*, pages 192–200, 1998.

[4] S. Chakravarthy, V. Krishnaprasad, E. Anwar, and S.-K. Kim. Composite events for active databases: Semantics, contexts and detection. In *Proc. VLDB*, pages 606–617, 1994.

[5] A. Demers, J. Gehrke, M. Hong, M. Riedewald, and W. White. Towards expressive publish/subcribe systems. In *Proc. EDBT*, 2006. To appear.

[6] A. Galton and J. C. Augusto. Two approaches to event definition. In *Proc. DEXA*, pages 547–556, 2002.

[7] N. H. Gehani, H. V. Jagadish, and O. Shmueli. Composite event specification in active databases: Model and implementation. In *Proc. VLDB*, pages 327–338, 1992.

[8] J. Kraemer and B. Seeger. A temporal foundation for continuous queries over data streams. In *Proc. COMAD*, pages 72–80, January 6-8, 2005.

[9] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.

[10] C. Liebig, M. Cila, and A. Buchmann. Event composition in time-dependent distributed systems. In *Proc. Int. Conf. on Cooperative Information Systems (CoopIS)*, pages 70–78. IEEE Computer Society, 1999.

[11] I. Motakis and C. Zaniolo. Formal semantics for composite temporal events in active database rules. *Journal of Systems Integration*, 7(3-4):291–325, 1997.

[12] P. R. Pietzuch, B. Shand, and J. Bacon. A framework for event composition in distributed systems. In *Proc. ACM/IFIP/USENIX International Middleware Conference*, pages 62–82, 2003.

[13] J. van Benthem. *The Logic of Time*. Kluwer Academic Publishers, 2nd edition edition, 1991.