On Time

Uri Abraham [*]

April 3, 2014

Abstract

We discuss the question of time representation in models of concurrency.

Time is the subject of this lecture. It is an important notion in almost any scientific area as well as in philosophy, theology, art, and literature. Time has kept the fascination of all thinkers, and the study of time in physics for example is of prime importance. In Engineering too, time is extremely important. Perhaps Mathematics is an exception in that time is not intrinsic to the field, since this eternal edifice seems to have an existence out of Time. Of course, as a human endeavor, the development of Mathematics is immersed in Time, but as a subject it seems to be beyond its domain. For Computer Science, being both engineering and mathematics, both a theoretical and an applied science, Time is an issue of prime importance, and therefore the question of its representation is dealt in the following lectures.

1 Time in philosophy

The beginning of philosophy is, roughly speaking, attributed to the thinkers who lived some two thousands and five hundred years ago in or around Greece and are collectively known as the presocratic philosophers. Not much is known about them, and the little, indirect remaining pronouncements are

[*]Departments of Mathematics and Computer Science, Ben-Gurion University, Beer-Sheva, Israel. Email: abraham@cs.bgu.ac.il

often cryptic and obscure. Heraclitus is known as having said that "All things are flowing", and Plato wrote that "Heraclitus somewhere says that all things are in process and nothing stays still, and comparing things to the current of a river he says that you would not step twice into the same river". It is not completely clear what Heraclitus meant to say, but it seems that he was concerned with the general question of change: how can change be? If A changes to B then necessarily A is not B (or else there would not be a change), and thus A and B are different implying that there is no change since a change refers to a single object. As computer scientists we must solve similar problems, albeit in a different setting: what does it mean that a state changes? How can a variable change? Computer science is the engineering of change at its most abstract level, and so understanding change and stability must be a concern of us as well.

Another presocratic philosopher whose riddles and paradoxes are still intriguing is Zeno, and of his four paradoxes preserved by Aristotle the Achilles and Tortoise is the most famous. (The Early Greek Philosophy, by John Burnet.)

> Achilles will never overtake the tortoise. He must first reach the place from which the tortoise started. By that time the tortoise will have got some way ahead. Achilles must then make up that, and again the tortoise will be ahead. He is always coming nearer, but he never makes up to it.

"This argument", writes Burnet, "is intended to bring out the absurd conclusions which follow from the assumption that all quantity is discrete, and what Zeno has really done is to establish the conception of continuous quantity by a reductio ad absurdum of the other hypothesis." Zeno's paradox is relevant and interesting for us, and the question of whether Time should be represented by a discrete or a continuous line is a main theme of this lecture.

How should Time be represented? Guided by "modern" views of Physics, we all tend to think that Time is to be represented by the real numbers, or perhaps the positive real numbers if we believe that Time has to have a first moment. Yet this view would probably not be shared by all ancient thinkers. Pythagoras is said to have said that all things are numbers, and his followers the Pythagoreans believed that understanding the natural numbers is understanding the universe itself. They would probably be inclined to use natural numbers to represent Time, and as we shell shortly see we will also adopt this usage of the natural numbers in our lectures. Pythagoras is another presocratic philosopher, better known as the first mathematician who has a theorem attached to his name. It is ironic that this theorem (undoubtedly the most famous in the history of mathematics) is also the reason why we cannot accept the view that all that can be measured is to be measured with natural numbers (or rational numbers which are an immediate by-product). Yet the development of the real numbers in mathematics took many centuries, and only in the nineteenth century was there a sufficiently satisfying theory that explained the real numbers (Dedekind, Cantor).

In class we discuss here the definitions of the real numbers, rational number and natural numbers.

2 Time representations

Consider a certain execution of a computer system, comprising of several processes that operate concurrently, and suppose that each of these processes is serial. Namely that its operations are executed one after the other serially. An execution of an operation is an event stretched in time, and so any event, E, has a beginning, $Begin(E)$, and (if terminating) an ending $End(E)$. Here $Begin(E)$ and $End(E)$ are "moments of time" perhaps represented by real numbers so that $Begin(E) \leq End(E)$, and $(Begin(E), End(E))$ is the temporal interval during which E is executed. Two operation executions by different processes may either overlap (that is their temporal intervals intersect) or else one of the intervals precedes the other. We use here the interval notation (a, b) for the set of all points x such that $a < x < b$.

Suppose we videotape a dance and then show it at twice the original pace or perhaps at a slow motion. The dance might appear unusual and perhaps even distorted and unnatural if not at its original tempo. If we record some music or speech and play it back at a much higher (or slower) speed, the result will appear to be very distorted and perhaps completely incomprehensible. With executions of systems the situation is different: if we have two systems that execute the same protocols and generate the same set of events, but with two machines that differ only in that the clock of one is at twice the pace as the clock of the other machine, then for many purposes one might consider the two system executions as essentially the same. That is, many correctness conditions are such that what really counts is not the speed of the execution and the absolute length of the operation executions, but rather the temporal precedence relation between the events that comprise the execution. A typical example is the correctness condition known as “mutual-exclusion” (explained in details later on). The requirement is that any two critical section executions do not overlap, and this requirement is expressed only with the precedence relation. The actual length of each temporal interval that corresponds to an operation execution is irrelevant. This discussion leads to the conclusion that the temporal precedence relation between the events is of prime importance and this leads to the investigation of interval orderings.

There are certainly cases in which the length of the temporal intervals is important. For example when considering real-time issues, one may want to specify that a certain response is never delayed by more than say 3 seconds. In such cases, the interval ordering approach is insufficient, and one may want to have real numbers to represent moments in some scale (for example the time in seconds). For example, any real number $t \geq 0$ can represent the number of seconds (with any degree of preciseness) that elapsed since the system was switched on. So for each event e we denote with $\mathit{Begin}(e), \mathit{End}(e)$ those real numbers with $\mathit{Begin}(e) \leq \mathit{End}(e)$ that mark the beginning and ending of event e. For example, $\mathit{Begin}(e) = 21.114$ says that event e began 21 seconds and 114 milliseconds after the system began its execution. An alternative is to take the temporal domain to be the set of natural numbers, counting the number of nanoseconds since the system was switched on.

A third possibility, one which is in between the “real time representation” and the abstract “interval ordering representation” is obtained when the end-points of the temporal intervals are abstractly represented as points on some linearly ordered set (of unspecified nature, not related to any scale). We will concentrate on the abstract interval ordering representation and the end-point representation without any scale attached to it.

3 Interval ordering

We begin with a short recall of some basic definitions: partial ordering, linear ordering, isomorphism of partial ordering etc. Recall that a partially ordered set (poset) is a pair $\mathcal{A} = (A, <_A)$ so that $<_A$ is an irreflexive and transitive relation on A. It is a linear ordering iff for every $a, b \in A$ we have $a = b$ or $a <_A b$ or $b <_A a$. When $(A, <_A)$ is a partial ordered set, we write $a \leq_A b$) for "$a <_A b \vee a = b$". Clearly $\leq_A$ is reflexive, transitive and antisymmetric. Sometimes a partial ordering is defined as a relation that satisfies these three properties.

We were rather pedantic in writing $\mathcal{A} = (A, <_A)$, thereby making a distinction between the universe A of the poset and the poset itself which consists of the universe and the ordering relation $<_A$ (which is, formally speaking, a set of ordered pairs, namely a subset of $A \times A$). In many places, and certainly in lectures and informal discussions, one may write $A = (A, <_A)$ and refer to A as both the universe (a set) and its poset. This "overloading" of meaning on the symbol A should not be problematic for the attentive reader.

A *convex set* of a poset $(A, <_A)$ is a subset $X \subseteq A$ such that for every $a, b \in X$ for every $c \in A$ if $a <_A c <_A b$ then $c \in X$ as well. When $(A, <_A)$ is a linear ordering an interval of A is a non-empty convex set. The collection of all non-empty intervals of a linear ordering $(A, <_A)$ is denoted $Interval(A, <_A)$ (or $Interval(A)$ for short). Most often intervals are defined by their end-points. For $a, b \in A$ with $a \leq_A b$ the closed interval $[a, b] = \{x \in A \mid a \leq_A x \leq_A b\}$ and the open intervals $(a, b) = \{x \in A \mid a <_A x <_A b\}$ are defined. Unbounded intervals $(a, \infty) = \{x \in A \mid a <_A x\}$ and $(-\infty, a)$ are also defined.

For completeness of our exposition we define the notions of isomorphism and embedding of partial orderings. Let $\mathcal{A}_1 = (A_1, <_1)$ and $\mathcal{A}_2 = (A_2, <_2)$ be partial ordered sets. A function $f : A_1 \to A_2$ is an isomorphism of $\mathcal{A}_1$ on $\mathcal{A}_2$ if it is a bijection (one-to-one and onto A_2) and

$$\text{for every } a, b \in A_1,\ a <_1 b \text{ if and only if } f(a) <_2 f(b). \tag{1}$$

If $f : A_1 \to A_2$ is an injection (one-to-one but not necessarily onto A_2) and it satisfies formula (1) then we say that f is an *embedding* of $\mathcal{A}_1$ into $\mathcal{A}_2$. In this case, f is clearly an isomorphism between $\mathcal{A}_1$ and its image $f[A_1]$ (viewed as a subposet of $\mathcal{A}_2$).

Exercise 3.1 *Prove that if f is an isomorphism as above then $f^{-1} : A_2 \to A_1$ is an isomorphism of $\mathcal{A}_2$ on $\mathcal{A}_1$. Prove also that the relation "$\mathcal{A}_1$ is isomorphic to $\mathcal{A}_2$" is an equivalence relation on the class of posets.*

Complete the proof within the brackets so that it appears in blue.

We continue our lecture with some observations that date back to Bertrand Russell and Norbert Wiener concerning "interval orderings" and time. Then we investigate finiteness conditions (introduced in this context by Lamport).

Modeling is not only an abstract description of reality, but also a simplification which disregards some non-essential features. For many applications in computer science the precise time of each event occurrence is not needed, and what really counts is the *precedence* relation. This relation is a partial ordering which holds for events X and Y if X ends before Y begins. So two events are incomparable if their temporal extensions overlap. Two simple and useful properties of this precedence relation reflect its nature: The Russell–Wiener property, and the finiteness property of Lamport.

Recall that for a linearly ordered set $(A, <_A)$ the collection of all non-empty intervals is denoted $Interval(A, <_A)$ (or $Interval(A)$ for short).

A main definition is that of interval orderings. (Fishburn [6] is devoted to this subject.)

Definition 3.2 *(Interval orderings) Let $(L, <_L)$ be a linear ordering and $S \subseteq Interval(L)$ be a set of intervals of L. Then define for $I, J \in S$, $I \prec_L J$ iff $\forall x \in I \, \forall y \in J (x <_L y)$. We say that I precedes J when $I \prec_L J$. It is easy to prove that $\prec_L$ is a partial ordering of S. We say that the pair $(S, \prec_L)$ is an interval ordering.*

More formally, an interval ordering is any partial ordering that is isomorphic *to $(S, \prec_L)$ where S is a set of intervals of some linear ordering $(L, <_L)$.*

This leads to the notion of representation.

Definition 3.3 *A* representation *of a partial ordering $(E, <_R)$ is a function μ so that, for some linear ordering $(L, <_L)$, $\mu : E \to Interval(L)$ maps E into the poset of intervals of L under the $\prec_L$ relation. That is, for all $a, b \in E$,*

$$a <_R b \text{ iff } \mu(a) \prec_L \mu(b). \quad (2)$$

An interval ordering is thus an ordering that has a representation.

We did not require that μ is one-to-one, but it is not hard to see that if μ is a representation, then there exists a representation that is one-to-one (possibly into intervals of a richer linear ordering).

For concreteness, the intervals in a representation may be assumed to be closed. That is, of the form $I = [a, b]$ where a is denoted $Begin(I)$, and b is denoted $End(I)$. Then we have $I \prec_L J$ iff $End(I) < Begin(J)$. Also, if $I \not\prec_L J$ denotes the negation of $I \prec_L J$, then we have that $I \not\prec_L J$ iff $Begin(J) \leq End(I)$.

This convention of closed intervals is not essential, and we will have occasions where half-closed intervals of the form $[a, b)$ will be used.

Two intervals I and J are said to overlap iff $I \cap J \neq \emptyset$. Since L is a linear ordering, for any intervals I and J either I and J are comparable ($I \prec_L J$ or $J \prec_L I$) or else I and J overlap. Thus overlapping coincides with incomparability.

Exercise 3.4 *Let $(L, <_L)$ be a linear ordering, and $I, J \subset L$ be two non-empty intervals (that is convex sets). Prove that either $I \prec_L J$ or $J \prec_L I$ or $I \cap J \neq \emptyset$.*

The following property of $\prec_L$ is of prime importance. Suppose that $A \prec_L B$ and $C \prec_L D$ are intervals such that B and C overlap; then $A \prec_L D$. To prove this pick any point $p \in B \cap C$. Then $\forall a \in A\ (a <_L p)$ follows from $A \prec_L B$ since $p \in B$. Similarly $\forall d \in D\ (p <_L d)$ follows. Hence, for any $a \in A$ and $d \in D$, $a <_L p <_L d$, and thus $a <_L d$. This means that $A \prec_L D$. Bertrand Russell and Norbert Wiener [10] proved that this property characterizes interval orderings, as we are going to show next (see also Fishburn [6] and Lamport [7].) In fact, instead of requiring that B and C overlap, we may simply ask that $\neg(C \prec_L B)$. It turns out that the relation $\neg(C \prec_L B)$ is so significant that it deserves a special symbol. So we define

$$x \,\not\succ_L y \text{ iff } \neg(y \prec_L x). \tag{3}$$

With this symbol, the property of $\prec_L$ discussed above is:

$$A \prec_L B \,\not\succ_L C \prec_L D \longrightarrow A \prec_L D.$$

This is easier to remember because of its resemblance with transitivity. We can read $X \not\succ Y$ as "X weakly precedes Y", or as "X can have an impact on Y".

Definition 3.5 (The Russell–Wiener property) *Let $(E, <_R)$ be a partial-ordering. The Russell–Wiener property is the following:*

$$\forall a, b, c, d \in E(a <_R b \wedge \neg(c <_R b) \wedge c <_R d \longrightarrow a <_R d) \tag{4}$$

We have seen, for any linear ordering $(L, <)$ and a set of intervals of L, that the Russell–Wiener property holds for the relation $\prec_L$.

Theorem 3.6 (Russell and Wiener) *A partial ordering satisfies the Russell–Wiener property (4) iff it has a representation.*

Proof. [Wiener [10]] We have already argued that an interval ordering $\prec_L$ satisfies the Russell–Wiener property.

We define the notion of "antichain" in a partial ordering. Let $(E, <_E)$ be any partial ordering. We say that $X \subseteq E$ is an *antichain* iff it is a non-empty set of pairwise incomparable elements:

$$\forall x \neq y \in X(\neg\, x <_R y \text{ and } \neg\, y <_R x)).$$

An antichain X is a maximal antichain if there is no antichain Y such that $X \subset Y$ ($X \subset Y$ means $X \subseteq Y$ and $X \neq Y$). Equivalently, an antichain X is maximal if for every $e \in E$ we have that $e \in X$ or $x <_E e$ for some $x \in X$ or $e <_E x$ for some $x \in X$.

Assume now an ordering $(E, <_R)$ that satisfies the Russell–Wiener property, and we shall find a representation.

A maximal (under inclusion) antichain is called a *moment* (it may be finite or infinite). An antichain X is thus a moment iff any y not in X is comparable with some $x \in X$.

Let L be the collection of all moments, and define for $X, Y \in L$,

$$X < Y \text{ iff } \exists x \in X, \exists y \in Y\ (x <_R y).$$

Exercise 3.7 *Prove that $<$ is a linear ordering of the set of moments.*

Now we define the representation map μ from E into intervals of L. Let $\mu(e)$ for $e \in E$ be the collection of all moments X such that $e \in X$.

$$\mu(e) = \{X \in \mathit{moments} \mid e \in X\}.$$

Since any $e \in E$ is contained in at least one maximal antichain, the representing set $\mu(e)$ is nonempty. [1]

We claim that

[1] It seems that we need here the axiom of choice in order to conclude that any e is contained in some antichain, but in applications finite antichains suffice because the finiteness condition implies that all antichains are finite.

1. $\mu(e)$ is convex (an interval in L), and
2. for all $e, f \in E$, $e <_R f$ iff $\mu(e) \prec_L \mu(f)$.

The first item is in the following exercise.

Exercise 3.8 *Prove the convexity of $\mu(e)$.*

Hint: Assume $X < Y < Z$, and $X, Z \in \mu(e)$. Then $e \in X$ and $e \in Z$. To prove that $Y \in \mu(e)$ we need to show that $e \in Y$. Assume for the sake of a contradiction that $e \notin Y$. Use the fact that Y is a maximal antichain and obtain a contradiction.

In order to prove the second part of the claim, assume first that $e <_R f$. To prove that $\mu(e) \prec_L \mu(f)$, we must prove for every $X \in \mu(e)$ and $Y \in \mu(f)$ that $X < Y$. Since $X \in \mu(e)$ and $Y \in \mu(f)$, $e \in X$ and $f \in Y$ by definition of μ, and so $X < Y$ by definition of $<$. Thus $\mu(e) \prec_L \mu(f)$.

Next we check that if $\neg\, e <_R f$ then $\neg(\mu(e) \prec_L \mu(f))$. Assume $\neg\, e <_R f$. If $f <_R e$ then $\mu(f) \prec_L \mu(e)$, as proved above, and hence $\neg(\mu(e) \prec_L \mu(f))$ by the transitivity and irreflexivity of $\prec_L$. If e and f are incomparable under $<_R$ then a maximal antichain X exists which includes both e and f. So $X \in \mu(e) \cap \mu(f)$. And again this shows $\neg(\mu(e) \prec_L \mu(f))$. □

This proof is the one given by Wiener [10]; it follows Russell's intuition which wants to base the concept of a temporal point (moment) on the more elementary notions of event and precedence. The proof in Fishburn [6] is more constructive.

Frank D. Anger [4] Studied the following property of a poset $(A, <)$ which we call the Anger property:

$$\text{For all } a, b, c, d \in A, \text{ if } (b \not< a,\ b < c,\ \text{and } d \not< c) \text{ then } d \not< a. \qquad (5)$$

Exercise 3.9 *Prove that Anger's property is equivalent to the Russell–Wiener property.*

Again, as for the precedence relation on the intervals, relation $b \not< a$ turns out to have an independent importance in the context that interest us, and hence we introduce a special symbol: whenever $(A, <)$ is a poset, define $a \not\succ b$ iff $b \not< a$. Intuitively, we can interpret $a \not\succ b$ as "a can have an impact on b". Under this interpretation the property of Anger expressed as $a \not\succ b < c \not\succ d$

$\Rightarrow a \not\succ d$ is easier to remember since it has the appearance of a certain type of transitivity.

Lamport has used the notation $a \longrightarrow b$ and $a \dashrightarrow b$ in a way that is similar to our $a < b$ and $a \not\succ b$. In this notation, a poset $(A, \longrightarrow b)$ satisfies the Russell–Wiener property iff $a \longrightarrow b \dashrightarrow c \longrightarrow d$ implies $a \longrightarrow d$. I sometimes find that these two arrow relations are more suggestive and hence a better choice than the $<$ and $\not\succ$ notations, but in many cases it seems to me that $<$ is the natural choice for the precedence relation and hence I've finally opted for the latter notation.

3.1 Another proof

We give here another proof of Theorem 3.6 which has the advantage of yielding a useful algorithm. Given a partially ordered set $R = (E, <_R)$ define for every $e \in E$, $(< e) = \{x \in E \mid x <_R e\}$. We say that $(< e)$ is the initial segment determined by e. Similarly, $(> e) = \{x \in E \mid e <_R x\}$ is the final segment determined by e. We shall also use notations such as $(\leq e) = (< e) \cup \{e\}$.

Lemma 3.10 *A partial ordering $R = (E, <_R)$ satisfies the Russell–Wiener property (4) iff for every $a, b \in E$ we have $(< a) \subseteq (< b)$ or $(< b) \subseteq (< a)$. (This property will be called the initial segments comparability property.)*

Exercise 3.11 *Prove the lemma.*

Exercise 3.12 *Prove that $R = (E, <_R)$ satisfies the Russell-Wiener property if and only if R^* satisfies that property as well (where $R^* = (E, <^*)$ is the inverse relation: $a <^* b$ iff $b < a$).*

It follows that a partial ordering $R = (E, <_R)$ satisfies the Russell-Wiener property if and only if for every $a, b \in E$ we have $(> a) \subseteq (> b)$ or $(> b) \subseteq (> a)$.

We reprove now theorem 3.6. Given $R = (E, <_R)$ that satisfies the Russell–Wiener property let $L = \{(< a) \mid a \in E\} \cup \{E\}$ be the set of all initial segments determined by members of E (and E itself) ordered by the proper inclusion relation $\subset$. Lemma 3.10 says that this is a linear ordering of L. For every $e \in E$ define $\mu(e)$ an interval of $(L, \subseteq)$ by

$$\mu(e) = \{(< a) \mid a \in E,\ (< e) \subseteq (< a) \text{ and } e \notin (< a)\}.$$

That is, $\mu(e)$ is the set of all initial segments $(< a)$ that contain all of $(< e)$ but not e itself. Since $(< e) \in \mu(e)$, $\mu(e)$ is nonempty. Note that if $(< a) \in \mu(e)$ then a and e are incomparable in $<_R$.

Exercise 3.13 *Prove that μ is a representation of $(E, <_R)$ into intervals of L (ordered by $\prec_L$).*

Hints. (1) Prove first that $\mu(e)$ is a convex set. That is, take any $a, b, c \in E$ such that $(< a) \subseteq (< b) \subseteq (< c)$ and $a, c \in \mu(e)$, and then deduce that $b \in \mu(e)$. (2) Then prove that μ is order preserving: if $e_1 <_R e_2$ then $\mu(e_1) \prec_L \mu(e_2)$. (3) And finally prove that if $\neg(e_1 <_R e_2)$ then $\neg(\mu(e_1) \prec_L \mu(e_2))$. (For this, we have to find $x \in \mu(e_1)$ and $y \in \mu(e_2)$ such that it is not the case that $x \subset y$. If $(< e_1) \not\subseteq (< e_2)$, then $x = (< e_1)$ and $y = (< e_2)$ work. If $(< e_1) \subseteq (< e_2)$, then $x = y = (< e_2)$ works.

In figure 3.1 at the end of this article there is an example. it shows a set of intervals on a timeline axis $A, \ldots, I$. Then a diagram shows the abstract precedence relation on these objects. Then the process of representing this precedence relation is described. For each of $a \in \{A, \ldots, I\}$ the initial segment $(< a)$ is described. The set of points $L = \{(< A), \ldots, (< I)\}$ is linearly ordered, and the representation μ is obtained.

We sum-up the theorems that were obtained so far in the following.

Theorem 3.14 *Let $R = (E, <)$ be some partial ordering. Then the following are all equivalent.*

1. *R is an interval ordering. That is, it has a representation into a set of intervals of some linear ordering.*
2. *R satisfies the Russell–Wiener property.*
3. *R satisfies the Agner property.*
4. *R satisfies the initial segment comparability property.*

3.2 Finiteness conditions and global-time structures

I think no one will contest the assumption that any system can only contain finitely many event occurrences that precede a given instant. In fact, since the number of concurrently operating agents is finite, only finitely many events can even begin before any given instant t, and hence all others begin after t.[2]

Leslie Lamport [8] proposed and showed the applicability of the following property.

Definition 3.15 (Finiteness Property) *A partial ordering* $(E, <)$ *is said to satisfy the finiteness condition (of Lamport) iff for every* $e \in E$ *the set* $\{x \in E \mid \neg(e < x)\}$ *is finite.*

A weaker property is the "finite predecessors property". A partial ordering $(E, <)$ *satisfies the finite predecessors property iff for every* $e \in E$ *the set* $\{x \in E \mid x < e\}$ *is finite.*

Observe that Lamport's finiteness condition implies the finite predecessors property and it also implies that every antichain is finite. (The finiteness condition is strictly stronger than the conjunction of these two properties: take for example the union of two infinite chains of order type of the natural numbers and with no relation between members of different chains. Then every antichain contains at most two elements and the finite predecessors property holds, and yet Lamport's finiteness property does not hold.)

The finite predecessors property enables induction. To prove that $\forall x \phi(x)$ holds in $(E, <)$ we assume towards a contradiction that for some $e \in E$ $\phi[^x_e]$ does not hold, and take e to be minimal with this property. That is, for any $e' < e$, $\phi[^{x'}_e]$ does hold. If we can derive a contradiction from this situation then $\forall x \phi$ is proved. (The notation $\phi[^x_e]$ is for the formula obtained when e is substituted for x in the formula ϕ.)

Lamport distinguishes between terminating and nonterminating events. Intuitively, terminating events have a bounded (finite) duration, while nonterminating events last forever. A nonterminating event can never be followed by another event. Formally we have the following.

---

[2]An amusing argument is known as Thompson's lamp. Nowadays these lamps are not so fashionable but we used to have table lamps with a button you click once to put it on and click again to shut it off. Now suppose someone plays with clicking this button in accelerating speed so that the interval between any two clicks is half the interval between the preceding clicks. So if the lamp is on for the first second it is off for the next half second and on again for the next quarter of a second and so on. What will be the state of the lamp Thompson asks after two seconds? Is this question an argument against such super-tasks?

Definition 3.16 (Lamport) Global-time structures: *Let* $\mathcal{S} = (E, <, T)$ *be a structure where* E *is a non-empty set (the universe),* $<$ *is a partial ordering of* E*, and* T *a unary predicate on* E *represented as a subset. Members of* E *are called events,* T *is a subset of* E *called the set of terminating events, and* $<$ *is called the precedence relation of* $\mathcal{S}$*. (We shall often write* x *is terminating instead of* $T(x)$*.) Suppose that relation* $\not\succ$ *is defined by* $a \not\succ b$ *iff* $\neg(b < a)$*. Then* $\mathcal{S}$ *is called a* global-time structure *iff:*

1. $\forall a, b, c, d\ (a < b \not\succ c < d \longrightarrow a < d)$*. (This is the Russell–Wiener property. So* $<$ *is an interval ordering.)*
2. *Every antichain is finite, and the finite predecessors property holds: that is for any* $a \in E$ *the set* $\{x \in E \mid x < a\}$ *is finite.*
3. *If* $\neg T(e)$ *then there is no* $x \in E$ *such that* $e < x$*. In other words, if an event* e *is followed by some event* x*, then* e *must be terminating.*
4. *For any terminating* $a \in E$*,* $\{x \in E \mid x \not\succ a\}$ *is finite. (The finiteness property of Lamport, for terminating events only: the set of events that follow a terminating event is co-finite (a subset is co-finite if its complement is finite).)*

Intuitively, the finiteness property of Lamport says that if a is terminating, then most of the events follow a, where "most" means that there is only a finite number of exceptions. So this property is meaningful (non trivial) only if the set of events is infinite. To the reader who is familiar with the notations of Lamport in [8] I should say that $\prec$ here corresponds to $\longrightarrow$ and $\not\succ$ corresponds to $\dashrightarrow$.

Claim: It follows from (3) that the nonterminating events form an antichain, and by (2) it is a finite antichain. Item (4) implies that there can be only countably many terminating events, and so it follows that a global-time structure is countable.

Proof of the claim that the number of events in a global-time structure is countable. For any terminating event e in E let $p(e)$ be the number of events x with $x \prec e$. The finite predecessor property implies that $p(e)$ is finite, and we define for every $n \in \mathbb{N}$ $E_n = \{e \in E \mid T(e) \wedge p(e) = n\}$. Then $\bigcup_{n \in \mathbb{N}} E_n$ is the set of all terminating events and it suffices to prove that every E_n is finite to conclude that the number of all terminating events is countable, and hence that the number of all events is countable (as the nonterminating events form an antichain which is finite). If some E_n were infinite, pick $a \in E_n$ and apply the finiteness condition (4) to obtain that $X = \{x \in E \mid x \not> a\}$ is finite. So there is some $e \in E_n \setminus X$ since E_n is infinite. But then $e > a$ and hence all predecessors of a—and a itself—are predecessors of e, which shows that $p(e) \geq n+1$, in contradiction to $e \in E_n$.

Global-time structures serve to model the following situations. There are finitely many processes $P_1, \ldots, P_n$ where every process is serial (which means that the events generated by P_i are linearly ordered). The set of events is the union $E = \bigcup_{1\leq i\leq n} P_i$ (where P_i also denotes the set of events that process P_i generates). Some of the events are terminating, and $T \subseteq E$ denotes the terminating events. The precedence relation $<$ on the events is $<$. In such global-time structures $(E, <, T)$ every antichain has size $\leq n$ because every P_i is serial.

Definition 3.17 *[Representations] Let $(L, <_L)$ be a linear ordering and $\mathcal{S} = (E, <_E, T)$ be a global-time structure. We say that μ is a* representation *of $\mathcal{S}$ in L iff μ is a representation as in Definition 3.3 in which terminating events are represented by bounded intervals, and nonterminating events are represented by right unbounded intervals. That is:*

1. *For each $e \in E$, $\mu(e)$ is an interval in L, and*

$$e_1 <_E e_2 \text{ iff } \mu(e_1) \prec_L \mu(e_2).$$

2. *If $T(e)$ holds, then $\mu(e)$ is of the form $[a, b]$, but if $\neg T(e)$ holds (that is e is nonterminating) then $\mu(e)$ has the form $[a, \infty) = \{x \in L \mid a \leq_L x\}$.*

Lemma 3.18 *Let $\mathcal{S} = (E, <_E, T)$ be a global-time structure and μ a representation of $\mathcal{S}$ into intervals of a linear ordering $(L, <_L)$. If $e \in E$ is a nonterminating event then, except for a finite number of events, $\mu(x) \subseteq \mu(e)$ holds for all (other) events.*

Proof. The lemma is obvious when E is finite, since in this case we may take all of E as the exceptional set. So we assume that E is infinite. The set of events x such that $x < e$ is finite (by the finite precedence property) and hence there is a terminating event a such that it is not the case that $a < e$. That is $Begin(e) \leq End(a)$. There is a finite set of events X such that if x is any event not in X then $a <_E x$ holds, and hence $Begin(e) < Begin(x)$.

It must be that $\mu(x) \subseteq \mu(e)$ for such terminating x's, or else there is some $\ell \in \mu(x) \backslash \mu(e)$. But then a simple argument shows that ℓ is above all members of $\mu(e)$ (namely $m <_L \ell$ for all $m \in \mu(e)$). To see that this is impossible take any x' such that $x < x'$ and conclude that $\mu(e) \prec \mu(x')$ and hence that $e < x'$ in contradiction to the assumption that e is terminating. □

Theorem 3.19 *Any global-time model has a representation as intervals of natural numbers. That is,* $(L, <_L)$ *can be chosen as* $(\mathbb{N}, <)$.

Proof. Consider a representation μ of $(E, <_E)$ into any linear ordering $(L, <_L)$ as obtained by Theorem 3.6. The theorem is obvious is L or E are finite. So we may assume that E and L are infinite. Make a list $\langle e_n \mid n \in \mathbb{N} \rangle$ of all terminating events in E (assuming E is infinite, for otherwise the claim is obvious). Define by induction on n a finite set P_n of points in L so that $P_n \subseteq P_{n+1}$. $P_0 = \{a_0\}$ where a_0 is any point in $\mu(e_0)$. At the nth stage of the construction consider e_n and the finite set of points P_n constructed so far. Define $P_{n+1} \supseteq P_n$ by adding points in the interval $\mu(e_n)$ in such a way that $P_{n+1} \cap \mu(e_n) \neq \emptyset$, and for every $m < n$ if $\mu(e_n) \cap \mu(e_m) \neq \emptyset$, then $P_{n+1} \cap \mu(e_n) \cap \mu(e_m) \neq \emptyset$. That is P_{n+1} adds to P_n points as necessary for this evidence. Now the points in $P = \bigcup_{i \in \mathbb{N}} P_n$ have order-type ω (that of $\mathbb{N}$), because P satisfies the finiteness property in L. Indeed, any $a \in P$ is in some $\mu(e_n)$, and there is some N such that for every $p \geq N$, $e_n <_E e_p$, and hence $\mu(e_n) \prec_L \mu(e_p)$. Thus the points x added at the pth stages of the construction (for $p \geq N$) are greater than a (x being in $\mu(e_p)$).

Consider the function $\mu'(e) = \mu(e) \cap P$ for $e \in E$. We claim that it is a representation such that $\mu'(e)$ is finite iff e is terminating. Clearly every $\mu'(e)$ is nonempty. The finiteness of $\mu'(e)$ for terminating events was argued already when we proved that for every terminating e there is N so that for all $p \geq N$ points at stage p are added strictly above $\mu(e)$. If e is nonterminating, then the previous lemma implies that for almost all m's $\mu(e_m) \subseteq \mu(e)$. As a new point is added to $\mu(e_m)$, this shows that if e is nonterminating then $\mu'(e)$ is infinite.

Clearly for any events x_1 and x_2 $x_1 < x_2$ implies $\mu'(x_1) \prec_P \mu'(x_1)$, but the other direction is also true, by construction. That is, if x_1 and x_2 are incomparable, then $\mu(x_1)$ and $\mu(x_2)$ intersect, and hence $\mu'(x_1)$ and $\mu'(x_2)$ intersect. This proves the claim that in choosing the linear temporal domain for representations there is no loss of generality by picking the set ω of natural numbers as our choice. □

The finiteness properties of Lamport enable induction "on the right-end points" for the terminating events as explained here. Let $\mathcal{S} = (E, <_E, T)$ be a global-time model. We have seen that there is a representation of $\mathcal{S}$ as intervals in $\mathbb{N}$, and fixing such a representation μ it is possible to prove by induction on n statements of the form: For every interval I in the range of μ, if n is the right-end point of the interval then $\phi(I)$ holds. It is possible and perhaps more elegant to make such an inductive proof without referring to any representation.

Define for every terminating event e a set $P(e) = \{x \in E \mid \neg(e <_E x)\}$. The finiteness property of Lamport says that $P(e)$ is finite, and hence proofs by induction on $|P(e)|$ can be carried out. That is, to prove a statement of the form: "For all terminating events e $\psi(e)$ holds" argue as follows. Assume that event e is a counterexample with least value of $|P(e)|$, and derive a contradiction by finding another counterexample e_0 such that for some event u

$$e_0 <_E u, \text{ and } \neg(e <_E u).$$

We prove that $P(e_0) \subseteq P(e)$. Suppose $x \in P(e_0)$ but $x \notin P(e)$. Then $e <_E x$, and the Russell–Wiener property applies to

$$e_0 < u, \ \neg(e < u), \ e < x$$

and gives $e_0 <_E x$, in contradiction to $x \in P(e_0)$.

3.3 Augmentations of posets

The result described in this section (due to Lamport but its proof, I think, was never published) is very important (especially for the discussion of linearizability) and the effort invested in its understanding is worthwhile.

Definition 3.20 *1. Given a partial ordering $\mathcal{S} = (E, <)$, a partial ordering $\mathcal{S}^* = (E, <^*)$ with the same universe E is said to* augment *$\mathcal{S}$ iff for all $a, b \in E$ if $a < b$ then $a <^* b$.*

2. A linearization *of a partially ordered set is an augmentation that is a linear ordering.*

Note that we require that a poset and its augmentations have the same universe, and it is only the ordering relation that is augmented. Some writers use the term "extension" for what we call here (following Pratt [9]) augmentation. We prefer to reserve the term "extension" to those cases when E is allowed to increase with new points.

A classical theorem (Szpilrajn's theorem) states that any poset has a linearly ordered augmentation. The main step in the proof of that theorem is in the following.

Lemma 3.21 *Let $S = (E, <)$ be a partial ordering and suppose that $a, b \in E$ are incomparable (i.e. $a \neq b$ and neither $a < b$ nor $b < a$). Then there is an augmentation $(E, <^*)$ such that $a <^* b$.*

Proof. We write $x \leq y$ for $x < y \vee x = y$. For every distinct $x, y \in E$ define $x <^* y$ iff $x < y$ or else $x \leq a$ and $b \leq y$. It is easy to check that $<^*$ is a partial ordering (that is an irreflexive and transitive relation over E) that extends $<$ and is such that $a <^* b$). Note that $<^*$ is the minimal augmentation of $<$ for which $a <^* b$. □

An augmentation of an interval ordering need not be an interval ordering again. Take the following simple example: $E = \{I, J, K, L\}$ is partially ordered so that $L < K$ but there are no more relations (this is clearly an interval ordering). Now, if we add as a new relation $I < J$, then the resulting poset with the relations $I < J$ and $L < K$ is not an interval ordering (because if J and L are incomparable then we should have I before J).

An augmentation that is a linear ordering is, of course, an interval ordering since any linear ordering is an interval ordering.

Let E be any set and μ a function defined on E which takes values that are intervals in some linear ordering (which, for concreteness we assume to be the set of rational numbers with their ordering). A function μ' defined on E and such that $\mu'(e)$ is a subinterval of $\mu(e)$ for every $e \in E$ is said to be a *refinement* of μ. If $\mathcal{S} = (E, <_S)$ is an interval ordering and μ a representation into intervals of the rational numbers, then any refinement μ' of μ naturally defines an augmentation of $\mathcal{S}$. That is, if $\mu'(x)$ is a subinterval of $\mu(x)$ for every event x, then the ordering $<^*$ defined by $x <^* x'$ iff $\mu'(x) \prec \mu'(x')$ is an augmentation of $<_S$. The following theorem of Lamport (stated in [7] without proof) shows that *any* augmentation of $\mathcal{S}$ can be obtained in this way (if it is an interval ordering). The theorem will be used when the notion of linearizability is discussed.

Theorem 3.22 *Let $\mathcal{S} = (E, <_S)$ and $\mathcal{T} = (E, <_T)$ be interval orderings where $\mathcal{S}$ satisfies Lamport's finiteness property. Assume that $\mathcal{T}$ augments $\mathcal{S}$ (hence the finiteness property holds for $\mathcal{T}$ as well). Suppose that μ_S is a representation of $\mathcal{S}$ that maps E into left-closed right-open intervals of rational numbers (or any dense linear ordering). Then there is a representation μ_T of $\mathcal{T}$ which is a refinement of μ_S (i.e. for any $e \in E$, $\mu_T(e)$ is a subinterval of $\mu_S(e)$).*

The requirement that the intervals are of the form $[a, b)$ (where a, b are rational numbers with $a < b$) is made to avoid trivial counterexamples. (For instance, if $\mathcal{S}$ contains two incomparable events a and b, and if $\mu_S(a) = \mu_S(b)$ are points, then the theorem does not hold.) Note that the intersection of any two left-closed right-open intervals is again of this type (when non-empty). Of course, our decision to take left-closed rather than right-closed intervals is arbitrary.

We use the following notation. If $I = [a, b)$ is a rational interval (where $a < b$) then we write $a = \min I$ and $b = \sup I$.

A simpler version of this theorem is when the augmentation is a linear ordering. In most applications it is this simpler version that is used, and hence we bring it first.

Theorem 3.23 *Let $\mathcal{S} = (E, <_S)$ and $\mathcal{T} = (E, <_T)$ be interval orderings where $\mathcal{S}$ satisfies Lamport's finiteness property, and $\mathcal{T}$ is actually a linear ordering that augments $\mathcal{S}$. Suppose that μ_S is a representation of $\mathcal{S}$ that maps E into left-closed right-open intervals of rational numbers. Then there is a representation μ_T of $\mathcal{T}$ which is a refinement of μ_S.*

Proof. It doesn't make much difference in the proof if we assume that E is an infinite set. So $\mathcal{T}$ is an infinite ordering with the property that the number of predecessors of each of its points is finite. The order-type of $\mathcal{T}$ is hence that of the natural numbers, and we can enumerate E in increasing $<_T$ ordering: $e_0 <_T e_1 <_T \cdots e_i <_T e_{i+1} \cdots$. By induction on $i \geq 0$ we shall define μ_i a refinement of μ such that the following holds.

1. $\mu_0 = \mu$ is the given representation of $\mathcal{S}$.
2. For $i > 0$ μ_i is a refinement of μ_{i-1}. For every $k \neq i$ $\mu_{i-1}(k) = \mu_i(k)$. That is, μ_i can only reduce the interval of e_{i-1} and it must not change any other interval. (It follows that for every $k \geq i$, $\mu_i(e_k) = \mu(e_k)$.)

3. $\mu_i(e_0) \prec \mu_i(e_1) \prec \cdots \prec \mu_i(e_{i-1})$.

4. For every $k \geq i$, $\sup \mu_i(e_k) > \sup \mu(e_k)$.

$\mu_0 = \mu$ is given. To define μ_1 we have to define $\mu_1(e_0)$ as a subinterval of $\mu(e_0)$, and then define $\mu_1(e_k) = \mu(e_k)$ for $k > 0$. The sole requirement is that for every $k \geq 1$ $\sup(I_0) < \sup(\mu(e_k))$.

1. First, we observe that there is no $e \in E$ with $\sup \mu(e) \leq \min \mu(e_0)$. For if there was such an event, then $e <_S e_0$ would follow from the fact that μ is a representation of $\mathcal{S}$, and then $e <_T e_0$ could be deduced from the fact that $\mathcal{T}$ augments $\mathcal{S}$, and a contradiction to the fact that e_0 is the first event in E in the $<_T$ ordering.

2. Next we observe that the set $U_0 = \{e \in E \mid \sup(e) \leq \sup(e_0)\}$ is finite and nonempty. It is finite since by the Lamport finiteness property of $\mathcal{S}$ there is a finite set of events $F \subseteq E$ such that if $e \in E \setminus F$ then $e_0 < e$ (and thence $\sup \mu(e_0) \leq \min \mu(e)$). But $U_0 \subseteq F$ and so U_0 is finite. As $e_0 \in U_0$, U_0 is nonempty.

3. It follows from item 1 that if $e \in U_0$ then $\min \mu(e_0) < \sup \mu(e) \leq \sup(e_0)$.

4. Since any finite set of numbers has a minimum, we can define $m = \min U_0$ and get that $\min \mu(e_0) < m \leq \sup(e_0)$.

5. Finally, we take any rational number q such that $\min \mu(e_0) < q < m$, and define $\mu_1(e_0) = [\min \mu(e_0), q)$.

6. By the minimality of m we get that for every $k \geq 1$ $q < m \leq \sup \mu(e_k)$, and hence $\mu_1(e_0)$ is as required.

We now assume that μ_{i-1} is defined where $i - 1 \geq 1$, and we describe the definition of μ_i. By the inductive assumption $\mu_{i-1}(e_0) \prec \cdots \prec \mu_{i-1}(e_{i-2})$, and, for $k \geq i-1$, $\mu_{i-1}(e_k) = \mu(e_k)$ and $\sup \mu_{i-1}(e_{i-2}) < \sup \mu(e_k)$.

Let $p = \max\{\sup \mu_{i-1}(e_{i-2}), \min \mu(e_{i-1})\}$. Then $p < \sup \mu(e_{i-1})$. It can be shown that for every $k \geq i-1$, $p < \sup \mu(e_k)$. We shall find a suitable $q \in \mathbb{Q}$ such that $p < q < \sup \mu(e_{i-1})$ and define $\mu_i(e_{i-1}) = [p, q)$. This will entail that $\mu_{i-1}(e_{i-2}) \prec \mu_i(e_{i-1})$, so that the care in finding q is in order to ensure that $q < \min \mu(e_k)$ for every $k > i-1$. Define $U_{i-1} = \{e_k \mid k > i-1 \wedge$

$\sup \mu(e_k) \leq \sup \mu_{i-1}(e_{i-1})\}$. As above, U_{i-1} is nonempty and the finiteness property of Lamport implies that it is finite. We let $m = \min U_{i-1}$. Then $p < m$ and by taking any q such that $p < q < m$ and defining $\mu_i(e_{i-1}) = p, q)$ we can complete the inductive definition.

Now we define the representation μ_T and complete the proof of Theorem 3.23. Define $\mu_T(e_j) = \mu_{j+1}(e_j)$. Then $\mu_T(e_j) = \mu_k(e_j)$ for every $k \geq j+1$, and it follows that $\mu_T(e_j) \prec \mu_T(e_{j+1})$ for every j. □

The proof of Theorem 3.22 is given next for completeness; its importance is mainly conceptual since for applications we only need the version of Theorem 3.23. The main lemma that we use in the proof of the theorem is the following. It shows a single step towards obtaining the required refinement.

Lemma 3.24 *Let $\mathcal{S} = (E, <_S)$ and $\mathcal{T} = (E, <_T)$ be interval orderings where $\mathcal{S}$ satisfies Lamport's finiteness property. Assume that $\mathcal{T}$ augments $\mathcal{S}$. Suppose that μ_S is a representation of $\mathcal{S}$ that maps E into left-closed right-open intervals of rational numbers. Assume that $a, b \in E$ are such that $a <_T b$. Then there is μ' which is a refinement of μ_S and is such that:*

1. *$\mu'(e)$ is also a left-closed right-open (sub-interval of $\mu(e)$) for every $e \in E$.*
2. *$\mu'(a) \prec \mu'(b)$.*
3. *The ordering of E induced by μ' (which is an augmentation of $<_S$) is a subordering of $<_T$.*

Proof. In the following proof we shall find some $m \in \mathbb{Q}$ and define $\mu'(a) = \{x \in \mu(a) \mid x < m\}$ and $\mu'(b) = \{x \in \mu(b) \mid m \geq x\}$. So obviously $\mu'(a) \prec \mu'(b)$. For any $e \in E \setminus \{a, b\}$ we shall leave $\mu'(e)$ unchanged and define $\mu'(e) = \mu(e)$. Our aim is to choose m in such a way that the resulting ordering of E that μ' induces is indeed augmented by $\mathcal{T}$. Two types of problems are possible if m is not properly chosen.

1. There may be some $u \in E$ so that it is not the case that $a <_T u$ and yet $\mu'(a) \prec \mu(u)$. To avoid this problem we must chose m so that $\mathit{Begin}(\mu(u)) < m$.
2. There may be some $v \in E$ so that it is not the case that $v <_T b$ and yet $\mu(v) \prec \mu'(b)$. To avoid this problem we must chose m so that $m < \mathit{End}(v)$.

The problem is therefore in choosing m to accommodate both types of requirements and in the intersection of $\mu(a)$ and $\mu(b)$.

Let $\not\succ$ be defined by $x \not\succ y$ iff $\neg(y <_T x)$ (as in Equation 3). Define

$$U = \{u \in E \mid u \not\succ a\}$$

and

$$V = \{v \in E \mid b \not\succ v \wedge v \not\succ b\}.$$

The property of Anger (see (5))

$$u \not\succ a < b \not\succ v \textit{ implies } u \not\succ v$$

shows that if $u \in U$ and $v \in V$ then $u \not\succ v$, and hence

$$\mathit{Begin}(\mu(u)) < \mathit{End}(\mu(v)). \tag{6}$$

(Or else $\mathit{End}(\mu(v)) \leq \mathit{Begin}(\mu(u))$ would imply that $v < u$.) U and V are nonempty as $a \in U$ and $b \in V$. By the finiteness property it follows that U and V are finite. Using this, let $u_0 \in U$ be with maximal $\mathit{Begin}(\mu(u_0))$. (That is, $\mathit{Begin}(\mu(u)) \leq \mathit{Begin}(\mu(u_0))$ for all $u \in U$.) Since $a \in U$, $\mathit{Begin}(\mu(a)) \leq \mathit{Begin}(\mu(u_0))$. Let $v_0 \in V$ be with minimal $\mathit{End}(\mu(v))$ among $v \in V$. Since $b \in V$, $\mathit{End}(\mu(v_0)) \leq \mathit{End}(\mu(b))$. It follows from (6) that $\mathit{Begin}(\mu(u_0)) < \mathit{End}(\mu(v_0))$. Pick any point m with $\mathit{Begin}(\mu(u_0)) < m < \mathit{End}(\mu(v_0))$. Then define the intervals I and J by

$$I = \mu(a) \cap (-\infty, m) \text{ and } J = \mu(b) \cap [m, \infty).$$

Clearly $I \prec J$. Note that these intervals are not empty since $\mathit{Begin}(\mu(a) < m < \mathit{End}(\mu(b)$. Observe that (by maximality of u_0 and minimality of v_0):

$$\forall u \in U \; I \not\prec \mu(u), \text{ and } \forall v \in V \; \mu(v) \not\prec J. \tag{7}$$

Now set $\mu'(a) = I, \mu'(b) = J$, and $\mu'(x) = \mu(x)$ for all other x's. This defines the new refinement μ', and we let $<_{S'}$ be the ordering relation on E induced by μ'. That is $e_1 <_{S'} e_2$ iff $\mu'(e_1) \prec \mu'(e_2)$.

Claim 3.25 *$\mathcal{T}$ is an augmentation of $<_{S'}$.*

Proof. Let $e_1, e_2 \in E$ be such that $\mu'(e_1) \prec \mu'(e_2)$. We have to prove that $e_1 <_T e_2$. In case $e_1 = a$ and $e_2 = b$ then it is an assumption that $a <_T b$.

If neither a nor b are in $\{a, b\}$, then $\mu'(e_1) = \mu(e_1)$ and $\mu'(e_2) = \mu(e_2)$ and hence $\mu(e_1) \prec \mu(b)$, which implies that $e_1 <_S e_2$ (as μ is a representation) and from this $e_1 <_T e_2$ follows since $<_T$ augments $<_S$. Ihe two remaining cases are:

1. a is one of e_1 and e_2 but b is not, there are two possibilities. If $a = e_2$, the fact that $Begin(\mu'(a)) = Begin(\mu(a))$ implies that $\mu'(e_1) = \mu(e_1) \prec \mu(a)$, and hence $e_1 <_T a$ follows. The more interesting possibility is that $a = e_1$. We have that $\mu'(a) \prec \mu(e_2)$ and we have to prove that $a <_T e_2$. If this is not the case, then $e_2 \not\succ a$ and hence $e_2 \in U$.

2. b is one of e_1 and e_2 but a is not.

□

This ends the proof of the main lemma, and now we show how Theorem 3.22 can be concluded. Let $\{(a_i, b_i) \mid i \geq 0\}$ be an enumeration of all pairs in $\mathcal{T}$ that are new (i.e., pairs (a, b) such that $a <_T b$ but $\neg(a <_S b)$). Observe that any $x \in E$ can appear only finitely many times as a_i or as b_i. (An easy argument show that this is so by the finiteness property). We use the preceding lemma to define by induction on i an interval ordering $\mathcal{S}_i = (E, <_i)$ with a representation μ_i (into intervals of $\mathbb{Q}$ that are left-closed right-open) such that the following hold.

1. $\mathcal{S}_0 = \mathcal{S}$ and $\mu_0 = \mu_S$ are as given in the theorem.

2. If $k < i$ then $\mathcal{S}_i$ augments $\mathcal{S}_k$ and $\mathcal{T}$ augments $\mathcal{S}_i$. μ_i is a refinement of μ_k.

3. $\mu_{i+1}(a_i) \prec \mu_{i+1}(b_i)$.

4. For every $e \in E$ there is an integer K such that for all $j \geq K$ $\mu_j(e) = \mu_K(e)$.

Suppose that $\mathcal{S}_i$ and its representation μ_i are defined, and $\mathcal{T}$ is an augmentation of $\mathcal{S}_i$. Then we consider the pair $a_i <_T b_i$, and use the lemma to define a refinement μ_{i+1} such that

1. $\mu_{i+1}(a_i) \prec \mu_{i+1}(b_i)$.

2. For every $e \in E \setminus \{a_i, b_i\}$, $\mu_{i+1}(e) = \mu_i(e)$.

3. We define $\mathcal{S}_{i+1}$ as the interval ordering induced by μ_{i+1}. $\mathcal{T}$ augments $\mathcal{S}_{i+1}$.

For every $e \in E$ we define $\mu_T(e) = \mu_K(e)$ where K is such that item 4 above holds. That is $\mu_T(e)$ is the stable interval attached to e. We want to show that μ_T is a refinement of μ_S and a representation of $\mathcal{T}$. Clearly, if $i < j$ then for every $e \in E$ $\mu_T(e) \subseteq \mu_j(e) \subseteq \mu_i(e)$. So, for any i, if $x <_i y$ then $\mu_i(x) \prec \mu_i(y)$, and hence $\mu_T(x) \prec \mu_T(y)$. Hence if $x <_T y$ then there is some i such that $x <_i y$, and so $\mu_T(x) \prec \mu_T(y)$. Now if $\mu_T(x) \prec \mu_T(y)$ then for some K large enought $\mu_T(x) = \mu_K(x)$ and $\mu_T(y) = \mu_K(y)$. Hence $\mu_K(x) \prec \mu_K(y)$ and so $x <_K y$. But as $\mathcal{T}$ is an augmentation of $\mathcal{S}_K$, $x <_T y$ follows.

This ends the proof of Theorem 3.22. □

Corollary 3.26 *Let $(E, <)$ be an interval ordering that satisfies Lamport's finiteness property, and suppose that $E_0 \subseteq E$ is a subset. If $<_0$ is an interval ordering of E_0 that extends the restriction of $<$ on E_0, then there is an interval ordering of E, $<^*$, that extends $<$ and is such that*

1. *the restriction of $<^*$ on E_0 is $<_0$, and*

2. *the restriction of $<^*$ on $E \setminus E_0$ is equal to the restriction of $<$ to $E \setminus E_0$.*

Proof. We saw (Theorem 3.19) that $(E, <)$ has a representation μ into intervals of the natural numbers, and hence into intervals of rational numbers. By Theorem 3.22, $<_0$ can be represented by a map μ_0 defined over E_0 so that $\mu_0(x)$ is a subinterval of $\mu(x)$ for every $x \in E_0$. Now we define μ^* by the formula $\mu^*(x) = mu(x)$ for $x \in E \setminus E_0$ and $\mu^*(x) = \mu_0(x)$ for $x \in E_0$. Then the ordering $x_1 <^* x_i$ iff $\mu(x_1) \prec \mu(x_2)$ is as required. □

We will use this corollary when discussing locality of linearizability. In that application we will conclude that if $(E, <)$ is an interval ordering that satisfies Lamport's finiteness property and $<_i$ is a linear ordering on E_i that extends $<$, where $E = \bigcup_{i<k} E_i$ is a partition of E, then there is a linearization of E that extends $<$ and is equal to $<_i$ on every E_i.

3.4 Global time

This is a very short section (almost an extended footnote) which aims to give some definitions concerning the question of global time. The reader may want to pursue this issue with the papers cited at the end of the section.

Basing their arguments on modern physics, philosophers and scientists have offered alternative views to the temporal picture of a single, linearly ordered time axis that is adopted here. Lamport's analysis of distributed systems ([8]) is based on two precedence relations on operation executions $A \longrightarrow B$ for "precedence", and $A \dashrightarrow B$ which can be read as "A can affect B".

Lamport's Axioms for $\longrightarrow$ and $\dashrightarrow$ precedence relations.

1. $\longrightarrow$ is irreflexive and transitive.
2. $a \longrightarrow b$ implies $a \dashrightarrow b$ and $\neg b \dashrightarrow a$.
3. $a \longrightarrow b \dashrightarrow c$ implies $a \dashrightarrow c$. And similarly $a \dashrightarrow b \longrightarrow c$ implies $a \dashrightarrow c$.
4. $a \longrightarrow b \dashrightarrow c \longrightarrow d$ implies $a \longrightarrow d$.

The following structures were defined in [8]:

Definition 3.27 *We say that* $(E, \longrightarrow, \dashrightarrow, T)$ *is a* Lamport's structure *iff it satisfies the four axioms listed above and moreover it satisfies the finiteness property: for every terminating* $x \in E$ *(i.e., x's such that* $T(x)$*) the following set is finite.*

$$\{y \in E \mid \neg x \longrightarrow y\}$$

Lamport discusses (and rejects) the *global-time* axiom.

Global-Time Axiom

$$\text{For every } a \text{ and } b,\ a \longrightarrow b \text{ iff } \neg b \dashrightarrow a.$$

The global-time axiom is not a consequence of Lamport's four axioms, and if it is added then the temporal picture is much simplified. I prefer this simpler setting both for describing semantics of protocols and for investigating concurrency. Simplicity is a strong reason, but there are other reasons described in the papers by Ben-David [5]; Abraham, Ben-David and Magidor [2]; and Abraham, Ben-David and Moran [3]. The essence of these investigations is that, for a large family of protocols, correctness under the global-time assumption implies correctness under the more general axioms of Lamport.

An interesting investigation of Lamport's structures is in Anger [4].

References

[1] U. Abraham. *Models for Concurrency.* Gordon and Breach, 1999.

[2] U. Abraham, S. Ben-David, and M. Magidor, On global-time and inter-process communication, in: M. Z. Kwiatkowska et al., eds, *BCS-FACS Workshop on Semantics for Concurrency*, Leicester (1990) 311-323 (*Workshops in Computing*, a Springer-Verlag Series edited by C.J. van Rijsbergen).

[3] U. Abraham, S. Ben-David, and S. Moran, On the limitation of the Global-Time assumption in distributed systems (extended abstract), in S. Toueg et al., (eds.) *Distributed Algorithms*, WDAG '91 Proceedings, Lecture Notes in Computer Science **579**, pp 1–8, Springer, 1991.

[4] F. Anger, On Lamport's interprocess communication, *ACM Transactions on Programming Languages and Systems* **11** (1989) 404 - 417.

[5] S. Ben-David, The global-time assumption and semantics for distributed systems, in: *Proc. 7th Ann. ACM Symp. on Principles of Distributed Computing* (1988) 223 - 232.

[6] P. C. Fishburn, *Interval orders and interval graphs,* Wiley, New York, 1986 (Wiley-Interscience series in discrete mathematics).

[7] L. Lamport, Time, clocks, and the ordering of events in a distributed system, *C. ACM*, **21**:7 (1978) 558-565.

[8] L. Lamport, On interprocess communication, Part I: Basic formalism; Part II: Algorithms, *Distributed Computing* **1** (1986) 77 - 101.

[9] V. Pratt, Modeling Concurrency with Partial Orders, *Internat. J. of Parallel Programming* **15** (1986) 33 - 71.

[10] N. Wiener, A contribution to the theory of relative position, *Proc. Camb. Philos. Soc.* **17** (1914) 441-449.

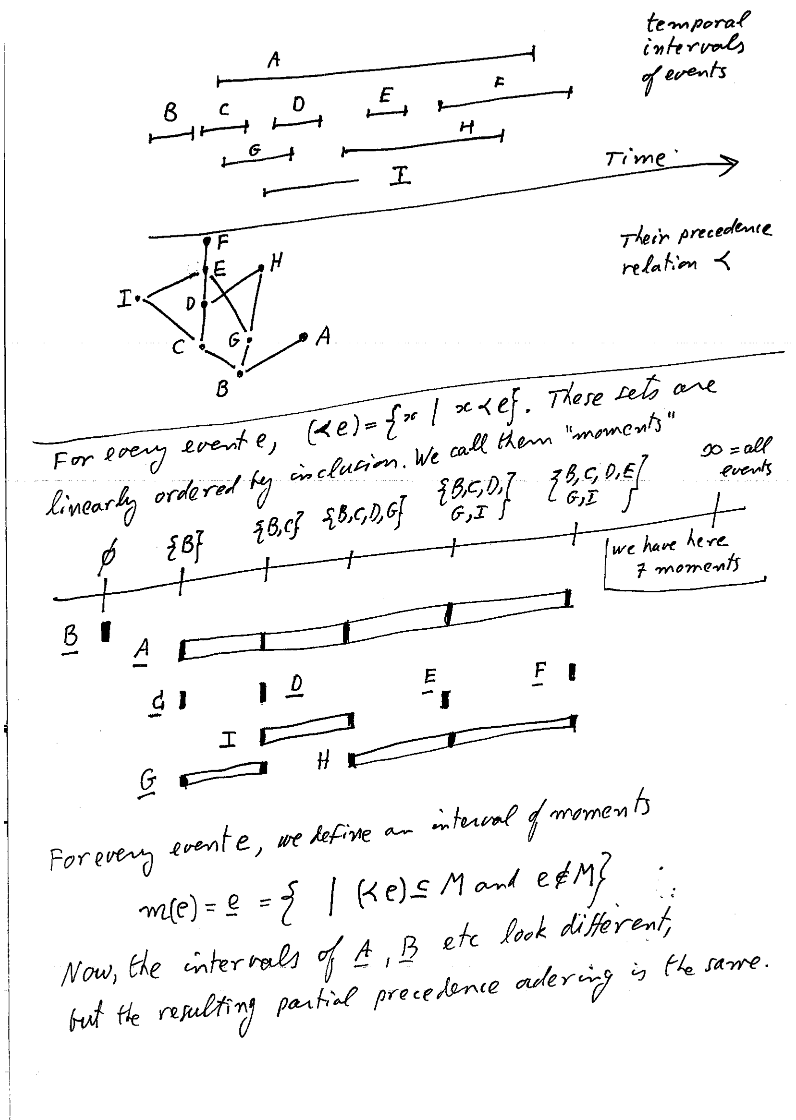

Figure 1: An example for Theorem 3.6