91

An Old-Fashioned Recipe for Real Time

Martín Abadi and Leslie Lamport

October 12, 1992

Systems Research Center

DEC's business and technology objectives require a strong research program. The Systems Research Center (SRC) and three other research laboratories are committed to filling that need.

SRC began recruiting its first research scientists in l984—their charter, to advance the state of knowledge in all aspects of computer systems research. Our current work includes exploring high-performance personal computing, distributed computing, programming environments, system modelling techniques, specification technology, and tightly-coupled multiprocessors.

Our approach to both hardware and software research is to create and use real systems so that we can investigate their properties fully. Complex systems cannot be evaluated solely in the abstract. Based on this belief, our strategy is to demonstrate the technical and practical feasibility of our ideas by building prototypes and using them as daily tools. The experience we gain is useful in the short term in enabling us to refine our designs, and invaluable in the long term in helping us to advance the state of knowledge about those systems. Most of the major advances in information systems have come through this strategy, including time-sharing, the ArpaNet, and distributed personal computing.

SRC also performs work of a more mathematical flavor which complements our systems research. Some of this work is in established fields of theoretical computer science, such as the analysis of algorithms, computational geometry, and logics of programming. The rest of this work explores new ground motivated by problems that arise in our systems research.

DEC has a strong commitment to communicating the results and experience gained through pursuing these activities. The Company values the improved understanding that comes with exposing and testing our ideas within the research community. SRC will therefore report results in conferences, in professional journals, and in our research report series. We will seek users for our prototype systems among those with whom we have common research interests, and we will encourage collaboration with university researchers.

Robert W. Taylor, Director

An Old-Fashioned Recipe for Real Time

Martín Abadi and Leslie Lamport

October 12, 1992

Authors' Abstract

Traditional methods for specifying and reasoning about concurrent systems work for real-time systems. Using TLA (the temporal logic of actions), we illustrate how they work with the examples of a queue and of a mutual-exclusion protocol. In general, two problems must be addressed: avoiding the real-time programming version of Zeno's paradox, and coping with circularities when composing real-time assumption/guarantee specifications. Their solutions rest on properties of machine closure and realizability.

Contents

1 Introduction

A new class of systems is often viewed as an opportunity to invent a new semantics. A number of years ago, the new class was distributed systems. More recently, it has been real-time systems. The proliferation of new semantics may be fun for semanticists, but developing a practical method for reasoning formally about systems is a lot of work. It would be unfortunate if every new class of systems required inventing new semantics, along with proof rules, languages, and tools.

Fortunately, no fundamental change to the old methods for specifying and reasoning about systems is needed for these new classes. It has long been known that the methods originally developed for shared-memory multiprocessing apply equally well to distributed systems [7, 11]. The first application we have seen of a clearly "off-the-shelf" method to a real-time algorithm was in 1983 [16], but there were probably earlier ones. Indeed, the "extension" of an existing temporal logic to real-time programs by Bernstein and Harter in 1981 [6] can be viewed as an application of that logic.

The old-fashioned methods handle real time by introducing a variable, which we call *now*, to represent time. This idea is so simple and obvious that it seems hardly worth writing about, except that few people appear to be aware that it works in practice. We therefore describe how to apply a conventional method to real-time systems.

Any formalism for reasoning about concurrent programs can be used to prove properties of real-time systems. However, in a conventional formalism based on a programming language, real-time assumptions are expressed by adding program operations that read and modify the variable *now*. The result can be a complicated program that is hard to understand and easy to get wrong. We take as our formalism TLA, the temporal logic of actions [13]. In TLA, programs and properties are represented as logical formulas. A real-time program can be written as the conjunction of its untimed version, expressed in a standard way as a TLA formula, and its timing assumptions, expressed in terms of a few standard parameterized formulas. This separate specification of timing properties makes real-time specifications easier to write and understand.

The method is illustrated with two examples. The first is a queue in which the sender and receiver synchronize by the use of timing assumptions instead of acknowledgements. We indicate how safety and liveness properties of the queue can be proved. The second example is an n-process mutual exclusion protocol, in which mutual exclusion depends on assumptions about the length of time taken by the operations. Its correctness is proved by a conventional invariance argument.

We also discuss two problems that arise when time is represented as a program variable—problems that seem to have received little attention—and present new solutions. The solutions are expressed in terms of TLA, but they can be applied to any formalism whose semantics is based on sequences of states or actions.

The first problem is how to avoid the real-time programming version of Zeno's paradox. If time becomes an ordinary program variable, then one can inadvertently write programs in which time behaves improperly. An obvious danger is deadlock, where time stops. A more insidious possibility is that time keeps advancing but is bounded, approaching closer and closer to some limit. One way to avoid such "Zeno" behaviors is to place an a priori lower bound on the duration of any action, but this can complicate the representation of some systems. We provide a more general and, we feel, a more natural solution.

The second problem is coping with the circularity that arises in open system specifications. The specification of an open system asserts that it operates correctly under some assumptions on the system's environment. A modular specification method requires a rule asserting that, if each component satisfies its specification, then it behaves correctly in concert with other components. This rule is circular, because a component's specification requires only that it behave correctly if its environment does, and its environment consists of all the other components. Despite its circularity, the rule is sound for specifications written in a particular style [1, 15, 17]. By examining an apparently paradoxical example, we discover how real-time specifications of open systems can be written in this style.

2 Closed Systems

We briefly review how to represent closed systems in TLA. A closed system is one that is self-contained and does not communicate with an environment. No one intentionally designs autistic systems; in a closed system, the environment is represented as part of the system. Open systems, in which the environment and system are separated, are discussed in Section 4.

We begin our review of TLA with an example. Section 2.2 summarizes the formal definitions. A more leisurely exposition appears in [13], and most definitions in the current paper are repeated in a list in the appendix. Section 2.3 reviews the concepts of safety [4] and machine closure [2] (also known as feasibility [5]) and relates them to TLA, and Section 2.4 defines a useful class of history variables [2].

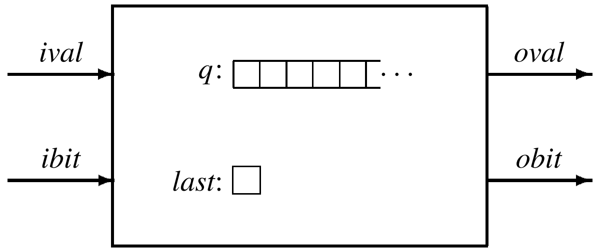

Figure 1: A simple queue.

Propositions and theorems are proved in the appendix.

2.1 The Lossy-Queue Example

We introduce TLA with the example of the lossy queue shown in Figure 1. The interface consists of two pairs of "wires", each pair consisting of a *val* wire that holds a message and a boolean-valued *bit* wire. A message m is sent over a pair of wires by setting the *val* wire to m and complementing the *bit* wire. Input to the queue arrives on the wire pair (*ival*, *ibit*), and output is sent on the wire pair (*oval*, *obit*). There is no acknowledgment protocol, so inputs are lost if they arrive faster than the queue processes them. The property guaranteed by this lossy queue is that the sequence of output messages is a subsequence of the sequence of input messages. In Section 3.1, we add timing constraints to rule out the possibility of lost messages.

A specification is a TLA formula Π describing a set of allowed behaviors. A property P is also a TLA formula. The specification Π satisfies property P iff (if and only if) every behavior allowed by Π is also allowed by P—that is, if Π implies P. Similarly, a specification Ψ implements Π iff every behavior allowed by Ψ is also allowed by Π, so implementation means implication.

The specification of the lossy queue is a TLA formula that mentions the four variables *ibit*, *obit*, *ival*, and *oval*, as well as two internal variables: q, which equals the sequence of messages received but not yet output; and *last*, which equals the value of *ibit* for the last received message. (The variable *last* is used to prevent the same message from being received twice.) These six variables are *flexible variables*; their values can change during a behavior. We also introduce a *rigid variable* Msg denoting the set of possible messages; it has the same value

$$
\begin{array}{lcl}
Init_Q & \triangleq & \land\ ibit, obit \in \{\mathsf{true}, \mathsf{false}\} \\
 & & \land\ ival, oval \in \mathsf{Msg} \\
 & & \land\ last = ibit \\
 & & \land\ q = \langle\langle\rangle\rangle \\
Inp & \triangleq & \land\ ibit' = \neg ibit \\
 & & \land\ ival' \in \mathsf{Msg} \\
 & & \land\ (obit, oval, q, last)' = (obit, oval, q, last) \\
EnQ & \triangleq & \land\ last \neq ibit \\
 & & \land\ q' = q \circ \langle\langle ival\rangle\rangle \\
 & & \land\ last' = ibit \\
 & & \land\ (ibit, obit, ival, oval)' = (ibit, obit, ival, oval) \\
DeQ & \triangleq & \land\ q \neq \langle\langle\rangle\rangle \\
 & & \land\ oval' = Head(q) \\
 & & \land\ q' = Tail(q) \\
 & & \land\ obit' = \neg obit \\
 & & \land\ (ibit, ival, last)' = (ibit, ival, last) \\
\mathcal{N}_Q & \triangleq & Inp \lor EnQ \lor DeQ \\
v & \triangleq & (ibit, obit, ival, oval, q, last) \\
\Pi_Q & \triangleq & Init_Q \land \Box[\mathcal{N}_Q]_v \\
\Phi_Q & \triangleq & \exists q, last : \Pi_Q
\end{array}
$$

Figure 2: The TLA specification of a lossy queue.

throughout a behavior. We usually refer to flexible variables simply as variables, and to rigid variables as *constants*.

The TLA specification is shown in Figure 2, using the following notation. A list of formulas, each prefaced by $\land$, denotes the conjunction of the formulas, and indentation is used to eliminate parentheses. The expression $\langle\langle\ \rangle\rangle$ denotes the empty sequence, $\langle\langle m\rangle\rangle$ denotes the singleton sequence having m as its one element, $\circ$ denotes concatenation, $Head(\sigma)$ denotes the first element of σ, and $Tail(\sigma)$ denotes the sequence obtained by removing the first element of σ. The symbol $\triangleq$ means *is defined to equal*.

The first definition is of the *predicate* $Init_Q$, which describes the initial state. This predicate asserts that the values of variables *ibit* and *obit* are arbitrary booleans, the values of *ival* and *oval* are elements of Msg, the values of *last* and *ibit* are equal, and the value of q is the empty sequence.

Next is defined the *action Inp*, which describes all state changes that represent the sending of an input message. (Since this is the specification of a closed system, it includes the environment's *Inp* action.) The first conjunct, $ibit' = \neg ibit$, asserts that the new value of *ibit* equals the complement of its old value. The second conjunct asserts that the new value of *ival* is an element of Msg. The third conjunct asserts that the value of the four-tuple (*obit*, *oval*, q, *last*) is unchanged; it is equivalent to the assertion that the value of each of the four variables *obit*, *oval*, q, and *last* is unchanged. The action *Inp* is always *enabled*, meaning that, in any state, a new input message can be sent.

Action *EnQ* represents the receipt of a message by the system. The first conjunct asserts that *last* is not equal to *ibit*, so the message on the input wire has not yet been received. The second conjunct asserts that the new value of q equals the sequence obtained by concatenating the old value of *ival* to the end of q's old value. The third conjunct asserts that the new value of *last* equals the old value of *ibit*. The final conjunct asserts that the values of *ibit*, *obit*, *ival*, and *oval* are unchanged. Action *EnQ* is enabled in a state iff the values of *last* and *ibit* in that state are unequal.

The action *DeQ* represents the operation of removing a message from the head of q and sending it on the output wire. It is enabled iff the value of q is not the empty sequence.

The action $\mathcal{N}_Q$ is the specification's *next-state relation*. It describes all allowed changes to the queue system's variables. Since the only allowed changes are the ones described by the actions *Inp*, *EnQ*, and *DeQ*, action $\mathcal{N}_Q$ is the disjunction of those three actions.

In TLA specifications, it is convenient to give a name to the tuple of all relevant variables. Here, we call it v.

Formula Π_Q is the internal specification of the lossy queue—the formula specifying all sequences of values that may be assumed by the queue's six variables, including the internal variables q and *last*. Its first conjunct asserts that $Init_Q$ is true in the initial state. Its second conjunct, $\Box[\mathcal{N}_Q]_v$, asserts that every step is either an $\mathcal{N}_Q$ step (a state change allowed by $\mathcal{N}_Q$) or else leaves v unchanged, meaning that it leaves all six variables unchanged.

Formula Φ_Q is the actual specification, in which the internal variables q and *last* have been hidden. A behavior satisfies Φ_Q iff there is some way to assign sequences of values to q and *last* such that Π_Q is satisfied. The free variables of Φ_Q are *ibit*, *obit*, *ival*, and *oval*, so Φ_Q specifies what sequences of values these four variables can assume. All the preceding definitions just represent one possible way of structuring the definition of Φ_Q; there are infinitely many ways to write formulas that are equivalent to Φ_Q and are therefore equivalent specifications.

TLA is an untyped logic; a variable may assume any value. Type correctness is expressed by the formula $\Box T$, where T is the predicate asserting that all relevant variables have values of the expected "types". For the internal queue specification, the type-correctness predicate is

$$T_Q \triangleq \begin{array}{l} \wedge\ \mathit{ibit}, \mathit{obit}, \mathit{last} \in \{\mathsf{true}, \mathsf{false}\} \\ \wedge\ \mathit{ival}, \mathit{oval} \in \mathsf{Msg} \\ \wedge\ q \in \mathsf{Msg}^* \end{array} \qquad (1)$$

where Msg^* is the set of finite sequences of messages. Type correctness of Π_Q is asserted by the formula $\Pi_Q \Rightarrow \Box T_Q$, which is easily proved [13]. Type correctness of Φ_Q follows from $\Pi_Q \Rightarrow \Box T_Q$ by the usual rules for reasoning about quantifiers.

Formulas Π_Q and Φ_Q are *safety properties*, meaning that they are satisfied by an infinite behavior iff they are satisfied by every finite initial portion of the behavior. Safety properties allow behaviors in which a system performs properly for a while and then the values of all variables are frozen, never to change again. In asynchronous systems, such undesirable behaviors are ruled out by adding *fairness* properties. We could strengthen our lossy-queue specification by conjoining the *weak fairness* property $\mathrm{WF}_v(\mathit{DeQ})$ and the *strong fairness* property $\mathrm{SF}_v(\mathit{EnQ})$ to Π_Q, obtaining

$$\exists\, q, \mathit{last} : (\mathit{Init}_Q \wedge \Box[\mathcal{N}_Q]_v \wedge \mathrm{WF}_v(\mathit{DeQ}) \wedge \mathrm{SF}_v(\mathit{EnQ})) \qquad (2)$$

Property $\mathrm{WF}_v(\mathit{DeQ})$ asserts that if action *DeQ* is enabled forever, then infinitely many *DeQ* steps must occur. This property implies that every message reaching the queue is eventually output. Property $\mathrm{SF}_v(\mathit{EnQ})$ asserts that if action *EnQ* is enabled infinitely often, then infinitely many *EnQ* steps must occur. It implies that if infinitely many inputs are sent, then the queue must receive infinitely many of them. The formula (2) implies the *liveness property* [4] that an infinite number of inputs produces an infinite number of outputs. This formula also implies the same safety properties as Φ_Q. A formula such as (2), which is the conjunction of an initial predicate, a term of the form $\Box[\mathcal{A}]_f$, and a fairness property, is said to be in *canonical form*.

2.2 The Semantics of TLA

We begin with some definitions. We assume a set of constant *values*, and we let $[\![F]\!]$ denote the semantic meaning of a formula F.

state A mapping from variables to values. We let $s.x$ denote the value that state s assigns to variable x.

state function An expression formed from variables, constants, and operators. The meaning of a state function is a mapping from states to values. For example, $x + 1$ is a state function such that $[\![x + 1]\!](s)$ equals $s.x + 1$, for any state s.

predicate A boolean-valued state function, such as $x > y + 1$.

transition function An expression formed from variables, primed variables, constants, and operators. The meaning of a transition function is a mapping from pairs of states to values. For example, $x + y' + 1$ is a transition function, and $[\![x + y' + 1]\!](s, t)$ equals the value $s.x + t.y + 1$, for any pair of states s, t.

action A boolean-valued transition function, such as $x > (y' + 1)$.

step A pair of states s, t. It is called an $\mathcal{A}$ *step* iff $[\![\mathcal{A}]\!](s, t)$ equals true, for an action $\mathcal{A}$. It is called a *stuttering* step iff $s = t$.

f' The transition function obtained from the state function f by priming all the free variables of f, so $[\![f']\!](s, t) = [\![f]\!](t)$ for any states s and t.

$[\mathcal{A}]_f$ The action $\mathcal{A} \vee (f' = f)$, for any action $\mathcal{A}$ and state function f.

$\langle\mathcal{A}\rangle_f$ The action $\mathcal{A} \wedge (f' \neq f)$, for any action $\mathcal{A}$ and state function f.

$Enabled\ \mathcal{A}$ For any action $\mathcal{A}$, the predicate such that $[\![Enabled\ \mathcal{A}]\!](s)$ equals $\exists t : [\![\mathcal{A}]\!](s, t)$, for any state s.

Informally, we often identify a formula and its meaning. For example we say that a predicate P is true in state s instead of $[\![P]\!](s) = \mathsf{true}$.

An RTLA (raw TLA) formula is an expression built from actions, classical operators (boolean operators and quantification over rigid variables), and the unary temporal operator $\Box$. The meaning of an RTLA formula is a boolean-valued function on *behaviors*, where a behavior is an infinite sequence of states. The meaning of the operator $\Box$ is defined by

$$[\![\Box F]\!](s_1, s_2, s_3, \ldots) \;\triangleq\; \forall\, n > 0 : [\![F]\!](s_n, s_{n+1}, s_{n+2}, \ldots)$$

Intuitively, $\Box F$ asserts that F is "always" true. The meaning of an action as an RTLA formula is defined in terms of its meaning as an action by letting $[\![\mathcal{A}]\!](s_1, s_2, s_3, \ldots)$ equal $[\![\mathcal{A}]\!](s_1, s_2)$. A predicate P is an action; P is true for a behavior iff it is true for the first state of the behavior, and $\Box P$ is true iff P is true in all states. For any action $\mathcal{A}$ and state function f, the formula $\Box[\mathcal{A}]_f$ is true for a behavior iff each step is an $\mathcal{A}$ step or else leaves f unchanged. The classical operators have their usual meanings.

A TLA formula is one that can be constructed from predicates and formulas $\Box[\mathcal{A}]_f$ using classical operators, $\Box$, and existential quantification over flexible variables. The semantics of actions, classical operators, and $\Box$ are defined as before. The approximate meaning of quantification over a flexible variable is that $\exists x : F$ is true for a behavior iff there is some sequence of values that can be assigned to x that makes F true. The precise definition appears in [13] and is recalled in the appendix. As usual, we write $\exists\, x_1, \ldots, x_n : F$ instead of $\exists\, x_1 : \ldots, \exists\, x_n : F$.

A *property* is a set of behaviors that is *invariant under stuttering*, meaning that it contains a behavior σ iff it contains every behavior obtained from σ by adding and/or removing stuttering steps. The set of all behaviors satisfying a TLA formula is a property, which we often identify with the formula.

For any TLA formula F, action $\mathcal{A}$, and state function f:

$$\begin{array}{lcl}
\Diamond F & \triangleq & \neg\Box\neg F \\
\mathrm{WF}_f(\mathcal{A}) & \triangleq & \Box\Diamond\neg(\mathit{Enabled}\ \langle\mathcal{A}\rangle_f) \ \vee\ \Box\Diamond\langle\mathcal{A}\rangle_f \\
\mathrm{SF}_f(\mathcal{A}) & \triangleq & \Diamond\Box\neg(\mathit{Enabled}\ \langle\mathcal{A}\rangle_f) \ \vee\ \Box\Diamond\langle\mathcal{A}\rangle_f
\end{array}$$

These are TLA formulas, since $\Diamond\langle\mathcal{A}\rangle_f$ equals $\neg\Box[\neg\mathcal{A}]_f$.

2.3 Safety and Fairness

A *finite behavior* is a finite sequence of states. We say that a finite behavior satisfies a property F iff it can be continued to an infinite behavior in F. A property F is a *safety property* [4] iff the following condition holds: F contains a behavior iff it is satisfied by every finite prefix of the behavior.[1] Intuitively, a safety property asserts that something "bad" does not happen. Predicates and formulas of the form $\Box[\mathcal{A}]_f$ are safety properties.

Safety properties form the closed sets for a topology on the set of all behaviors. Hence, if two TLA formulas F and G are safety properties, then $F \wedge G$ is also a safety property. The *closure* $\mathcal{C}(F)$ of a property F is the smallest safety property containing F. It can be shown that $\mathcal{C}(F)$ is expressible in TLA, for any TLA formula F.

If Π is a safety property and L an arbitrary property, then the pair (Π, L) is *machine closed* iff every finite behavior satisfying Π can be extended to an infinite behavior in $\Pi \wedge L$. Proposition 1 below shows that machine closure generalizes the concept of fairness. The *canonical form* for a TLA formula is

$$\exists\, x : (\mathit{Init} \;\wedge\; \Box[\mathcal{N}]_v \;\wedge\; L) \tag{3}$$

where $(\mathit{Init} \wedge \Box[\mathcal{N}]_v,\, L)$ is machine closed and x is a tuple of variables called the *internal variables* of the formula. The state function v will usually be the tuple of all variables appearing free in *Init*, $\mathcal{N}$, and L (including the variables of x). A behavior satisfies (3) iff there is some way of choosing values for x such that (a) *Init* is true in the initial state, (b) every step is either an $\mathcal{N}$ step or leaves all the variables in v unchanged, and (c) the entire behavior satisfies L.

An action $\mathcal{A}$ is said to be a *subaction* of a safety property Π iff for every finite behavior $s_1, \ldots, s_n$ satisfying Π with $\mathit{Enabled}\ \mathcal{A}$ true in state s_n, there exists a state s_{n+1} such that (s_n, s_{n+1}) is an $\mathcal{A}$ step and $s_1, \ldots, s_{n+1}$ satisfies Π. By this definition, $\mathcal{A}$ is a subaction of $\mathit{Init} \wedge \Box[\mathcal{N}]_v$ iff[2]

$$\mathit{Init} \wedge \Box[\mathcal{N}]_v \;\Rightarrow\; \Box((\mathit{Enabled}\ \mathcal{A}) \Rightarrow \mathit{Enabled}\ (\mathcal{A} \wedge [\mathcal{N}]_v))$$

Two actions $\mathcal{A}$ and $\mathcal{B}$ are *disjoint for* a safety property Π iff no behavior satisfying Π contains an $\mathcal{A} \wedge \mathcal{B}$ step. By this definition, $\mathcal{A}$ and $\mathcal{B}$ are disjoint for $\mathit{Init} \wedge \Box[\mathcal{N}]_v$ iff

$$\mathit{Init} \wedge \Box[\mathcal{N}]_v \;\Rightarrow\; \Box\neg \mathit{Enabled}\ (\mathcal{A} \wedge \mathcal{B} \wedge [\mathcal{N}]_v)$$

The following result shows that the conjunction of WF and SF formulas is a fairness property. It is a special case of Proposition 4 of Section 4.

[1] One sometimes defines $s_1, \ldots, s_n$ to satisfy F iff the behavior $s_1, \ldots, s_n, s_n, s_n, \ldots$ is in F. Since properties are invariant under stuttering, this alternative definition leads to the same definition of a safety property.

[2] We let $\Rightarrow$ have lower precedence than the other boolean operators.

Proposition 1 *If* Π *is a safety property and* L *is the conjunction of a finite or countably infinite number of formulas of the form* $\mathrm{WF}_w(\mathcal{A})$ *and/or* $\mathrm{SF}_w(\mathcal{A})$ *such that each* $\langle \mathcal{A} \rangle_w$ *is a subaction of* Π*, then* (Π, L) *is machine closed.*

In practice, each w will usually be a tuple of variables changed by the corresponding action $\mathcal{A}$, so $\langle \mathcal{A} \rangle_w$ will equal $\mathcal{A}$.[3] In the informal exposition, we often omit the subscript and talk about $\mathcal{A}$ when we really mean $\langle \mathcal{A} \rangle_w$.

Machine closure for more general classes of properties can be proved with the following two propositions, which are proved in the appendix. To apply the first, one must prove that $\exists x : \Pi$ is a safety property. By Proposition 2 of [2, page 265], it suffices to prove that Π has finite internal nondeterminism (fin), with x as its internal state component. Here, fin means roughly that there are only a finite number of sequences of values for x that can make a finite behavior satisfy Π.

Proposition 2 *If* (Π, L) *is machine closed,* x *is a tuple of variables that do not occur free in* L*, and* $\exists x : \Pi$ *is a safety property, then* $((\exists x : \Pi), L)$ *is machine closed.*

Proposition 3 *If* (Π, L_1) *is machine closed and* $\Pi \wedge L_1$ *implies* L_2*, then* (Π, L_2) *is machine closed.*

2.4 History-Determined Variables

A *history-determined* variable is one whose current value can be inferred from the current and past values of other variables. For the precise definition, let

$$Hist(h, f, g, v) \;\stackrel{\Delta}{=}\; (h = f) \;\wedge\; \Box[(h' = g) \wedge (v' \neq v)]_{\langle h, v \rangle} \tag{4}$$

where f and v are state functions and g is a transition function. A variable h is a history-determined variable for a formula Π iff Π implies $Hist(h, f, g, v)$, for some f, g, and v such that h occurs free in neither f nor v, and h' does not occur free in g.

If f and v do not depend on h, and g does not depend on h', then $\exists h : Hist(h, f, g, v)$ is identically true. Therefore, if h does not occur free in formula Φ, then $\exists h : (\Phi \wedge Hist(h, f, g, v))$ is equivalent to Φ. In other words, conjoining $Hist(h, f, g, v)$ to Φ does not change the behavior of its variables, so it makes h a "dummy variable" for Φ—in fact, it is a special kind of history variable [2, page 270].

[3] More precisely, $T \wedge \mathcal{A}$ will imply $w' \neq w$, where T is the type-correctness invariant.

As an example, we add to the lossy queue's specification Φ_Q a history variable *hin* that records the sequence of values transmitted on the input wire. Let

$$H_{in} \stackrel{\Delta}{=} \begin{array}[t]{l} \wedge\ hin = \langle\!\langle\ \rangle\!\rangle \\ \wedge\ \Box[\begin{array}[t]{l} \wedge\ hin' = hin \circ \langle\!\langle ival' \rangle\!\rangle \\ \wedge\ (ival, ibit)' \neq (ival, ibit) \end{array}\ ]_{(hin, ival, ibit)} \end{array} \tag{5}$$

Then H_{in} equals $Hist(hin, \langle\!\langle\ \rangle\!\rangle, hin \circ \langle\!\langle ival' \rangle\!\rangle, (ival, ibit))$, so *hin* is a history-determined variable for $\Phi_Q \wedge H_{in}$, and $\exists\, hin : (\Phi_Q \wedge H_{in})$ equals Φ_Q.

If h is a history-determined variable for a property Π, then Π is fin, with h as its internal state component. Hence, if Π is a safety property, then $\exists h : \Pi$ is also a safety property.

3 Real-Time Closed Systems

We now use TLA to specify and reason about timing properties of closed systems. Section 3.1 explains how time and timing properties can be represented with TLA formulas, and Section 3.2 describes how to reason about these formulas. The problem of Zeno specifications is addressed in Section 3.3. Our method of specifying and reasoning about timing properties is illustrated in Section 3.4 with the example of a real-time mutual exclusion protocol.

3.1 Time and Timers

In real-time TLA specifications, real time is represented by the variable *now*. Although it has a special interpretation, *now* is just an ordinary variable of the logic. The value of *now* is always a real number, and it never decreases—conditions expressed by the TLA formula

$$RT \stackrel{\Delta}{=} (now \in \mathbf{R}) \wedge \Box[now' \in (now, \infty)]_{now}$$

where $\mathbf{R}$ is the set of real numbers and (r, ∞) is $\{t \in \mathbf{R} : t > r\}$.

It is convenient to make time-advancing steps distinct from ordinary program steps. This is done by strengthening the formula RT to

$$RT_v \stackrel{\Delta}{=} (now \in \mathbf{R}) \wedge \Box[(now' \in (now, \infty)) \wedge (v' = v)]_{now}$$

This property differs from *RT* only in asserting that v does not change when *now* advances. Simple logical manipulation shows that RT_v is equivalent to $RT \wedge \Box[now' = now]_v$, and

$$Init \wedge \Box[\mathcal{N}]_v \wedge RT_v \;=\; Init \wedge \Box[\mathcal{N} \wedge (now' = now)]_v \wedge RT$$

Real-time constraints are imposed by using *timers* to restrict the increase of *now*. A *timer for* Π is a state function t such that Π implies $\Box(t \in \mathbf{R} \cup \{\pm\infty\})$. Timer t is used as an upper-bound timer by conjoining the formula

$$MaxTime(t) \;\stackrel{\Delta}{=}\; (now \leq t) \;\wedge\; \Box[now' \leq t']_{now}$$

to a specification. This formula asserts that *now* is never advanced past t. Timer t is used as a lower-bound timer for an action $\mathcal{A}$ by conjoining the formula

$$MinTime(t,\mathcal{A},v) \;\stackrel{\Delta}{=}\; \Box[\mathcal{A} \Rightarrow (t \leq now)]_v$$

to a specification. This formula asserts that an $\langle\mathcal{A}\rangle_v$ step cannot occur when *now* is less than t.[4]

A common type of timing constraint asserts that an $\mathcal{A}$ step must occur within δ seconds of when the action $\mathcal{A}$ becomes enabled, for some constant δ. After an $\mathcal{A}$ step, the next $\mathcal{A}$ step must occur within δ seconds of when action $\mathcal{A}$ is re-enabled. There are at least two reasonable interpretations of this requirement.

The first interpretation is that the $\mathcal{A}$ step must occur if $\mathcal{A}$ has been continuously enabled for δ seconds. This is expressed by $MaxTime(t)$ when t is a state function satisfying

$$\begin{array}{rcl} VTimer(t, A, \delta, v) & \stackrel{\Delta}{=} & \wedge\; t = \textbf{if } Enabled\ \langle\mathcal{A}\rangle_v \textbf{ then } now + \delta \\ & & \qquad\qquad\qquad\quad \textbf{else } \infty \\ & & \wedge\; \Box[\;\wedge\; t' = \textbf{if } (Enabled\ \langle\mathcal{A}\rangle_v)' \\ & & \qquad\qquad\quad \textbf{then if } \langle\mathcal{A}\rangle_v \vee \neg Enabled\ \langle\mathcal{A}\rangle_v \\ & & \qquad\qquad\qquad\quad \textbf{then } now + \delta \\ & & \qquad\qquad\qquad\quad \textbf{else } t \\ & & \qquad\qquad\quad \textbf{else } \infty \\ & & \quad\;\; \wedge\; v' \neq v\;\;]_{\langle t, v\rangle} \end{array}$$

[4]Unlike the usual timers in computer systems that represent an increment of time, our timers represent an absolute time. To allow the type of strict time bound that would be expressed by replacing $\leq$ with $<$ in the definition of *MaxTime* or *MinTime*, we could introduce, as additional possible values for timers, the set of all "infinitesimally shifted" real numbers r^-, where $t \leq r^-$ iff $t < r$, for any reals t and r.

Such a t is called a *volatile δ-timer*.

Another interpretation of the timing requirement is that an $\mathcal{A}$ step must occur if $\mathcal{A}$ has been enabled for a total of δ seconds, though not necessarily continuously enabled. This is expressed by $MaxTime(t)$ when t satisfies

$$
\begin{array}{rcl}
PTimer(t, A, \delta, v) & \triangleq & \land\ t = now + \delta \\
& & \land\ \Box[\ \land\ t' = \textbf{if } Enabled\ \langle \mathcal{A} \rangle_v \\
& & \qquad\qquad\quad \textbf{then if } \langle \mathcal{A} \rangle_v \textbf{ then } now + \delta \\
& & \qquad\qquad\qquad\qquad\quad \textbf{else } t \\
& & \qquad\qquad\quad \textbf{else } t + (now' - now) \\
& & \qquad \land\ (v, now)' \neq (v, now)\ ]_{\langle t, v, now \rangle}
\end{array}
$$

Such a t is called a *persistent δ-timer*. We can use δ-timers as lower-bound timers as well as upper-bound timers.

Observe that $VTimer(t, \mathcal{A}, \delta, v)$ has the form $Hist(t, f, g, v)$ and that $PTimer(t, \mathcal{A}, \delta, v)$ has the form $Hist(t, f, g, (v, now))$, where $Hist$ is defined by (4). Thus, if formula Π implies that a variable t satisfies either of these formulas, then t is a history-determined variable for Π.

As an example of the use of timers, we make the lossy queue of Section 2.1 nonlossy by adding the following timing constraints.

- Values must be put on a wire at most once every δ_{snd} seconds. There are two conditions—one on the input wire and one on the output wire. They are expressed by using δ_{snd}-timers t_{Inp} and t_{DeQ}, for the actions *Inp* and *DeQ*, as lower-bound timers.
- A value must be added to the queue at most Δ_{rcv} seconds after it appears on the input wire. This is expressed by using a Δ_{rcv}-timer T_{EnQ}, for the enqueue action, as an upper-bound timer.
- A value must be sent on the output wire within Δ_{snd} seconds of when it reaches the head of the queue. This is expressed by using a Δ_{snd}-timer T_{DeQ}, for the dequeue action, as an upper-bound timer.

The timed queue will be nonlossy if $\Delta_{rcv} < \delta_{snd}$. In this case, we expect the *Inp*, *EnQ*, and *DeQ* actions to remain enabled until they are "executed", so it doesn't matter whether we use volatile or persistent timers. We use volatile timers because they are a little easier to reason about.

The timed version Π^t_Q of the queue's internal specification Π_Q is obtained by conjoining the timing constraints to Π_Q:

$$
\begin{array}{rcl}
\Pi^t_Q & \stackrel{\Delta}{=} & \land\ \Pi_Q\ \land\ RT_v \\
& & \land\ \mathit{VTimer}(t_{\mathit{Inp}}, \mathit{Inp}, \delta_{snd}, v)\ \land\ \mathit{MinTime}(t_{\mathit{Inp}}, \mathit{Inp}, v) \\
& & \land\ \mathit{VTimer}(t_{\mathit{DeQ}}, \mathit{DeQ}, \delta_{snd}, v)\ \land\ \mathit{MinTime}(t_{\mathit{DeQ}}, \mathit{DeQ}, v) \\
& & \land\ \mathit{VTimer}(T_{\mathit{EnQ}}, \mathit{EnQ}, \Delta_{rcv}, v)\ \land\ \mathit{MaxTime}(T_{\mathit{EnQ}}) \\
& & \land\ \mathit{VTimer}(T_{\mathit{DeQ}}, \mathit{DeQ}, \Delta_{snd}, v)\ \land\ \mathit{MaxTime}(T_{\mathit{DeQ}})
\end{array}
\tag{6}
$$

The external specification Φ^t_Q of the timed queue is obtained by existentially quantifying first the timers and then the variables q and *last*.

Formula Π^t_Q of (6) is not in the canonical form for a TLA formula. A straightforward calculation, using the type-correctness invariant (1) and the equivalence of $(\Box F) \land (\Box G)$ and $\Box(F \land G)$, converts the expression (6) for Π^t_Q to the canonical form given in Figure 3.[5] Observe how each subaction $\mathcal{A}$ of the original formula has a corresponding timed version $\mathcal{A}^t$. Action $\mathcal{A}^t$ is obtained by conjoining $\mathcal{A}$ with the appropriate relations between the old and new values of the timers. If $\mathcal{A}$ has a lower-bound timer, then $\mathcal{A}^t$ also has a conjunct asserting that it is not enabled when *now* is less than this timer. (The lower-bound timer t_{Inp} for *Inp* does not affect the enabling of other subactions because *Inp* is disjoint from all other subactions; a similar remark applies to the lower-bound timer t_{DeQ}.) There is also a new action, *QTick*, that advances *now*.

Formula Π^t_Q is the TLA specification of a program that satisfies each maximum-delay constraint by preventing *now* from advancing before the constraint has been satisfied. Thus, the program "implements" timing constraints by stopping time, an apparent absurdity. However, the absurdity results from thinking of a TLA formula, or the abstract program that it represents, as a *prescription of how* something is accomplished. A TLA formula is really a *description of what* is supposed to happen. Formula Π^t_Q says only that an action occurs before *now* reaches a certain value. It is just our familiarity with ordinary programs that makes us jump to the conclusion that *now* is being changed by the system.

3.2 Reasoning About Time

Formula Π^t_Q is a safety property; it is satisfied by a behavior in which no variables change values. In particular, it allows behaviors in which time stops. We can rule

[5] Further simplification of this formula is possible, but it requires an invariant. In particular, the fourth conjunct of DeQ^t can be replaced by $T'_{\mathit{EnQ}} = T_{\mathit{EnQ}}$.

$$
\begin{array}{lcl}
Init^t_Q & \triangleq & \land\ Init_Q \\
 & & \land\ now \in \mathbf{R} \\
 & & \land\ t_{Inp} = now + \delta_{snd} \\
 & & \land\ t_{DeQ} = T_{EnQ} = T_{DeQ} = \infty \\
Inp^t & \triangleq & \land\ Inp \\
 & & \land\ t_{Inp} \leq now \\
 & & \land\ t'_{Inp} = now' + \delta_{snd} \\
 & & \land\ T'_{EnQ}\ =\ \textbf{if}\ last' \neq ibit'\ \textbf{then}\ now' + \Delta_{rcv}\ \textbf{else}\ \infty \\
 & & \land\ (t_{DeQ}, T_{DeQ})'\ =\ \textbf{if}\ q = \langle\!\langle\ \rangle\!\rangle\ \textbf{then}\ (\infty, \infty)\ \textbf{else}\ (t_{DeQ}, T_{DeQ}) \\
 & & \land\ now' = now \\
EnQ^t & \triangleq & \land\ EnQ \\
 & & \land\ T'_{EnQ} = \infty \\
 & & \land\ (t_{DeQ}, T_{DeQ})'\ =\ \textbf{if}\ q = \langle\!\langle\ \rangle\!\rangle\ \textbf{then}\ (now + \delta_{snd},\ now + \Delta_{snd}) \\
 & & \qquad\qquad\qquad\qquad\qquad\quad \textbf{else}\ \ (t_{DeQ}, T_{DeQ}) \\
 & & \land\ (t_{Inp}, now)' = (t_{Inp}, now) \\
DeQ^t & \triangleq & \land\ DeQ \\
 & & \land\ t_{DeQ} \leq now \\
 & & \land\ (t_{DeQ}, T_{DeQ})'\ =\ \textbf{if}\ q' = \langle\!\langle\ \rangle\!\rangle\ \textbf{then}\ (\infty, \infty) \\
 & & \qquad\qquad\qquad\qquad\qquad\quad \textbf{else}\ \ (now + \delta_{snd},\ now + \Delta_{snd}) \\
 & & \land\ T'_{EnQ}\ =\ \textbf{if}\ last' = ibit'\ \textbf{then}\ \infty\ \textbf{else}\ T_{EnQ} \\
 & & \land\ (t_{Inp}, now)' = (t_{Inp}, now) \\
QTick & \triangleq & \land\ now' \in (now, \min(T_{DeQ}, T_{EnQ})] \\
 & & \land\ (v, t_{Inp}, t_{DeQ}, T_{DeQ}, T_{EnQ})' = (v, t_{Inp}, t_{DeQ}, T_{DeQ}, T_{EnQ}) \\
vt & \triangleq & (v, now, t_{Inp}, t_{DeQ}, T_{DeQ}, T_{EnQ}) \\
\Pi^t_Q & \triangleq & \land\ Init^t_Q \\
 & & \land\ \Box[Inp^t\ \lor\ EnQ^t\ \lor\ DeQ^t\ \lor\ QTick]_{vt}
\end{array}
$$

Figure 3: The canonical form for Π^t_Q, where $(r, s]$ denotes the set of reals u such that $r < u \leq s$.

out such behaviors by conjoining to Π^t_Q the liveness property

$$NZ \;\triangleq\; \forall\, t \in \mathbf{R} : \Diamond(now > t)$$

which asserts that *now* gets arbitrarily large. However, when reasoning only about real-time properties, this is not necessary. For example, suppose we want to show that our timed queue satisfies a real-time property expressed by formula Ψ^t, which is also a safety property. If Π^t_Q implies Ψ^t, then $\Pi^t_Q \wedge NZ$ implies $\Psi^t \wedge NZ$. Conversely, we don't expect conjoining a liveness property to add safety properties; if $\Pi^t_Q \wedge NZ$ implies Ψ^t, then Π^t_Q by itself should imply Ψ^t—a point discussed in Section 3.3 below. Hence, there is no need to introduce the liveness property *NZ*.

A safety property we might want to prove for the timed queue is that it does not lose any inputs. To express this property, let *hin* be the history variable, determined by H_{in} of (5), that records the sequence of input values; and let *hout* and H_{out} be the analogous history variable and property for the outputs. The assertion that the timed queue loses no inputs is expressed by

$$\Pi^t_Q \wedge H_{in} \wedge H_{out} \;\Rightarrow\; \Box(hout \preceq hinp)$$

where $\alpha \preceq \beta$ iff α is an initial prefix of β. This is a standard invariance property. The usual method for proving such properties leads to the following invariant

$$\begin{aligned}
&\wedge\; T_Q \;\wedge\; (t_{Inp}, now \in \mathbf{R}) \;\wedge\; (T_{EnQ}, t_{DeQ}, T_{DeQ} \in \mathbf{R} \cup \{\infty\})\\
&\wedge\; now \leq \min(T_{EnQ}, T_{DeQ})\\
&\wedge\; (last = ibit) \;\Rightarrow\; (T_{EnQ} = \infty) \;\wedge\; (hinp = hout \circ q)\\
&\wedge\; (last \neq ibit) \;\Rightarrow\; (T_{EnQ} < t_{Inp}) \;\wedge\; (hinp = hout \circ q \circ \langle\!\langle ival \rangle\!\rangle)\\
&\wedge\; (q = \langle\!\langle\, \rangle\!\rangle) \;\equiv\; (T_{DeQ} = \infty)
\end{aligned}$$

and to the necessary assumption $\Delta_{rcv} < \delta_{snd}$. (Recall that T_Q is the type-correctness predicate (1) for Π_Q.)

Property *NZ* is needed to prove that real-time properties imply liveness properties. The desired liveness property for the timed queue is that the sequence of input messages up to any point eventually appears as the sequence of output messages. It is expressed by

$$\Pi^t_Q \wedge NZ \;\Rightarrow\; \forall\, \sigma : \Box((hinp = \sigma) \Rightarrow \Diamond(hout = \sigma))$$

This formula is proved by first showing

$$\Pi^t_Q \wedge NZ \;\Rightarrow\; \mathrm{WF}_v(EnQ) \wedge \mathrm{WF}_v(DeQ) \tag{7}$$

and then using a standard liveness argument to prove

$$\Pi^t_Q \wedge \mathrm{WF}_v(EnQ) \wedge \mathrm{WF}_v(DeQ) \;\Rightarrow\; \forall\, \sigma : \Box((hinp = \sigma) \Rightarrow \Diamond(hout = \sigma))$$

The proof that $\Pi^t_Q \wedge NZ$ implies $\mathrm{WF}_v(EnQ)$ is by contradiction. Assume EnQ is forever enabled but never occurs. An invariance argument then shows that Π^t_Q implies that T_{EnQ} forever equals its current value, preventing *now* from advancing past that value; and this contradicts NZ. The proof that $\Pi^t_Q \wedge NZ$ implies $\mathrm{WF}_v(DeQ)$ is similar.

3.3 The NonZeno Condition

The timed queue specification Π^t_Q asserts that a DeQ action must occur between δ_{snd} and Δ_{snd} seconds of when it becomes enabled. What if $\Delta_{snd} < \delta_{snd}$? If an input occurs, it eventually is put in the queue, enabling DeQ. At that point, the value of *now* can never become more than Δ_{snd} greater than its current value, so the program eventually reaches a "time-blocked state". In a time-blocked state, only the *QTick* action can be enabled, and it cannot advance *now* past some fixed time. In other words, eventually a state is reached in which every variable other than *now* remains the same, and *now* either remains the same or keeps advancing closer and closer to some upper bound.

We can attempt to correct such pathological specifications by requiring that *now* increase without bound. This is easily done by conjoining the liveness property NZ to the safety property Π^t_Q, but that doesn't accomplish anything. Since $\Pi^t_Q \wedge NZ$ rules out behaviors in which *now* is bounded, it allows only behaviors in which there is no input, if $\Delta_{snd} < \delta_{snd}$. Such a specification is no better than the original specification Π^t_Q. The fact that the safety property allows the possibility of reaching a time-blocked state indicates an error in the specification. One does not add timing constraints on output actions with the intention of forbidding input.

We call a safety property *Zeno* if it allows the system to reach a state from which *now* must remain bounded. More precisely, a safety property Π is *nonZeno* iff every finite behavior satisfying Π can be completed to an infinite behavior satisfying Π in which *now* increases without bound. In other words, Π is nonZeno iff the pair (Π, NZ) is machine closed.[6]

Zeno specifications can be a source of incompleteness for proof methods. Only nonZeno behaviors are physically meaningful, so a real-time system with specifi-

[6]An arbitrary property Π is nonZeno iff $(\mathcal{C}(\Pi), \Pi \wedge NZ)$ is machine closed. We restrict our attention to real-time constraints for safety specifications.

cation Π satisfies a property Ψ if $\Pi \wedge NZ \Rightarrow \Psi$. Most methods for proving safety properties use only safety properties as hypotheses, so they can prove $\Pi \wedge NZ \Rightarrow \Psi$ for safety properties Π and Ψ only by proving $\Pi \Rightarrow \Psi$. NonZenoness of Π means that $\Pi \wedge NZ \Rightarrow \Psi$ holds iff $\Pi \Rightarrow \Psi$ does. However, if Π is Zeno, then $\Pi \wedge NZ \Rightarrow \Psi$ could hold even though $\Pi \Rightarrow \Psi$ does not, and these methods will be unable to prove that the system with specification Π satisfies Ψ. NonZenoness is therefore required for completeness.

The following result can be used to ensure that a real-time specification written in terms of volatile δ-timers is nonZeno.

Theorem 1 *Let v be the tuple of variables free in Init or $\mathcal{N}$. The property*

$$\begin{array}{l} \wedge\ Init \wedge \Box[\mathcal{N}]_v \wedge RT_v \\ \wedge\ \forall i \in I : VTimer(t_i, \mathcal{A}_i, \delta_i, v)\ \wedge\ MinTime(t_i, \mathcal{A}_i, v) \\ \wedge\ \forall j \in J : VTimer(T_j, \mathcal{A}_j, \Delta_j, v)\ \wedge\ MaxTime(T_j) \end{array}$$

is nonZeno if now does not appear in v, I and J are finite sets, and for all $i \in I$ and $j \in J$:

1. *$\langle \mathcal{A}_j \rangle_v$ is a subaction of $Init \wedge \Box[\mathcal{N}]_v$ whose free variables appear in v,*
2. *$\langle \mathcal{A}_i \rangle_v$ and $\langle \mathcal{A}_j \rangle_v$ are disjoint for $Init \wedge \Box[\mathcal{N}]_v$ if $i \neq j$,*
3. *δ_i and Δ_j are positive reals and, if $i = j$, then $\delta_i \leq \Delta_j$,*
4. *the t_i and T_j are distinct variables different from now and from the variables in v.*

We can apply the theorem to prove that the specification Π^t_Q is nonZeno if $\delta_{snd} \leq \Delta_{snd}$. The hypotheses of the theorem are checked as follows.

1. Actions $\langle DeQ \rangle_v$ and $\langle EnQ \rangle_v$ imply $\mathcal{N}_Q$, so they are subactions of Π_Q.
2. The conjunction of any two of the actions $\langle Inp \rangle_v$, $\langle DeQ \rangle_v$, and $\langle EnQ \rangle_v$ equals false, so the actions are pairwise disjoint for Π_Q.[7]
3. The hypothesis $\delta_{snd} \leq \Delta_{snd}$ is used here.

[7] Actually, the type-correctness predicate T_Q is needed to prove that $\langle Inp \rangle_v \wedge \langle DeQ \rangle_v$ equals false.

4. Trivial.

The theorem is valid for persistent as well as volatile timers. Any combination of *VTimer* and *PTimer* formulas may occur, except that a single $\mathcal{A}_k$ cannot have a persistent lower-bound timer t_k and a volatile upper-bound timer T_k. All of these results are corollaries of the following theorem, which in turn is a consequence of Theorem 4 of Section 4.

Theorem 2 *Let*

- Π *be a safety property of the form* $Init \wedge \Box[\mathcal{N}]_w$,
- t_i *and* T_j *be timers for* Π *and let* $\mathcal{A}_k$ *be an action, for all* $i \in I$, $j \in J$, *and* $k \in I \cup J$, *where* I *and* J *are sets, with* J *finite,*
- $$\Pi^t \;\stackrel{\Delta}{=}\; \Pi \wedge RT_v \wedge \forall\, i \in I : MinTime(t_i, \mathcal{A}_i, v) \;\wedge\; \forall\, j \in J : MaxTime(T_j)$$

If 1. $\langle \mathcal{A}_i \rangle_v$ *and* $\langle \mathcal{A}_j \rangle_v$ *are disjoint for* Π, *for all* $i \in I$ *and* $j \in J$ *with* $i \neq j$,

now

2. (a) *does not occur free in* v,
   (b) $(now' = r) \wedge (v' = v)$ *is a subaction of* Π, *for all* $r \in \mathbf{R}$,
3. *for all* $j \in J$:
   (a) $\langle \mathcal{A}_j \rangle_v \wedge (now' = now)$ *is a subaction of* Π.
   (b) $\Pi \;\Rightarrow\; VTimer(T_j, \mathcal{A}_j, \Delta_j, v)$, *or* $\Pi \;\Rightarrow\; PTimer(T_j, \mathcal{A}_j, \Delta_j, v)$, *where* $\Delta_j \in (0, \infty)$,
   (c) $\Pi^t \;\Rightarrow\; \Box(Enabled\ \langle \mathcal{A}_j \rangle_v = Enabled\ (\langle \mathcal{A}_j \rangle_v \wedge (now' = now)))$
   (d) $(v' = v) \;\Rightarrow\; (Enabled\ \langle \mathcal{A}_j \rangle_v = (Enabled\ \langle \mathcal{A}_j \rangle_v)')$
4. $\Pi^t \;\Rightarrow\; \Box(t_k \leq T_k)$, *for all* $k \in I \cap J$,

then (Π^t, NZ) *is machine closed*

Most nonaxiomatic approaches, including both real-time process algebras and more traditional programming languages with timing constraints, essentially use δ-timers for actions. Theorem 2 implies that they automatically yield nonZeno specifications.

Theorem 2 can be further generalized in two ways. First, J can be infinite—if Π^t implies that only a finite number of actions $\mathcal{A}_j$ with $j \in J$ are enabled before time r, for any $r \in \mathbf{R}$. For example, by letting $\mathcal{A}_j$ be the action that sends message number j, we can apply the theorem to a program that sends messages number 1 through n at time n, for every integer n. This program is nonZeno even though the number of actions per second that it performs is unbounded. Second, we can extend the theorem to the more general class of timers obtained by letting the Δ_j be arbitrary real-valued state functions, rather than just constants—if all the Δ_j are bounded from below by a positive constant Δ.

Theorem 2 can be proved using Propositions 1 and 3 and ordinary TLA reasoning. By these propositions, it suffices to display a formula L that is the conjunction of fairness conditions on subactions of Π^t such that $\Pi^t \wedge L$ implies NZ. A suitable L is $\mathrm{WF}_{\langle now, v\rangle}(\mathcal{C})$, where $\mathcal{C}$ is an action that either (a) advances *now* by $\min_{j \in J} \Delta_j$ if allowed by the upper-bound timers T_j, or else as far as they do allow, or (b) executes an $\langle \mathcal{A}_j \rangle_v$ action for which $now = T_j$. The proof in the appendix of Theorem 4, which implies Theorem 2, generalizes this approach.

Theorem 2 does not cover all situations of interest. For example, one can require of our timed queue that the first value appear on the output line within ϵ seconds of when it is placed on the input line. This effectively places an upper bound on the sum of the times needed for performing the *EnQ* and *DeQ* actions; it cannot be expressed with δ-timers on individual actions. For these general timing constraints, nonZenoness must be proved for the individual specification. The proof uses the method described above for proving Theorem 2: we add to the timed program Π^t a liveness property L that is the conjunction of any fairness properties we like, including fairness of the action that advances *now*, and prove that $\Pi^t \wedge L$ implies NZ. NonZenoness then follows from Propositions 1 and 3.

There is another possible approach to proving nonZenoness. One can make granularity assumptions—lower bounds both on the amount by which *now* is incremented and on the minimum delay for each action. Under these assumptions, nonZenoness is equivalent to the absence of deadlock, which can be proved by existing methods. Granularity assumptions are probably adequate—after all, what harm can come from pretending that nothing happens in less than 10^{-100} nanoseconds? However, they can be unnatural and cumbersome. For example, distributed algorithms often assume that only message delays are significant, so the time required for local actions is ignored. The specification of such an algorithm should place no lower bound on the time required for a local action, but that would violate any granularity assumptions. We believe that any proof of deadlock freedom based on granularity can be translated into a proof of nonZenoness using the method outlined above.

So far, we have been discussing nonZenoness of the internal specification, where both the timers and the system's internal variables are visible. Timers are defined by adding history-determined variables, so existentially quantifying over them preserves nonZenoness by Proposition 2. We expect most specifications to be fin [2, page 263], so nonZenoness will also be preserved by existentially quantifying over the system's internal variables. This is the case for the timed queue.

3.4 An Example: Fischer's Protocol

As another example of real-time closed systems, we treat a simplified version of a real-time mutual exclusion protocol proposed by Fischer [9], [12, page 2]. The example was suggested by Schneider [18]. The protocol consists of each process i executing the following code, where angle brackets denote instantaneous atomic actions:

$$\begin{array}{rl} a\text{:} & \textbf{await}\ \langle x = 0 \rangle; \\ b\text{:} & \langle x := i \rangle; \\ c\text{:} & \textbf{await}\ \langle x = i \rangle; \\ cs\text{:} & \text{critical section} \end{array}$$

There is a maximum delay Δ_b between the execution of the test in statement a and the assignment in statement b, and a minimum delay δ_c between the assignment in statement b and the test in statement c. The problem is to prove that, with suitable conditions on Δ_b and δ_c, this protocol guarantees mutual exclusion (at most one process can enter its critical section).

As written, Fischer's protocol permits only one process to enter its critical section one time. The protocol can be converted to an actual mutual exclusion algorithm. The correctness proof of the protocol is easily extended to a proof of such an algorithm.

The TLA specification of the protocol is given in Figure 4. The formula Π_F describing the untimed version is standard TLA. We assume a finite set Proc of processes. Variable x represents the program variable x, and variable pc represents the control state. The value of pc will be an array indexed by Proc, where $pc[i]$ equals one of the strings "a", "b", "c", "cs" when control in process i is at the corresponding statement. The initial predicate $Init_F$ asserts that $pc[i]$ equals "a" for each process i, so the processes start with control at statement a. No assumption on the initial value of x is needed to prove mutual exclusion.

Next come the definitions of the three actions corresponding to program statements

$$
\begin{array}{lcl}
Init_F & \triangleq & \forall\, i \in \mathsf{Proc} : pc[i] = \text{"a"} \\[4pt]
Go(i, u, v) & \triangleq & \land\ pc[i] = u \\
 & & \land\ pc'[i] = v \\
 & & \land\ \forall\, j \in \mathsf{Proc} : (j \neq i) \Rightarrow (pc'[j] = pc[j]) \\[4pt]
\mathcal{A}_i & \triangleq & Go(i, \text{"a"}, \text{"b"}) \land (x = x' = 0) \\[4pt]
\mathcal{B}_i & \triangleq & Go(i, \text{"b"}, \text{"c"}) \land (x' = i) \\[4pt]
\mathcal{C}_i & \triangleq & Go(i, \text{"c"}, \text{"cs"}) \land (x = x' = i) \\[4pt]
\mathcal{N}_F & \triangleq & \exists\, i \in \mathsf{Proc} : (\mathcal{A}_i \lor \mathcal{B}_i \lor \mathcal{C}_i) \\[4pt]
\Pi_F & \triangleq & Init_F \land \Box[\mathcal{N}_F]_{(x,pc)} \\[4pt]
\Pi_F^t & \triangleq & \land\ \Pi_F \land RT_{(x,pc)} \\
 & & \land\ \forall\, i \in \mathsf{Proc} : \land\ VTimer(T_b[i],\ \mathcal{B}_i,\ \Delta_b,\ (x, pc)) \\
 & & \qquad\qquad\qquad\quad \land\ MaxTime(T_b[i]) \\
 & & \land\ \forall\, i \in \mathsf{Proc} : \land\ VTimer(t_c[i],\ Go(i, \text{"c"}, \text{"cs"}),\ \delta_c,\ (x, pc)) \\
 & & \qquad\qquad\qquad\quad \land\ MinTime(t_c[i],\ \mathcal{C}_i,\ (x, pc)) \\[4pt]
\Phi_F^t & \triangleq & \exists\, T_b, t_c : \Pi_F^t
\end{array}
$$

Figure 4: The TLA specification of Fischer's real-time mutual exclusion protocol.

a, b, and c. They are defined using the formula Go, where $Go(i, u, v)$ asserts that control in process i changes from u to v, while control remains unchanged in the other processes. Action $\mathcal{A}_i$ represents the execution of statement a by process i; actions $\mathcal{B}_i$ and $\mathcal{C}_i$ have the analogous interpretation. In this simple protocol, a process stops when it gets to its critical section, so there are no other actions. The program's next-state action $\mathcal{N}_F$ is the disjunction of all these actions. Formula Π_F asserts that all processes start at statement a, and every step consists of executing the next statement of some process.

Action $\mathcal{B}_i$ is enabled by the execution of action $\mathcal{A}_i$. Therefore, the maximum delay of Δ_b between the execution of statements a and b can be expressed by an upper-bound constraint on a volatile Δ_b-timer for action $\mathcal{B}_i$. The variable T_b is an array of such timers, where $T_b[i]$ is the timer for action $\mathcal{B}_i$.

The constant δ_c is the minimum delay between when control reaches statement c and when that statement is executed. Therefore, we need an array t_c of lower-bound timers for the actions $\mathcal{C}_i$. The delay is measured from the time control reaches statement c, so we want $t_c[i]$ to be a δ_c-timer on an action that becomes enabled when process i reaches statement c and is not executed until $\mathcal{C}_i$ is. A suitable choice for this action is $Go(i, \text{“c”}, \text{“cs”})$.

Adding these timers and timing constraints to the untimed formula Π_F yields formula Π^t_F of Figure 4, the specification of the real-time protocol with the timers visible. The final specification, Φ^t_F, is obtained by quantifying over the timer variables T_b and t_c. Since $\mathcal{B}_j$ is a subaction of Π_F and $pc[i] = \text{“c”}$ is disjoint from $\mathcal{B}_j$, for all i and j in Proc, Theorem 2 implies that Π^t_F is nonZeno if Δ_b is positive. Proposition 2 can then be applied to prove that Φ^t_F is nonZeno.

Mutual exclusion asserts that two processes cannot be in their critical sections at the same time. It is expressed by the predicate

$$Mutex \ \stackrel{\Delta}{=}\ \forall\, i, j \in \mathsf{Proc} : (pc[i] = pc[j] = \text{“cs”}) \Rightarrow (i = j)$$

The property to be proved is

$$Assump \wedge \Phi^t_F \ \Rightarrow\ \Box Mutex \tag{8}$$

where $Assump$ expresses the assumptions about the constants Proc, Δ_b, and δ_c needed for correctness. Since the timer variables do not occur in $Mutex$ or $Assump$, (8) is equivalent to

$$Assump \wedge \Pi^t_F \ \Rightarrow\ \Box Mutex$$

The standard method for proving this kind of invariance property leads to the invariant

$$\begin{array}{l}
\wedge\ now \in \mathbf{R} \\
\wedge\ \forall\, i \in \mathsf{Proc} : \\
\qquad \wedge\ T_b[i], t_c[i] \ \in\ \mathbf{R} \cup \{\infty\} \\
\qquad \wedge\ pc[i] \in \{\text{“a”}, \text{“b”}, \text{“c”}, \text{“cs”}\} \\
\qquad \wedge\ (pc[i] = \text{“cs”}) \ \Rightarrow\ \wedge\ x = i \\
\qquad \qquad \qquad \qquad \qquad \ \ \wedge\ \forall\, j \in \mathsf{Proc} : pc[j] \neq \text{“b”} \\
\qquad \wedge\ (pc[i] = \text{“c”}) \ \Rightarrow\ \wedge\ x \neq 0 \\
\qquad \qquad \qquad \qquad \quad \ \wedge\ \forall\, j \in \mathsf{Proc} : (pc[j] = \text{“b”}) \Rightarrow (t_c[i] > T_b[j]) \\
\qquad \wedge\ (pc[i] = \text{“b”}) \ \Rightarrow\ (T_b[i] < now + \delta_c) \\
\qquad \wedge\ now \leq T_b[i]
\end{array}$$

and the assumption

$$Assump \ \stackrel{\Delta}{=}\ (0 \notin \mathsf{Proc}) \wedge (\Delta_b, \delta_c \in \mathbf{R}) \wedge (\Delta_b < \delta_c)$$

4 Open Systems

A closed system is solipsistic. An open system interacts with an environment, where system steps are distinguished from environment steps. Sections 4.1 and 4.2 reformulate a number of concepts introduced in [1] that are needed for treating open systems in TLA. Some new results appear in Section 4.3. The following two sections explain how reasoning about open systems is reduced to reasoning about closed systems, and how open systems are composed.

4.1 Receptiveness and Realizability

To describe an open system in TLA, one defines an action μ such that μ steps are attributed to the system and $\neg\mu$ steps are attributed to the environment. A specification should constrain only system steps, not environment steps.

For safety properties, the concept of constraining is formalized as follows: if μ is an action and Π a safety property, then Π *constrains at most* μ iff, for any finite behavior $s_1, \ldots, s_n$ and state s_{n+1}, if $s_1, \ldots, s_n$ satisfies Π and (s_n, s_{n+1}) is a $\neg\mu$ step, then $s_1, \ldots, s_{n+1}$ satisfies Π. The generalization to arbitrary properties of constraining at most μ is μ*-receptiveness*. Intuitively, Π is μ-receptive iff every behavior in Π can be achieved by an implementation that performs only μ steps—the environment being able to perform any $\neg\mu$ step. The concept of receptiveness is due to Dill [8]. The generalization to μ-receptiveness is developed in [1].[8] A safety property is μ-receptive iff it constrains at most μ.

The generalization of machine closure to open systems is *machine realizability*. Intuitively, (Π, L) is μ-machine realizable iff an implementation that performs only μ steps can ensure that any finite behavior satisfying Π is completed to an infinite behavior satisfying $\Pi \wedge L$. Formally, (Π, L) is defined to be μ-machine realizable iff (Π, L) is machine closed and $\Pi \wedge L$ is μ-receptive. For μ equal to true, machine realizability reduces to machine closure.

[8]To translate from the semantic model of [1] into that of TLA, we let agents be pairs of states and identify an action μ with the set of all agents that are μ steps. A TLA behavior $s_1, s_2, \ldots$ corresponds to the sequence $s_1 \xrightarrow{\alpha_2} s_2 \xrightarrow{\alpha_3} s_3 \xrightarrow{\alpha_4} \ldots$, where α_i equals (s_{i-1}, s_i). With this translation, the definitions in [1] differ from the ones given here and in the appendix mainly by attributing the choice of initial state to the environment rather than to the system, requiring initial conditions to be assumptions about the environment rather than guarantees by the system.

4.2 The $\mathrel{-\!\!\triangleright}$ Operator

A common way of specifying an open system is in terms of assumptions and guarantees [10], requiring the system to guarantee a property M if its environment satisfies an assumption E. An obvious formalization of such a specification is the property $E \Rightarrow M$. However, this property contains behaviors in which the system violates M and then the environment later violates E. Because the system cannot predict what the environment will do, such behaviors cannot occur in any actual implementation. A behavior σ generated by any implementation satisfies the additional property that if any finite prefix of σ satisfies E, then it satisfies M. We can therefore formalize the assumption/guarantee specification by the property $E \mathrel{-\!\!\triangleright} M$, defined by: $\sigma \in E \mathrel{-\!\!\triangleright} M$ iff $\sigma \in (E \Rightarrow M)$ and, for every finite prefix ρ of σ, if ρ satisfies E then ρ satisfies M. If E and M are safety properties, then $E \mathrel{-\!\!\triangleright} M$ is as well.

For safety properties, the operator $\mathrel{-\!\!\triangleright}$ is the implication operator of an intuitionistic logic [3]. Most valid propositional formulas without negation remain valid when $\Rightarrow$ is replaced by $\mathrel{-\!\!\triangleright}$, if all the formulas that appear on the left of a $\mathrel{-\!\!\triangleright}$ are safety properties. For example, the following formulas are valid if Φ and Π are safety properties.

$$\begin{aligned} \Phi \mathrel{-\!\!\triangleright} (\Pi \mathrel{-\!\!\triangleright} \Psi) \;&\equiv\; (\Phi \wedge \Pi) \mathrel{-\!\!\triangleright} \Psi \\ (\Phi \mathrel{-\!\!\triangleright} \Psi) \wedge (\Pi \mathrel{-\!\!\triangleright} \Psi) \;&\equiv\; (\Phi \vee \Pi) \mathrel{-\!\!\triangleright} \Psi \end{aligned} \tag{9}$$

For any TLA formulas Φ and Π, the property $\Phi \mathrel{-\!\!\triangleright} \Pi$ is expressible as a TLA formula.

4.3 Proving Machine Realizability

Propositions 1–3, which concern machine closure, have generalizations for machine realizability. Proposition 1 is the special case of Proposition 4 in which Φ and μ are identically true. Proposition 3 is similarly a special case of Proposition 5 if (true, L_2) is machine closed—that is, if L_2 is a liveness property. This is sufficient for our purposes, since NZ is a liveness property. The generalization of Proposition 2 is omitted; it would be analogous to Proposition 10 of [1].

Proposition 4 is stated in terms of μ-invariance, which generalizes the ordinary concept of invariance. A predicate P is a μ-invariant of a formula Π iff, in any behavior satisfying Π, no μ-step makes P false. This condition is expressed by the TLA formula $\Pi \Rightarrow \Box[(\mu \wedge P) \Rightarrow P']_P$.

Proposition 4 *If* Π *and* Φ *are safety properties,* Π *constrains at most* μ*, and* L *is the conjunction of a finite or countably infinite number of formulas of the form* $\mathrm{WF}_w(\mathcal{A})$ *and/or* $\mathrm{SF}_w(\mathcal{A})$*, where, for each such formula,*

1. $\langle \mathcal{A} \rangle_w$ *is a subaction of* $\Pi \wedge \Phi$,
2. $\Pi \wedge \Phi \Rightarrow \Box[\langle \mathcal{A} \rangle_w \Rightarrow \mu]_w$,
3. *if* $\mathcal{A}$ *appears in a formula* $\mathrm{SF}_w(\mathcal{A})$*, then* $\mathit{Enabled}\ \langle \mathcal{A} \rangle_w$ *is a* $\neg\mu$*-invariant of* $\Pi \wedge \Phi$,

then $(\Phi \mathrel{-\!\!\triangleright} \Pi,\ \Phi \Rightarrow L)$ *is* μ*-machine realizable.*

Proposition 5 *If* Φ *and* Π *are safety properties,* $(\Phi \mathrel{-\!\!\triangleright} \Pi, L_1)$ *and* (true, L_2) *are* μ*-machine realizable, and* $\Phi \wedge \Pi \wedge L_1$ *implies* L_2*, then* $(\Phi \mathrel{-\!\!\triangleright} \Pi, L_2)$ *is* μ*-machine realizable.*

4.4 Reduction to Closed Systems

Consider a specification $E \mathrel{-\!\!\triangleright} M$, where E and M are safety properties. We expect the system's requirement to restrict only system steps, meaning that M constrains at most μ. This implies that $E \mathrel{-\!\!\triangleright} M$ also constrains at most μ. We also expect the environment assumption E not to constrain system steps; formally, *E does not constrain μ* iff it constrains at most $\neg\mu$ and it is satisfied by every (finite behavior consisting only of an) initial state.[9]

Suppose E and M have the following form:

$$\begin{aligned} E &\triangleq \Box[\mu \vee \mathcal{N}_E]_v \\ M &\triangleq \mathit{Init} \wedge \Box[\neg\mu \vee \mathcal{N}_M]_v \end{aligned}$$

Then E does not constrain μ and M constrains at most μ. If the system's next-state action $\mathcal{N}_M$ implies μ, and the environment's next-state action $\mathcal{N}_E$ implies $\neg\mu$, then a simple calculation shows that

$$E \wedge M \equiv \mathit{Init} \wedge \Box[\mathcal{N}_E \vee \mathcal{N}_M]_v \tag{10}$$

Conjunction represents parallel composition, so $E \wedge M$ is the formula describing the closed system consisting of the open system together with its environment. Observe

[9]The asymmetry between *constrains at most* and *does not constrain* arises because we assign the system responsibility for the initial state.

that $E \wedge M$ has precisely the form we expect for a closed system comprising two components with next-state actions $\mathcal{N}_E$ and $\mathcal{N}_M$.

We can make the inverse transformation from a closed system specification Π to the corresponding assumption/guarantee specification $E \twoheadrightarrow M$ such that Π equals $E \wedge M$, where E does not constrain μ and M constrains at most μ. This is possible because any safety property Π can be written as such a conjunction.

Implementation means implication. A system with guarantee M implements a system with guarantee $\widehat{M}$, under environment assumption E, iff $E \twoheadrightarrow M$ implies $E \twoheadrightarrow \widehat{M}$. When E and M are safety properties, $E \twoheadrightarrow M$ implies $E \twoheadrightarrow \widehat{M}$ iff $E \wedge M$ implies $E \wedge \widehat{M}$. Thus, proving that one open system implements another is equivalent to proving the implementation relation for the corresponding closed systems. Implementation for open systems therefore reduces to implementation for closed systems.

4.5 Composition

The distinguishing feature of open systems is that they can be composed. The proof that the composition of two specifications implements a third specification is based on the following result, which is a reformulation of Theorem 2 of [1] for safety properties.

Theorem 3 *If E, E_1, E_2,, M_1, and M_2 are safety properties and μ_1 and μ_2 are actions such that*

1. *E_1 does not constrain μ_1 and E_2 does not constrain μ_2,*
2. *M_1 constrains at most μ_1 and M_2 constrains at most μ_2,*

then the following proof rule is valid:

$$\frac{E \wedge M_1 \wedge M_2 \;\Rightarrow\; E_1 \wedge E_2}{(E_1 \twoheadrightarrow M_1) \wedge (E_2 \twoheadrightarrow M_2) \;\Rightarrow\; (E \twoheadrightarrow M_1 \wedge M_2)}$$

This theorem is essentially the same as Theorem 1 of [3]; the proof is omitted.

5 Real-Time Open Systems

In Section 3, we saw how we can represent time by the variable *now* and introduce timing constraints with timers. To extend the method to open systems, we need only decide how to separate timing properties into environment assumptions and system guarantees. An examination of a paradoxical example in Section 5.1 leads to the general form described in Section 5.2, where the concept of nonZenoness is generalized.

5.1 A Paradox

Consider the two components Π_1 and Π_2 of Figure 5. Let the specification of Π_1 be $P_y \mathrel{-\!\!\triangleright} P_x$, which asserts that it writes a "good" sequence of outputs on x if its environment writes a good sequence of inputs on y. Let $P_x \mathrel{-\!\!\triangleright} P_y$ be the specification of Π_2, so Π_2 writes a good sequence of outputs on y if its environment writes a good sequence of inputs on x. If P_x and P_y are safety properties, then it appears that we should be able to apply Theorem 3, our composition principle, to deduce that the composite system Π_{12} satisfies $P_x \wedge P_y$, producing good sequences of values on x and y. (We can define μ_1 and μ_2 so that writing on x is a μ_1 action and writing on y is a μ_2 action.)

Now, suppose P_x and P_y both assert that the value 0 is written by noon. These can be regarded as safety properties, since they assert that an undesirable event never occurs—namely, noon passing without a 0 having been written. Hence, the composition principle apparently asserts that Π_{12} sends 0's along both x and y by noon. However, the specifications of Π_1 and Π_2 are satisfied by systems that wait for a 0 to be input, whereupon they immediately output a 0. The composition of those two systems does nothing.

This paradox depends on the ability of a system to respond instantaneously to an input. It is tempting to rule out such systems—perhaps even to outlaw specifications like these. We show that this Draconian measure is unnecessary. Indeed, if the specification of Π_2 is strengthened to assert that a 0 must unconditionally be written on y by noon, then there is no paradox, and the composition does guarantee that a 0 is written on both x and y by noon. All paradoxes disappear when one carefully examines how the specifications must be written.

To resolve the paradox, we examine more closely the specifications S_1 and S_2 of Π_1 and Π_2. For simplicity, let the only possible output actions be the setting of x and y to 0. The untimed version of S_1 then asserts that, if the environment does nothing but set y to 0, then the system does nothing but set x to 0. This is expressed

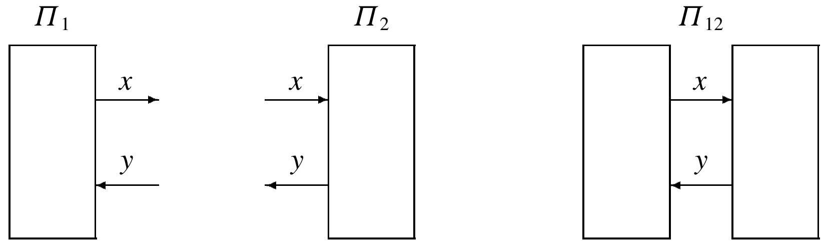

Figure 5: The composition of two systems.

in TLA by letting

$$\begin{aligned} \mathcal{M}_x &\triangleq (x' = 0) \land (y' = y) \qquad & \nu_1 &\triangleq x' \neq x \\ \mathcal{M}_y &\triangleq (y' = 0) \land (x' = x) \end{aligned}$$

and defining the untimed version of specification S_1 to be

$$\Box[\nu_1 \lor \mathcal{M}_y]_{(x,y)} \mathrel{-\!\!\triangleright} \Box[\neg\nu_1 \lor \mathcal{M}_x]_{(x,y)} \tag{11}$$

To add timing constraints, we must first decide whether the system or the environment should change *now*. Since the advancing of *now* is a mythical action that does not have to be performed by any device, either decision is possible. Somewhat surprisingly, it turns out to be more convenient to let the system advance time. With the convention that initial conditions appear in the system guarantee, we define:

$$\begin{aligned} \mathcal{N}_x &\triangleq \mathcal{M}_x \land (\mathit{now}' = \mathit{now}) \qquad & \mathit{MT}_x &\triangleq \mathit{MaxTime}(T_x) \\ \mathcal{N}_y &\triangleq \mathcal{M}_y \land (\mathit{now}' = \mathit{now}) & \mathit{MT}_y &\triangleq \mathit{MaxTime}(T_y) \\ T_x &\triangleq \textbf{if } x \neq 0 \textbf{ then } 12 \textbf{ else } \infty & \mu_1 &\triangleq \nu_1 \lor (\mathit{now}' \neq \mathit{now}) \\ T_y &\triangleq \textbf{if } y \neq 0 \textbf{ then } 12 \textbf{ else } \infty \end{aligned}$$

$$\begin{aligned} E_1 &\triangleq \Box[\mu_1 \lor \mathcal{N}_y]_{(x,y,\mathit{now})} \\ M_1 &\triangleq (\mathit{now} = 0) \land \Box[\neg\mu_1 \lor \mathcal{N}_x]_{(x,y,\mathit{now})} \land \mathit{RT}_{(x,y)} \land \mathit{MT}_x \end{aligned}$$

Adding timing constraints to (11) the same way we did for closed systems then leads to the following timed version of specification S_1.

$$(E_1 \land \mathit{MT}_y) \mathrel{-\!\!\triangleright} M_1 \tag{12}$$

However, this does not have the right form for an open system specification because MT_y constrains the advance of *now*, so the environment assumption constrains μ_1. The conjunct MT_y must be moved from the environment assumption to the system guarantee. Using (9), we rewrite (12) as:

$$S_1 \;\triangleq\; E_1 \mathrel{-\!\!\triangleright} (MT_y \mathrel{-\!\!\triangleright} M_1)$$

This has the expected form for an open system specification, with an environment assumption E_1 that does not constrain μ_1 and a system guarantee $MT_y \mathrel{-\!\!\triangleright} M_1$ that constrains at most μ_1.

The specification S_2 of the second component in Figure 5 is similar, where μ_2, E_2, M_2, and S_2 are obtained from μ_1, E_1, M_1, and S_1 by substituting 2 for 1, x for y, and y for x.

We now compose specifications S_1 and S_2. The definitions of M_1 and M_2 and the observation that $P \mathrel{-\!\!\triangleright} Q$ implies $P \Rightarrow Q$ yield

$$(MT_x \vee MT_y) \wedge (MT_y \mathrel{-\!\!\triangleright} M_1) \wedge (MT_x \mathrel{-\!\!\triangleright} M_2) \;\Rightarrow\; M_1 \wedge M_2 \tag{13}$$

The definitions of M_1 and M_2 and simple temporal reasoning yield

$$E \wedge M_1 \wedge M_2 \;\Rightarrow\; E_1 \wedge E_2 \tag{14}$$

where

$$E \;\triangleq\; \Box[\mu_1 \vee \mu_2]_{\langle x, y, now\rangle}$$

Combining (13) and (14) proves

$$E \wedge (MT_x \vee MT_y) \wedge (MT_y \mathrel{-\!\!\triangleright} M_1) \wedge (MT_x \mathrel{-\!\!\triangleright} M_2) \;\Rightarrow\; E_1 \wedge E_2$$

We can therefore apply Theorem 3, substituting $E \wedge (MT_x \vee MT_y)$ for E, $MT_y \mathrel{-\!\!\triangleright} M_1$ for M_1, and $MT_x \mathrel{-\!\!\triangleright} M_2$ for M_2, to deduce

$$S_1 \wedge S_2 \;\Rightarrow\; (E \wedge (MT_x \vee MT_y) \mathrel{-\!\!\triangleright} (MT_y \mathrel{-\!\!\triangleright} M_1) \;\wedge\; (MT_x \mathrel{-\!\!\triangleright} M_2))$$

Using the implication-like properties of $\mathrel{-\!\!\triangleright}$, this simplifies to

$$S_1 \;\wedge\; S_2 \;\Rightarrow\; (E \mathrel{-\!\!\triangleright} (MT_y \mathrel{-\!\!\triangleright} M_1) \;\wedge\; (MT_x \mathrel{-\!\!\triangleright} M_2)) \tag{15}$$

All one can conclude about the composition from (15) is: either x and y are both 0 when *now* reaches 12, or neither of them is 0 when *now* reaches 12. There is no paradox.

As another example, we replace S_2 by the specification $E_2 \mathrel{-\!\!\triangleright} M_2$. This specification, which we call S_3, asserts that the system sets y to 0 by noon, regardless of whether the environment sets x to 0. The definitions imply

$$MT_y \wedge E \wedge (MT_y \mathrel{-\!\!\triangleright} M_1) \wedge M_2 \;\Rightarrow\; E_1 \wedge E_2$$

and Theorem 3 yields

$$S_1 \;\wedge\; S_3 \quad\Rightarrow\quad (E \mathrel{-\!\!\triangleright} (MT_x \mathrel{-\!\!\triangleright} M_1) \;\wedge\; M_2)$$

Since M_2 implies MT_x, this simplifies to

$$S_1 \;\wedge\; S_3 \quad\Rightarrow\quad (E \mathrel{-\!\!\triangleright} M_1 \;\wedge\; M_2)$$

The composition of S_1 and S_3 does guarantee that both x and y equal 0 by noon.

5.2 Timing Constraints in General

Our no-longer-paradoxical example suggests that the form of a real-time open system specification should be

$$E \mathrel{-\!\!\triangleright} (P \mathrel{-\!\!\triangleright} M) \tag{16}$$

where M describes the system's timing constraints and the advancing of *now*, and P describes the upper-bound timing constraints for the environment. Since the environment's lower-bound timing constraints do not constrain the advance of *now*, they can remain in E. As we observed in Section 4.4, proving that one open specification implements another reduces to the proof for the corresponding closed systems. Since $E \mathrel{-\!\!\triangleright} (P \mathrel{-\!\!\triangleright} M)$ is equivalent to $(E \wedge P) \mathrel{-\!\!\triangleright} M$, the closed system corresponding to (16) is the expected one, $E \wedge P \wedge M$.

For the specification (16) to be reasonable, its closed system version, $E \wedge P \wedge M$, should be nonZeno. However, this is not sufficient. Consider a specification guaranteeing that the system produces a sequence of outputs until the environment sends a *stop* message, where the n^{th} output must occur by time $(n-1)/n$. There is no timing assumption on the environment; it need never send a *stop* message. This is an unreasonable specification because *now* cannot reach 1 until the environment sends its *stop* message, so the advance of time is contingent on an optional action of the environment. However, the corresponding closed system specification is nonZeno, since time can always be made to advance without bound by having the environment send a *stop* message.

If advancing *now* is a μ action, then a system that controls μ actions can guarantee time to be unbounded while satisfying a safety specification S iff the pair (S, NZ) is μ-machine realizable. We therefore take this condition to be the definition of nonZenoness for an open system specification S.

For specifications in terms of δ-timers, nonZenoness can be proved with generalizations to open systems of the theorems in Section 3.3. The following is the generalization of the strongest of them, Theorem 2. It is applied to a specification of the form (16) by substituting $E \wedge P$ for E.

Theorem 4 *With the notation and hypotheses of Theorem 2, if E and M are safety properties such that $\Pi = E \wedge M$, and*

5. *M constrains at most μ,*
6. *(a) $\langle \mathcal{A}_k \rangle_v \Rightarrow \mu$, for all $k \in I \cup J$,*
   (b) $(now' \neq now) \Rightarrow \mu$

then $(E \mathrel{-\!\!\triangleright} M^t,\ NZ)$ is μ-machine realizable, where

$$\begin{aligned} M^t \;\triangleq\;& M \;\wedge\; RT_v \;\wedge \\ & \quad \forall\, i \in I : MinTime(t_i, \mathcal{A}_i, v) \;\wedge\; \forall\, j \in J : MaxTime(T_j) \end{aligned}$$

The proof of Theorem 4, which appears in the appendix, is similar to the proof of Theorem 2 sketched in Section 3.3. It uses Propositions 4 and 5 instead of Propositions 1 and 3. Since machine realizability implies machine closure, Theorem 2 follows from Theorem 4 by letting E and μ equal true and M equal Π.

Theorem 4 applies to the internal specifications, where all variables are visible. For closed systems, existential quantification is handled with Proposition 2. For open systems, the generalization of this proposition—the analog of Proposition 10 of [1]—is needed.

6 Conclusion

6.1 What We Did

We started with a simple idea—specifying and reasoning about real-time systems by representing time as an ordinary variable. This idea led to an exposition that most readers probably found quite difficult. What happened to the simplicity?

About half of the exposition is a review of concepts unrelated to real time. All the fundamental concepts described in Sections 2 and 4, including machine closure, machine realizability, and the $\mathrel{-\!\!\triangleright}$ operator, have appeared before [1, 2]. These concepts are subtle, but they are important for understanding any concurrent system; they were not invented for real-time systems.

We chose to formulate these concepts in TLA. Like any language, TLA seems complicated on first encounter. We believe that a true measure of simplicity of a formal language is the simplicity of its formal description. The complete syntax and formal semantics of TLA are given in about 1-1/2 pages of figures in [13].

We never claimed that specifying and reasoning about concurrent systems is easy. Verifying concurrent systems is difficult and error prone. Our assertions that one formula follows from another, made so casually in the exposition, must be backed up by detailed calculations. We have omitted the proofs for our examples, which, done with the same detail as the proofs in the appendix, occupy some twenty pages.

We did claim that existing methods for specifying and reasoning about concurrent systems could be applied to real-time systems. Now, we can examine how hard they were to apply.

We found few obstacles in the realm of closed systems. The second author has more than fifteen years of experience in the formal verification of concurrent algorithms, and we knew that old-fashioned methods could be applied to real-time systems. However, TLA is relatively new, and we were pleased by how well it worked. The formal specification of Fischer's protocol in Figure 4, obtained by conjoining timing constraints to the untimed protocol, is as simple and direct as we could have hoped for. Moreover, the formal correctness proofs of this protocol and of the queue example, using the method of reasoning described in [13], were straightforward. Perhaps the most profound discovery was the relation between nonZenoness and machine closure.

Open systems made up for any lack of difficulty with closed systems. State-based approaches to open systems were a fairly recent development, and we had little experience with them. Studying real-time systems taught us a great deal, and led to a number of changes from the approach in [1]. For example, we now write specifications with $\mathrel{-\!\!\triangleright}$ instead of $\Rightarrow$, and we put initial conditions in the system guarantee rather than in the environment assumption. Many alternative ways of writing real-time specifications seemed plausible; choosing one that works was surprisingly hard. Even the simple idea of putting the environment's timing assumptions to the left of a $\mathrel{-\!\!\triangleright}$ in the system's guarantee came only after numerous failed efforts. Although the basic ideas we need to handle real-time open systems seem to be in place, we still have much to learn before reasoning about open systems becomes routine.

6.2 Beyond Real Time

Real-time systems introduce a fundamentally new problem: adding physical continuity to discrete systems. Our solution is based on the observation that, when reasoning about a discrete system, we can represent continuous processes by discrete actions. If we can pretend that the system progresses by discrete atomic actions, we can pretend that those actions occur at a single instant of time, and that the continuous change to time also occurs in discrete steps. If there is no system action between noon and $\sqrt{2}$ seconds past noon, we can pretend that time advances by those $\sqrt{2}$ seconds in a single action.

Physical continuity arises not just in real-time systems, but in "real-pressure" and "real-temperature" process-control systems. Such systems can be described in the same way as real-time systems: pressure and temperature as well as time are represented by ordinary variables. The continuous changes to pressure and temperature that occur between system actions are represented by discrete changes to the variables. The fundamental assumption is that the real, physical system is accurately represented by a model in which the system makes discrete, instantaneous changes to the physical parameters it affects.

The observation that continuous parameters other than time can be modeled by program variables has probably been known for years. However, the first published work we know of that uses this idea, by Marzullo, Schneider, and Budhiraja [14], appeared only recently.

Appendix

A Definitions

A.1 States and Behaviors

Σ: The set of all states (assignments of values to variables).
Σ^*: The set of finite behaviors (finite sequences of states).
Σ^∞: The set of behaviors (infinite sequences of states).
λ: The empty sequence.
$|\sigma|$ [for $\sigma \in \Sigma^* \cup \Sigma^\infty$]: The length of σ.
σ_n [for $\sigma \in \Sigma^* \cup \Sigma^\infty$ and $0 < n \leq |\sigma|$]: The n^{th} state in the sequence σ.
$\sigma|_n$ [for $\sigma \in \Sigma^* \cup \Sigma^\infty$ and $0 \leq n \leq |\sigma|$]: The finite behavior $\sigma_1, \ldots, \sigma_n$; equal to λ when $n = 0$.
$\sigma \circ s$ [for $\sigma \in \Sigma^*$ and $s \in \Sigma$]: The finite behavior $\sigma_1, \ldots, \sigma_{|\sigma|}, s$.
$\sigma \simeq \tau$ [for $\sigma, \tau \in \Sigma^*$ or $\sigma, \tau \in \Sigma^\infty$]: σ and τ are equal except for stuttering steps (repetitions of identical states).
$\sigma \simeq_x \tau$ [for x a tuple of variables and either $\sigma, \tau \in \Sigma^*$ or $\sigma, \tau \in \Sigma^\infty$]:
There exist behaviors $\widehat{\sigma}$ and $\widehat{\tau}$ such that
- $\wedge$ $\widehat{\sigma} \simeq \sigma$ and $\widehat{\tau} \simeq \tau$
- $\wedge$ $|\widehat{\sigma}| = |\widehat{\tau}|$
- $\wedge$ for every integer $n \leq |\widehat{\sigma}|$, the states $\widehat{\sigma}_n$ and $\widehat{\tau}_n$ are equal except for the values they assign to the variables of x.

A.2 Predicates and Actions

predicate: A subset of Σ.
state function: A mapping from Σ to the set of values.
action: A subset of $\Sigma \times \Sigma$.
transition function: A mapping from $\Sigma \times \Sigma$ to the set of values.
P' [for P a predicate]: The action $\{(s, t) : t \in P\}$.
f' [for f a state function]: The transition function such that $f'(s, t) = f(t)$.
$Enabled\ \mathcal{A}$ [for $\mathcal{A}$ an action]: The predicate $\{s \in \Sigma : (\exists t \in \sigma : (s, t) \in \mathcal{A})\}$.
$[\mathcal{A}]_f$ [for $\mathcal{A}$ an action and f a state function]: $\mathcal{A} \vee (f' = f)$
$\langle \mathcal{A} \rangle_f$ [for $\mathcal{A}$ an action and f a state function]: $\neg[\neg\mathcal{A}]_f$ (equals $\mathcal{A} \wedge (f' \neq f)$).

A.3 Properties

property: A subset Π of Σ^∞ such that for any σ, τ in Σ^∞, if $\sigma \simeq \tau$ then $\sigma \in \Pi$ iff $\tau \in \Pi$.

$\Box\Pi$ [for Π a property]:
The property $\{\sigma \in \sigma^\infty : (\forall n > 0 : \sigma_n, \sigma_{n+1}, \ldots \in \Pi)\}$.

$\Box[\mathcal{A}]_f$ [for $\mathcal{A}$ an action and f a state function]:
The property $\{\sigma \in \Sigma^\infty : (\forall n > 0 : (\sigma_n, \sigma_{n+1}) \in [\mathcal{A}]_f)\}$.

$\Diamond\Pi$ [for Π a property or an action of the form $\langle\mathcal{A}\rangle_f$]: $\neg\Box\neg\Pi$

$\Pi \leadsto \Phi$ [for Π and Φ properties]: $\Box(\Pi \Rightarrow \Diamond\Phi)$

$\mathrm{WF}_f(A)$ [for f a state function and $\mathcal{A}$ an action]:
The property $(\Box\Diamond\langle\mathcal{A}\rangle_f) \vee (\Box\Diamond\neg \mathit{Enabled}\ \langle\mathcal{A}\rangle_f)$.

$\mathrm{SF}_f(A)$ [for f a state function and $\mathcal{A}$ an action]:
The property $(\Box\Diamond\langle\mathcal{A}\rangle_f) \vee (\Diamond\Box\neg \mathit{Enabled}\ \langle\mathcal{A}\rangle_f)$.

$\exists x : \Pi$ [for x a tuple of variables and Π a property]:
The property $\{\sigma \in \Sigma^\infty : (\exists \rho \in \Pi : \rho \simeq_x \sigma)\}$.

A.4 Closure and Safety

$\sigma \models \Pi$ [for $\sigma \in \Sigma^*$ and Π a property]:
There exists $\tau \in \Pi$ and n such that $\sigma = \tau|_n$.

$\mathcal{C}(\Pi)$ [for Π a property]: The property $\{\sigma \in \Sigma^\infty : (\forall n \geq 0 : \sigma|_n \models \Pi\}$.

safety property: A property Π such that $\Pi = \mathcal{C}(\Pi)$.

(Π, L) *machine closed* [for Π and L properties]: $\Pi = \mathcal{C}(\Pi \wedge L)$

$\Phi \mathrel{-\!\!\triangleright} \Pi$ [for Φ and Π properties]: The property consisting of all behaviors σ such that
$\wedge$ $\sigma \in \Phi$ implies $\sigma \in \Pi$
$\wedge$ for all $n \geq 0$: $\sigma|_n \models \Phi$ implies $\sigma|_n \models \Pi$.

A.5 Realizability

Π *constrains at most* μ [for Π a safety property and μ an action]:
For any $s \in \Sigma$ and $\rho \in \Sigma^*$ with $|\rho| > 0$, if $\rho \models \Pi$ and $(\rho_{|\rho|}, s) \notin \mu$, then $\rho \circ s \models \Pi$.

Π *does not constrain* μ [for Π a safety property and μ an action]:
Π constrains at most $\neg\mu$, and every behavior of length 1 satisfies Π.

μ-*strategy* [for μ an action]: A mapping f from Σ^* to $\Sigma \cup \{\perp\}$ such that $f(\rho) \neq \perp$ implies $(\rho_{|\rho|}, f(\rho)) \in \mu$, for any sequence $\rho \in \Sigma^*$ with $\rho \neq \lambda$.

$\mathcal{O}_\mu(f)$ [for μ an action and f a μ-strategy]: The set of behaviors σ such that

$\wedge\ f(\lambda) = \sigma_1$
$\wedge$ for all $n > 0$: $(\sigma_n, \sigma_{n+1}) \in \mu$ implies $\sigma_{n+1} = f(\sigma|_n)$
$\wedge$ for infinitely many n: $f(\sigma|_n) = \perp$ or $(\sigma_n, \sigma_{n+1}) \in \mu$.

$\mathcal{O}_\mu(f, \rho)$ [for μ an action, f a μ-strategy, and $\rho \in \Sigma^*$]: If $\rho = \lambda$ then $\mathcal{O}_\mu(f)$, else the set of behaviors σ such that
$\wedge\ \tau = \sigma|_{|\tau|}$
$\wedge$ for all $n \geq |\tau|$: $(\sigma_n, \sigma_{n+1}) \in \mu$ implies $\sigma_{n+1} = f(\sigma|_n)$
$\wedge$ for infinitely many n: $f(\sigma|_n) = \perp$ or $(\sigma_n, \sigma_{n+1}) \in \mu$.

Π *is μ-receptive* [for Π a property and μ an action]: For every behavior $\sigma \in \Pi$ there exists a μ-strategy f such that $\sigma \in \mathcal{O}_\mu(f)$ and $\mathcal{O}_\mu(f) \subseteq \Pi$.

(Π, L) *is μ-machine realizable* [for Π and L properties and μ an action]:
(Π, L) is machine closed and $\Pi \wedge L$ is μ-receptive.

B Proofs

B.1 Proof Notation

We use hierarchically structured proofs. The theorem to be proved is statement $\langle 0\rangle 1$. The proof of statement $\langle i\rangle j$ is either an ordinary paragraph-style proof or the sequence of statements $\langle i+1\rangle 1$, $\langle i+1\rangle 2$, ... and their proofs. A proof may be preceded by a proof sketch, which informally describes the proof. Within a proof, $\langle k\rangle l$ denotes the most recent statement with that number. A statement has the form

ASSUME: *Assump* PROVE: *Goal*

which is abbreviated to *Goal* if there is no assumption. The assertion Q.E.D. in statement number $\langle i+1\rangle k$ of the proof of statement $\langle i\rangle j$ denotes the goal of statement $\langle i\rangle j$. The statement

CASE: *Assump*

is an abbreviation for

ASSUME: *Assump* PROVE: Q.E.D.

Within the proof of statement $\langle i\rangle j$, assumption $\langle i\rangle$ denotes that statement's assumption, and $\langle i\rangle.k$ denotes the assumption's k^{th} item.

We recommend that proofs be read hierarchically, from the top level down. To read the proof of a long level-k step: (i) read the level-$(k+1)$ statements that comprise its proof, together with the proof of the final Q.E.D. step (which is usually a short paragraph), (ii) read the proof of each level-$(k+1)$ step, in any desired order.

B.2 Proof of Proposition 2

ASSUME: 1. (Π, L) is machine-closed.
2. x a tuple of variables, none of which occurs free in L.
3. ρ is a finite behavior satisfying $\exists x : \Pi$.

PROVE: There exists an infinite behavior τ satisfying $(\exists x : \Pi) \wedge L$ such that ρ is a prefix of τ.

$\langle 1 \rangle 1$. Choose a finite behavior σ satisfying Π such that $\sigma \simeq_x \rho$.
  $\langle 2 \rangle 1$. Choose a behavior ϕ such that $\phi \in \exists x : \Pi$ and ρ a prefix of ϕ.
   PROOF: Assumption $\langle 0 \rangle.1$.
  $\langle 2 \rangle 2$. Choose a behavior ψ such that $\psi \in \Pi$ and $\psi \simeq_x \phi$.
   PROOF: $\langle 2 \rangle 1$ and the definition of $\exists x : \Pi$.
  $\langle 2 \rangle 3$. There exists an initial prefix σ of ψ such that $\sigma \simeq_x \rho$.
   PROOF: $\psi \simeq_x \phi$ (by $\langle 2 \rangle 2$) and ρ a prefix of ϕ (by $\langle 2 \rangle 1$).
  $\langle 2 \rangle 4$. Q.E.D.
   PROOF: $\langle 2 \rangle 2$ and $\langle 2 \rangle 3$.
$\langle 1 \rangle 2$. Choose η satisfying $\Pi \wedge L$ such that σ is a prefix of η.
  PROOF: The existence of τ follows from $\langle 1 \rangle 1$ and assumption $\langle 0 \rangle.1$.
$\langle 1 \rangle 3$. Choose τ such that $\tau \simeq_x \eta$ and ρ is a prefix of τ.
  PROOF: τ exists because σ is a prefix of η (by $\langle 1 \rangle 2$) and $\sigma \simeq_x \rho$ (by $\langle 1 \rangle 1$).
$\langle 1 \rangle 4$. τ satisfies $\exists x : \Pi$.
  PROOF: $\langle 1 \rangle 2$, $\langle 1 \rangle 3$, and the definition of $\exists x : \Pi$.
$\langle 1 \rangle 5$. τ satisfies L.
  PROOF: $\langle 1 \rangle 2$, $\langle 1 \rangle 3$, and assumption $\langle 0 \rangle.2$.
$\langle 1 \rangle 6$. Q.E.D.
  PROOF: Steps $\langle 1 \rangle 4$ and $\langle 1 \rangle 5$ imply that τ satisfies $(\exists x : \Pi) \wedge L$, and ρ is a prefix of τ by $\langle 1 \rangle 3$.

B.3 Proof of Proposition 3

ASSUME: 1. (Π, L_1) is machine closed.
2. $\Pi \wedge L_1$ implies L_2.

PROVE: (Π, L_2) is machine closed.

PROOF:
$$\begin{aligned} \Pi &= \mathcal{C}(\Pi \wedge L_1) && [\langle 0 \rangle.1] \\ &\subseteq \mathcal{C}(\Pi \wedge L_2) && [\langle 0 \rangle.2 \text{ and monotonicity of closure}] \\ &\subseteq \mathcal{C}(\Pi) && [\text{monotonicity of closure}] \\ &= \Pi && [\langle 0 \rangle.1, \text{ which implies } \Pi \text{ closed}] \end{aligned}$$

This proves that $\Pi = \mathcal{C}(\Pi \wedge L_2)$. □

B.4 Two Lemmas About Machine-Realizability

Lemma 1 *(Π, L) is μ-machine realizable iff for every finite behavior τ satisfying Π, there exists a μ-strategy f such that $\tau \models \mathcal{O}_\mu(f)$ and $\mathcal{O}_\mu(f) \subseteq \Pi \wedge L$.*

PROOF: This is proved in Proposition 9 of [1]. □

Lemma 2 *(Π, L) is μ-machine realizable iff for every finite behavior τ satisfying Π, there exists a μ-strategy f such that $\mathcal{O}_\mu(f, \tau) \subseteq \Pi \wedge L$.*

$\langle 1\rangle 1$. ASSUME: 1. For all finite τ such that $\tau \models \Pi$, there exists f_τ such that
a. f_τ is a μ-strategy.
b. $\mathcal{O}_\mu(f_\tau, \tau) \subseteq \Pi \wedge L$
2. ρ a finite behavior such that $\rho \models \Pi$.
PROVE: There exists a μ-strategy g such that
$\wedge\ \rho \models \mathcal{O}_\mu(g)$
$\wedge\ \mathcal{O}_\mu(g) \subseteq \Pi \wedge L$

$\langle 2\rangle 1$. Choose g such that, for all $\sigma \in \Sigma^*$:

$$g(\sigma) \ \triangleq\ \textbf{if}\ (\sigma = \rho|_n) \wedge (n < |\rho|)\ \textbf{then if}\ (n = 0) \vee (\rho_n, \rho_{n+1}) \in \mu\ \textbf{then}\ \rho_{n+1}\ \textbf{else}\ \perp\ \textbf{else}\ f_\eta(\sigma),\ \text{where } \eta \text{ is the longest common prefix of } \rho \text{ and } \sigma$$

PROOF: The requisite f_η exists by assumption $\langle 0\rangle.2$, since $\rho \models \Pi$ implies $\eta \models \Pi$, for any prefix η of ρ.

$\langle 2\rangle 2$. g is a μ-strategy.
PROOF: $\langle 0\rangle.1$a and the the definition of g ($\langle 2\rangle 1$).

$\langle 2\rangle 3$. $\mathcal{O}_\mu(g) \subseteq \Pi \wedge L$
PROOF: Assumption $\langle 0\rangle.1$b, since $\kappa \in \mathcal{O}_\mu(g)$ implies that $\kappa \in \mathcal{O}_\mu(f_\eta, \eta)$, where η is the longest common prefix of κ and ρ.

$\langle 2\rangle 4$. $\rho \models \mathcal{O}_\mu(g)$
PROOF: The definition of g implies that ρ is a prefix of an element of $\mathcal{O}_\mu(g)$.

$\langle 2\rangle 5$. Q.E.D.
PROOF: $\langle 2\rangle 2$, $\langle 2\rangle 3$, and $\langle 2\rangle 4$.

$\langle 1\rangle 2$. ASSUME: (Π, L) is μ-machine realizable and $\tau \models \Pi$.
PROVE: There exists a μ-strategy f such that $\mathcal{O}_\mu(f, \tau) \subseteq \Pi \wedge L$.

$\langle 2\rangle 1$. Choose a μ-strategy f such that $\tau \in \mathcal{O}_\mu(f)$ and $\mathcal{O}_\mu(f) \subseteq \Pi \wedge L$.
PROOF: Assumption $\langle 2\rangle$ and Lemma 1.

⟨2⟩2. $\mathcal{O}_\mu(f, \tau) \subseteq \Pi \wedge L$

PROOF: ⟨2⟩1.2 and the definition of $\mathcal{O}_\mu(f, \tau)$.

⟨2⟩3. Q.E.D.

PROOF: ⟨2⟩1 and ⟨2⟩2.

⟨1⟩3. Q.E.D.

PROOF: By ⟨1⟩1, ⟨1⟩2 and Lemma 1.

B.5 Proof of Proposition 4

LET: $\Psi \triangleq \Phi \wedge \Pi$

$L \triangleq \wedge\ \forall i \in I : \mathrm{WF}_{w_i}(\mathcal{A}_i)$
$\quad\ \ \wedge\ \forall j \in J : \mathrm{SF}_{w_j}(\mathcal{A}_j)$

ASSUME: 1. Φ and Π are safety properties.
2. Π constrains at most μ.
3. $I \cup J$ a finite or countably infinite set, which we take to be a set of natural numbers.
4. For all k in $I \cup J$:
a. $\langle \mathcal{A}_k \rangle_{w_k}$ is a subaction of Ψ.
b. $\Psi \Rightarrow \Box[\langle \mathcal{A}_k \rangle_{w_k} \Rightarrow \mu]_{w_k}$
5. Forall j in J: $Enabled\ \langle \mathcal{A}_j \rangle_{w_j}$ is a $\neg\mu$ invariant of Ψ.

PROVE: $(\Phi \mathrel{-\!\!\triangleright} \Pi, \Phi \Rightarrow L)$ is μ-machine realizable.

By Lemma 2, it suffices to

ASSUME: 6. $\tau \models \Phi \mathrel{-\!\!\triangleright} \Pi$

PROVE: There exists a μ-strategy f such that

$$\mathcal{O}_\mu(f, \tau) \subseteq (\Phi \mathrel{-\!\!\triangleright} \Pi) \wedge (\Phi \Rightarrow L).$$

PROOF SKETCH: Defining the strategy f is tantamount to constructing a scheduler that executes the actions $\langle \mathcal{A}_k \rangle_{w_k}$ to satisfy the fairness requirements. Action $\langle \mathcal{A}_k \rangle_{w_k}$ is never considered until k steps have occurred, so only finitely many actions are considered at any point. The next action scheduled is the enabled action that has been under consideration for the longest time without having been executed. In the event of ties, the lowest-numbered action is chosen. The strategy f will never violate Ψ, and it stops making moves if Ψ is ever made false by the environment. In the following definitions, *nextact*(ρ) is the k such that the scheduler chooses to execute $\langle \mathcal{A}_k \rangle_{w_k}$ after the finite behavior ρ.

LET: $\mathcal{B}_k \triangleq \langle \mathcal{A}_k \rangle_{w_k}$

$n \triangleq |\tau|$

$lasttime(\rho, k) \triangleq \max(\{l : 1 < l \leq |\rho| : (\rho_{l-1}, \rho_l) \in \mathcal{B}_k\} \cup \{k\})$

$$
\begin{array}{lll}
\textit{able}(\rho) & \triangleq & \textbf{if } \rho = \lambda \textbf{ then } \text{the set of naturals} \\
& & \qquad\qquad\;\; \textbf{else } \{k : \rho_{|\rho|} \in \mathit{Enabled}\ \mathcal{B}_k\} \\
\textit{oldestlast}(\rho) & \triangleq & \min\{\textit{lasttime}(\rho, k) : k \in \textit{able}(\rho)\} \\
\textit{nextact}(\rho) & \triangleq & \min\{k \in \textit{able}(\rho) : \textit{lasttime}(\rho, k) = \textit{oldestlast}(\rho)\}
\end{array}
$$

$\langle 1 \rangle 1.$ Choose f such that for all ρ in Σ^*:

$$
\begin{array}{l}
\textbf{if } (\textit{able}(\rho) = \emptyset) \lor (\rho \notin \Psi) \\
\quad \textbf{then } f(\rho) = \bot \\
\quad \textbf{else } \begin{array}[t]{l} \land\ (\rho \neq \lambda) \Rightarrow (\rho_{|\rho|}, f(\rho)) \in \mathcal{B}_{\textit{nextact}(\rho)} \\ \land\ \rho \circ f(\rho) \in \Psi \end{array}
\end{array}
$$

PROOF SKETCH: We must show that, for any ρ, the requisite $f(\rho)$ exists.

$\langle 2 \rangle 1.$ ASSUME: $(\textit{able}(\rho) \neq \emptyset) \ \land\ \rho \models \Psi$

PROVE: $\exists s \in \Sigma : \begin{array}[t]{l} \land\ (\rho \neq \lambda) \ \Rightarrow\ (\rho_{|\rho|}, s) \in \mathcal{B}_{\textit{nextact}(\rho)} \\ \land\ \rho \circ s \in \Psi \end{array}$

$\langle 3 \rangle 1.$ CASE: $\rho = \lambda$

PROOF: Assumption $\langle 2 \rangle$ and the definition of $\rho \models \Psi$.

$\langle 4 \rangle 1.$ CASE: $\rho \neq \lambda$

$\langle 5 \rangle 1.$ $\rho_{|\rho|} \in \mathit{Enabled}\ \mathcal{B}_{\textit{nextact}(\rho)}$

$\langle 6 \rangle 1.$ $\textit{nextact}(\rho) \in \textit{able}(\rho)$

PROOF: Assumption $\langle 2 \rangle$ and the definitions of *nextact* and *oldestlast*.

$\langle 6 \rangle 2.$ Q.E.D.

PROOF: Case assumption $\langle 4 \rangle$ and the definition of $\textit{able}(\rho)$.

$\langle 5 \rangle 2.$ Q.E.D.

PROOF: $\langle 5 \rangle 1$ and Assumption $\langle 0 \rangle$.4a.

$\langle 4 \rangle 2.$ Q.E.D.

PROOF: $\langle 3 \rangle 1$ and $\langle 4 \rangle 1$.

$\langle 2 \rangle 2.$ Q.E.D.

PROOF: The existence of the required f follows easily from $\langle 2 \rangle 1$.

$\langle 1 \rangle 2.$ f is a μ-strategy.

ASSUME: $(\rho \neq \lambda) \ \land\ (f(\rho) \neq \bot)$

PROVE: $(\rho_{|\rho|}, f(\rho)) \in \mu$

PROOF: $\langle 1 \rangle 1$ and assumption $\langle 1 \rangle$ imply

$$
\begin{array}{l}
\land\ (\rho_{|\rho|}, f(\rho)) \in \mathcal{B}_{\textit{nextact}(\rho)} \\
\land\ \rho \circ f(\rho) \models \Psi
\end{array}
$$

and the result follows from $\langle 0 \rangle$.4b.

$\langle 1 \rangle 3.$ $\mathcal{O}_\mu(f, \tau) \subseteq (\Phi \mathrel{-\!\!\triangleright} \Pi)$

PROOF SKETCH: This is a straightforward induction proof; $\langle 2 \rangle 1$ is the base case and $\langle 2 \rangle 2$ is the induction step.

$\langle 2 \rangle 1.$ ASSUME: $\sigma \in \mathcal{O}_\mu(f, \tau)$

PROVE: $\sigma|_0 \models \Phi \rhd\!\!\!\!- \Pi$

PROOF: $\sigma|_0 = \tau|_0$, and $\tau \models (\Phi \rhd\!\!\!\!- \Pi)$ by assumption $\langle 2\rangle$ and the definition of $\mathcal{O}_\mu(f, \tau)$.

$\langle 2\rangle 2$. ASSUME: 1. $\sigma \in \mathcal{O}_\mu(f, \tau)$
2. $\sigma|_i \models \Phi \rhd\!\!\!\!- \Pi$
3. $i \geq 0$

PROVE: $\sigma|_{i+1} \models \Phi \rhd\!\!\!\!- \Pi$

$\langle 3\rangle 1$. CASE: $i \geq |\tau|$

$\langle 4\rangle 1$. CASE: $(\sigma_i, \sigma_{i+1}) \notin \mu$

PROOF: Follows easily from assumption $\langle 2\rangle$.2 and assumption $\langle 0\rangle$.2, since Π constrains at most μ implies $\Phi \rhd\!\!\!\!- \Pi$ constrains at most μ.

$\langle 4\rangle 2$. CASE: $(\sigma_i, \sigma_{i+1}) \in \mu$

$\langle 5\rangle 1$. $\sigma_{i+1} = f(\sigma|_i)$

PROOF: Assumption $\langle 2\rangle$.1.

$\langle 5\rangle 2$. Q.E.D.

PROOF: $\langle 5\rangle 1$, assumption $\langle 2\rangle$.2, and the **else** clause in $\langle 1\rangle 1$ (the definition of f).

$\langle 4\rangle 3$. Q.E.D.

PROOF: $\langle 4\rangle 1$ and $\langle 4\rangle 2$.

$\langle 3\rangle 2$. CASE: $i < |\tau|$

PROOF: By assumption $\langle 0\rangle$.6 and assumption $\langle 2\rangle$.1, since $i < |\tau|$ implies $\sigma|_{i+1} = \tau|_{i+1}$.

$\langle 3\rangle 3$. Q.E.D.

PROOF: $\langle 3\rangle 1$ and $\langle 3\rangle 2$.

$\langle 2\rangle 3$. Q.E.D.

PROOF: $\langle 2\rangle 1$, $\langle 2\rangle 2$, mathematical induction, and assumption $\langle 0\rangle$.1.

$\langle 1\rangle 4$. $\mathcal{O}_\mu(f, \tau) \subseteq (\Phi \Rightarrow L)$

PROOF SKETCH: The proof is by contradiction. We define *IE* and *IO* to be the sets of indices of actions that should be executed infinitely often and that actually are executed infinitely often, respectively. We show that the existence of a q in *IE* but not in *IO* leads to a contradiction.

ASSUME: 1. $\sigma \in \mathcal{O}_\mu(f, \tau)$
2. $\sigma \in \Phi$
3. $\exists k \in I \cup J$: a.$\vee$ $k \in I \wedge \sigma \notin \mathrm{WF}_{w_k}(\mathcal{A}_k)$
b.$\vee$ $k \in J \wedge \sigma \notin \mathrm{SF}_{w_k}(\mathcal{A}_k)$

PROVE: flase

LET: $IE \triangleq \{i \in I : \Diamond\Box \mathit{Enabled}\ \mathcal{B}_i\} \cup \{j \in J : \Box\Diamond \mathit{Enabled}\ \mathcal{B}_j\}$

$IO \triangleq \{k \in I \cup J : (\sigma_n, \sigma_{n+1}) \in \mathcal{B}_k$ for infinitely many $n\}$

$\langle 2\rangle 1.$ $\sigma \in \Phi \wedge \Pi$

PROOF: $\langle 1\rangle 3$ and assumption $\langle 1\rangle.1$, which imply $\sigma \in \Phi \mathrel{-\!\!\triangleright} \Pi$, and $\langle 1\rangle.2$.

$\langle 2\rangle 2.$ Choose q such that

1. $q \in IE$
2. $q \notin IO$

PROOF: Assumption $\langle 1\rangle.3$.

$\langle 2\rangle 3.$ Choose m_1 such that

1. $m_1 > q$
2. $\forall n \geq m_1 : (\sigma_{n-1}, \sigma_n) \notin \mathcal{B}_q$
3. $m_1 > |\tau|$

PROOF: $\langle 2\rangle 2.2$

$\langle 2\rangle 4.$ Choose m_2 such that

1. $m_2 > m_1$
2. $\forall k \in I \cup J : (k \notin IO) \wedge (k < m_1) \Rightarrow \forall n > m_2 : (\sigma_n, \sigma_{n+1}) \notin \mathcal{B}_k$

PROOF: Definition of IO.

$\langle 2\rangle 5.$ Choose m_3 such that

1. $m_3 > m_2$
2. $\forall k \in I \cup J : (k \in IO) \wedge (k < m_1) \Rightarrow$
   $\exists n > m_2 : (n < m_3) \wedge (\sigma_n, \sigma_{n+1}) \in \mathcal{B}_k$

PROOF: Definition of IO.

$\langle 2\rangle 6.$ Choose m_4 such that

1. $m_4 > m_3$
2. $\sigma_{m_4} \in \mathit{Enabled}\ \mathcal{B}_q$
3. $(q \in I) \Rightarrow \forall n > m_4 : \sigma_n \in \mathit{Enabled}\ \mathcal{B}_q$

PROOF: $\langle 2\rangle 2.1$ and the definition of IE.

$\langle 2\rangle 7.$ $\forall n > m_4 :$ 1.$\wedge$ $\sigma_n \in \mathit{Enabled}\ \mathcal{B}_q$
2.$\wedge$ $(\sigma_{n-1}, \sigma_n) \notin \mu$

PROOF SKETCH: By time m_4, $\mathcal{B}_q$ is the next action scheduled for execution. Since $\mathcal{B}_q$ is never again executed (by $\langle 2\rangle 3$), there can be no further μ steps. Since only a μ step can disable $\mathcal{B}_q$ (by assumption $\langle 0\rangle.5$), $\mathcal{B}_q$ must be enabled forever. The formal proof is by induction on n; $\langle 3\rangle 1$ is the base case and $\langle 3\rangle 2$ is the induction step.

$\langle 3\rangle 1.$ $\sigma_{m_4} \in \mathit{Enabled}\ \mathcal{B}_q$

PROOF: $\langle 2\rangle 6$.

$\langle 3\rangle 2.$ ASSUME: 1. $\sigma_n \in \mathit{Enabled}\ \mathcal{B}_q$
2. $(\sigma_{n-1}, \sigma_n) \notin \mu$
3. $n \geq m_4$

PROVE: $\wedge$ $\sigma_{n+1} \in \mathit{Enabled}\ \mathcal{B}_q$
$\wedge$ $(\sigma_n, \sigma_{n+1}) \notin \mu$

$\langle 4\rangle 1.$ $(\sigma_n, \sigma_{n+1}) \notin \mu$

ASSUME: $(\sigma_n, \sigma_{n+1}) \in \mu$

PROVE: flase

$\langle 5\rangle 1.$ $\lor\ nextact(\sigma|_n) = q$

$\lor\ \exists k \in I \cup J : \land\ nextact(\sigma|_n) = k$

$\land\ k < m_1 \ \land\ k \notin IO$

$\land\ k \in able(\sigma|_n)$

$\langle 6\rangle 1.$ $q \in able(\sigma|_n)$

Assumption $\langle 3\rangle .1$.

$\langle 6\rangle 2.$ ASSUME: $k \geq m_1$

PROVE: $lasttime(\sigma|_n, k) > lasttime(\sigma|_n, q)$

PROOF: $lasttime(\sigma|_n, k) \geq m_1$ by definition of *lasttime*, and $m_1 > lasttime(\sigma|_n, q)$ by assumption $\langle 3\rangle .1$.

$\langle 6\rangle 3.$ ASSUME: $(k > m_1) \land (k \in IO)$

PROVE: $lasttime(\sigma|_n, k) > lasttime(\sigma|_n, q)$

PROOF: $lasttime(\sigma|_n, k) \geq m_2 > m_1$ by $\langle 2\rangle 5$, and $m_1 > lasttime(\sigma|_n, q)$ by assumption $\langle 3\rangle .1$.

$\langle 6\rangle 4.$ $\lor\ oldestlast(\sigma|_n) = q$

$\lor\ \exists k \in I \cup J : \land\ oldestlast(\sigma|_n) = k$

$\land\ k < m_1 \ \land\ k \notin IO$

$\land\ k \in able(\sigma|_n)$

PROOF: $\langle 6\rangle 1$, $\langle 6\rangle 2$, and $\langle 6\rangle 3$.

$\langle 6\rangle 5.$ Q.E.D.

PROOF: $\langle 6\rangle 4$ and definition of *nextact*.

$\langle 5\rangle 2.$ $\lor\ (\sigma_n, \sigma_{n+1}) \in \mathcal{B}_q$

$\lor\ \exists k \in I \cup J : \land\ (\sigma_n, \sigma_{n+1}) \in \mathcal{B}_k$

$\land\ k < m_1 \ \land\ k \notin IO$

PROOF: $\langle 5\rangle 1$, the definition of f ($\langle 1\rangle 1$), assumption $\langle 4\rangle$, and assumption $\langle 1\rangle .1$, since $n > |\tau|$ by assumption $\langle 3\rangle .3$ and $\langle 2\rangle 3$–$\langle 2\rangle 6$.

$\langle 5\rangle 3.$ Q.E.D.

$\langle 6\rangle 1.$ CASE: $(\sigma_n, \sigma_{n+1}) \in \mathcal{B}_q$

PROOF: Contradicts $\langle 2\rangle 3$, since $n > m_1$ by assumption $\langle 3\rangle .3$ and $\langle 2\rangle 4$–$\langle 2\rangle 6$.

$\langle 6\rangle 2.$ CASE: $\exists k \in I \cup J : \land\ (\sigma_n, \sigma_{n+1}) \in \mathcal{B}_k$

$\land\ k < m_1 \ \land\ k \notin IO$

PROOF: Contradicts $\langle 2\rangle 4$, since $n > m_2$ by assumption $\langle 3\rangle .3$, $\langle 2\rangle 5$, and $\langle 2\rangle 6$.

$\langle 6\rangle 3.$ Q.E.D.

PROOF: $\langle 6\rangle 1$ and $\langle 6\rangle 2$.

$\langle 4\rangle 2.$ $\sigma_{n+1} \in \mathit{Enabled}\ \mathcal{B}_q$
PROOF: Assumption $\langle 3\rangle.1$, $\langle 4\rangle 1$, $\langle 2\rangle 1$, and assumption $\langle 0\rangle.5$.
$\langle 4\rangle 3.$ Q.E.D.
PROOF: $\langle 4\rangle 1$ and $\langle 4\rangle 2$.
$\langle 3\rangle 3.$ Q.E.D.
PROOF: $\langle 3\rangle 1$, $\langle 3\rangle 2$, and mathematical induction.
$\langle 2\rangle 8.$ Q.E.D.
$\langle 3\rangle 1.$ $\forall n > m_4 : f(\sigma|_n) \neq \bot$
PROOF: $\langle 2\rangle 7.1$, assumption $\langle 1\rangle.1$, assumption $\langle 1\rangle.2$, and $\langle 1\rangle 3$.
$\langle 3\rangle 2.$ Q.E.D.
PROOF: $\langle 3\rangle 1$ and $\langle 2\rangle 7.2$ contradict assumption $\langle 1\rangle.1$.
$\langle 1\rangle 5.$ Q.E.D.
PROOF: $\langle 1\rangle 2$, $\langle 1\rangle 3$, and $\langle 1\rangle 4$.

B.6 Proof of Proposition 5

ASSUME: 1. Φ and Π are safety properties.
2. $(\Phi \mathrel{-\!\!\triangleright} \Pi, L_1)$ is μ-machine realizable.
3. (true, L_2) is μ-machine realizable.
4. $\Phi \wedge \Pi \wedge L_1$ implies L_2.
PROVE: $(\Phi \mathrel{-\!\!\triangleright} \Pi, L_2)$ is μ-machine realizable.
$\langle 1\rangle 1.$ ASSUME: ρ is a finite behavior such that $\rho \models (\Phi \mathrel{-\!\!\triangleright} \Pi)$.
PROVE: There exists a μ-strategy f such that
$$\mathcal{O}_\mu(f, \rho) \subseteq (\Phi \mathrel{-\!\!\triangleright} \Pi) \wedge L_2.$$
$\langle 2\rangle 1.$ Choose μ-strategy h such that $\mathcal{O}_\mu(h, \rho) \subseteq (\Phi \mathrel{-\!\!\triangleright} \Pi) \wedge L_1$.
PROOF: Assumption $\langle 0\rangle.2$, assumption $\langle 1\rangle$, and Lemma 2.
$\langle 2\rangle 2.$ For all τ, choose a μ-strategy g_τ such that $\mathcal{O}_\mu(g_\tau, \tau) \subseteq L_2$.
PROOF: Assumption $\langle 0\rangle.3$ and Lemma 2.
$\langle 2\rangle 3.$ Let $f(\tau) = h(\tau)$ for $\tau \models \Phi$, and otherwise let $f(\tau) = g_\eta(\tau)$ where η is the shortest prefix of τ such that $\eta \not\models \Phi$. Then f is a μ-strategy.
PROOF: By steps $\langle 2\rangle 1$ and $\langle 2\rangle 2$ (which say that h and g_η are μ-strategies).
$\langle 2\rangle 4.$ $\mathcal{O}_\mu(f, \rho) \subseteq (\Phi \mathrel{-\!\!\triangleright} \Pi) \wedge L_2$
ASSUME: $\tau \in \mathcal{O}_\mu(f, \rho)$
PROVE: $\tau \in (\Phi \mathrel{-\!\!\triangleright} \Pi) \wedge L_2$
$\langle 3\rangle 1.$ CASE: $\tau \in \Phi$
$\langle 4\rangle 1.$ $\tau \in \mathcal{O}_\mu(h, \rho)$

PROOF: By assumption $\langle 2 \rangle$, since the case assumption $\langle 3 \rangle$ and the definition of f (step $\langle 2 \rangle 3$) imply $f(\tau|_n) = h(\tau|_n)$ for all n.

$\langle 4 \rangle 2.$ $\tau \in (\Phi \mathrel{-\!\!\triangleright} \Pi) \wedge L_1$

PROOF: By step $\langle 4 \rangle 1$ and the choice of h (step $\langle 2 \rangle 1$).

$\langle 4 \rangle 3.$ $\tau \in \Phi \wedge \Pi \wedge L_1$

PROOF: By step $\langle 4 \rangle 2$ and case assumption $\langle 3 \rangle$.

$\langle 4 \rangle 4.$ Q.E.D.

PROOF: By step $\langle 4 \rangle 3$ and assumption $\langle 0 \rangle .4$.

$\langle 3 \rangle 2.$ CASE: $\tau \notin \Phi$

$\langle 4 \rangle 1.$ Let n be the least integer such that $\tau|_n \notin \Phi$.

PROOF: Such an n exists by case assumption $\langle 3 \rangle$ and assumption $\langle 0 \rangle .1$.

$\langle 4 \rangle 2.$ $\tau \in \mathcal{O}_\mu(g_{\tau|_n}, \tau|_n)$

PROOF: By step $\langle 4 \rangle 1$, assumption $\langle 2 \rangle$, and the definition of f (step $\langle 2 \rangle 3$).

$\langle 4 \rangle 3.$ $\tau \in L_2$

PROOF: By steps $\langle 4 \rangle 2$ and the choice of $g_{\tau|_n}$ (step $\langle 2 \rangle 2$).

$\langle 4 \rangle 4.$ Q.E.D.

PROOF: By steps $\langle 4 \rangle 3$ and case assumption $\langle 3 \rangle$.

$\langle 3 \rangle 3.$ Q.E.D.

PROOF: By steps $\langle 3 \rangle 1$ and $\langle 3 \rangle 2$.

$\langle 2 \rangle 5.$ Q.E.D.

PROOF: By steps $\langle 2 \rangle 3$ and $\langle 2 \rangle 4$.

$\langle 1 \rangle 2.$ Q.E.D.

PROOF: The result follows immediately from step $\langle 1 \rangle 1$ and Lemma 2.

B.7 Proof of Theorem 4

The proof of Theorem 4 is rather long and is presented in two steps. Section B.7.1 contains a high-level view of the proof; Section B.7.2 contains the complete proof, omitting definitions and complete subproofs that appear in the high-level view. If a formula F is a conjunction, we may number the conjuncts and let $F.i$ denote the i^{th} one.

The proof uses the following two lemmas.

Lemma 3 (Rule WF1) *For any predicates P, Q, and I; actions $\mathcal{N}$ and $\mathcal{A}$; and state function f; if*

1. $P \land I \land I' \land [\mathcal{N}]_f \;\Rightarrow\; (P' \lor Q')$
2. $P \land I \land I' \land \langle \mathcal{N} \land \mathcal{A} \rangle_f \;\Rightarrow\; Q'$
3. $P \land I \;\Rightarrow\; \mathit{Enabled}\ \langle \mathcal{A} \rangle_f$.

then $\mathrm{WF}_f(\mathcal{A}) \land \Box I \land \Box[\mathcal{N}]_f \;\Rightarrow\; (P \leadsto Q)$

PROOF: This is a simple generalization of rule WF1 of [13]; the proof of soundness is omitted.

Lemma 4 *For any actions $\mathcal{A}$ and $\mathcal{B}$, state function f, variable x, predicate P, and property Π:*

1. (a) $(\mathit{Enabled}\ (x' = f)) = \mathsf{true}$.
   (b) *If $\mathcal{A}$ and $\mathcal{B}$ have no primed variables in common, then* $(\mathit{Enabled}\ \mathcal{A} \land \mathcal{B}) = (\mathit{Enabled}\ \mathcal{A}) \land (\mathit{Enabled}\ \mathcal{B})$.
2. $(\mathit{Enabled}\ \mathcal{A}) \land (\neg \mathit{Enabled}\ \mathcal{B}) \;\Rightarrow\; (\mathit{Enabled}\ \mathcal{A} \land \neg \mathcal{B})$
3. $(\mathit{Enabled}\ P \land \mathcal{A}) = (P \land \mathit{Enabled}\ \mathcal{A})$
4. $(\mathit{Enabled}\ \exists x : \mathcal{A}) = (\exists x : \mathit{Enabled}\ \mathcal{A})$
5. *If* $\mathcal{A} \Rightarrow \mathcal{B}$ *then* $(\mathit{Enabled}\ \mathcal{A}) \Rightarrow (\mathit{Enabled}\ \mathcal{B})$.

PROOF: These properties all follow easily from the definitions.

B.7.1 High-Level Proof of the Theorem

LET: $\Pi \triangleq E \land M$

$M^t \triangleq M \land RT_v \land \forall i \in I : \mathit{MinTime}(t_i, \mathcal{A}_i, v) \land \forall j \in J : \mathit{MaxTime}(T_j)$

$\Pi^t \triangleq E \land M^t$

ASSUME: 0. a. $\Pi = \mathit{Init} \land \Box[\mathcal{N}]_w$ for some predicate *Init*, action $\mathcal{N}$, and state function w.

b. t_i is a timer for Π, for all i in I.

c. T_j is a timer for Π, for all j in J.

d. J is a finite set.

1. $\langle \mathcal{A}_i \rangle_v$ and $\langle \mathcal{A}_j \rangle_v$ are disjoint for Π, for all $i \in I$ and $j \in J$ with $i \neq j$.

now

2. a. does not occur free in v.

b. $(\mathit{now}' = r) \land (v' = v)$ is a subaction of Π, for all $r \in \mathbf{R}$.

3. For all $j \in J$:

a. $\langle \mathcal{A}_j \rangle_v \land (\mathit{now}' = \mathit{now})$ is a subaction of Π.

b. $\Pi \;\Rightarrow\; \mathit{VTimer}(T_j, \mathcal{A}_j, \Delta_j, v)$ or

$\Pi \;\Rightarrow\; \mathit{PTimer}(T_j, \mathcal{A}_j, \Delta_j, v)$, where $\Delta_j \in (0, \infty)$.

c. $\Pi^t \;\Rightarrow\; \Box(\mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v = \mathit{Enabled}\ (\langle \mathcal{A}_j \rangle_v \,\wedge\, (\mathit{now}' = \mathit{now})))$

d. $(v' = v) \;\Rightarrow\; (\mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v = (\mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v)')$

4. $\Pi^t \;\Rightarrow\; \Box(t_k \leq T_k)$, for all $k \in I \cap J$.
5. E and M are safety properties, and M constrains at most μ.
6. a. $\langle \mathcal{A}_k \rangle_v \Rightarrow \mu$, for all $k \in I \cup J$.

   b. $(\mathit{now}' \neq \mathit{now}) \Rightarrow \mu$

PROVE: $(E \mathrel{-\!\!\triangleright} M^t,\ \mathit{NZ})$ is μ-machine realizable.

PROOF SKETCH: We show that a fairness condition $\mathrm{WF}_{(\mathit{now},v)}(\mathcal{C})$ implies that *now* increases without bound, for a subaction $\langle \mathcal{C} \rangle_{(\mathit{now},v)}$ of Π^t, and then apply Propositions 4 and 5. To prove that *now* increases without bound, we prove that $\mathit{now} = r$ leads to $\mathit{now} \geq (r + \Delta)$ for any number r, where Δ is the minimum of the Δ_j. The action $\mathcal{C}$ advances *now* or, if that is impossible because $\mathit{now} = T_j$ for some upper-bound timer T_j, it performs the action $\mathcal{A}_j$ (thereby advancing T_j).

We begin by naming the next-state relations of the RT_v, *MaxTime*, *MinTime*, *VTimer*, and *PTimer* formulas and defining some actions and predicates, including $\mathcal{C}$. A $\mathcal{B}_j$ step is defined to be an $\mathcal{A}_j$ step that leaves *now* unchanged. An $\mathcal{A}^t_j$ or $\mathcal{B}^t_j$ step is an $\mathcal{A}_j$ or $\mathcal{B}_j$ step that satisfies the lower-bound timing constraint, if there is one. The state function T is the smallest T_j that restricts the advance of *now*.

LET:
$$\mathit{RTact}_v \;\triangleq\; [(\mathit{now}' \in (\mathit{now}, \infty)) \wedge (v' = v)]_{\mathit{now}}$$
$$\mathit{MaxTact}(t) \;\triangleq\; [\mathit{now}' \leq t']_{\mathit{now}}$$
$$\mathit{MinTact}(t, \mathcal{A}, v) \;\triangleq\; [\mathcal{A} \Rightarrow (t \leq \mathit{now})]_v$$
$$\begin{aligned}\mathit{VTact}(t, \mathcal{A}, \delta, v) \;\triangleq\; [\wedge\ t' \;=\; &\textbf{if}\ (\mathit{Enabled}\ \langle \mathcal{A} \rangle_v)' \\ &\textbf{then if}\ \langle \mathcal{A} \rangle_v \vee \neg \mathit{Enabled}\ \langle \mathcal{A} \rangle_v \\ &\qquad\textbf{then}\ \mathit{now} + \delta \\ &\qquad\textbf{else}\ \ t \\ &\textbf{else}\ \ \infty \\ \wedge\ v' \neq v&]_{(t,v)}\end{aligned}$$
$$\begin{aligned}\mathit{PTact}(t, \mathcal{A}, \delta, v) \;\triangleq\; [\wedge\ t' \;=\; &\textbf{if}\ \mathit{Enabled}\ \langle \mathcal{A} \rangle_v \\ &\textbf{then if}\ \langle \mathcal{A} \rangle_v\ \textbf{then}\ \mathit{now} + \delta \\ &\qquad\qquad\quad\ \textbf{else}\ \ t \\ &\textbf{else}\ \ t + (\mathit{now}' - \mathit{now}) \\ \wedge\ (\mathit{now}, v)' \neq (\mathit{now}&, v)]_{(t,v,\mathit{now})}\end{aligned}$$
$$P_j \;\triangleq\; (j \in I) \Rightarrow (t_j \leq \mathit{now})$$
$$\mathcal{A}^t_j \;\triangleq\; \mathcal{A}_j \wedge P_j$$
$$\mathcal{B}_j \;\triangleq\; \mathcal{A}_j \wedge (\mathit{now}' = \mathit{now})$$
$$\mathcal{B}^t_j \;\triangleq\; \mathcal{A}^t_j \wedge (\mathit{now}' = \mathit{now})$$

$$
\begin{array}{lcl}
T & \triangleq & \min\{T_j : j \in J \wedge \mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v\} \\
\Delta & \triangleq & \min\{\Delta_j : j \in J\} \\
\mathit{NowT} & \triangleq & (\mathit{now}' = \min(\mathit{now} + \Delta, T)) \wedge (v' = v) \\
\mathcal{C} & \triangleq & \vee\ (T \neq \mathit{now})\ \wedge\ \mathit{NowT} \\
 & & \vee\ (T = \mathit{now})\ \wedge\ \exists\, j \in J : (T_j = T) \wedge \mathcal{B}_j^t
\end{array}
$$

$\langle 1 \rangle 1.$ Choose J_P and J_V such that:

1. $J = J_P \cup J_V$
2. $J_P \cap J_V = \emptyset$
3. $\forall j \in J_V : \Pi \Rightarrow \mathit{VTimer}(T_j, \mathcal{A}_j, \Delta_j, v)$
4. $\forall j \in J_P : \Pi \Rightarrow \mathit{PTimer}(T_j, \mathcal{A}_j, \Delta_j, v)$

PROOF: J_P and J_V exist by assumption $\langle 0 \rangle.3b$.

$\langle 1 \rangle 2.$ $\Pi^t \Rightarrow \Box \mathit{Inv}$, where

$$
\begin{array}{lcl}
\mathit{Inv} & \triangleq & 1.\wedge\ \forall j \in J : T_j \in [\mathit{now}, \infty] \\
 & & 2.\wedge\ \mathit{now} \in \mathbf{R} \\
 & & 3.\wedge\ \forall j \in J : (\mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v = \\
 & & \qquad\qquad \mathit{Enabled}\ (\langle \mathcal{A}_j \rangle_v \wedge (\mathit{now}' = \mathit{now})) \\
 & & 4.\wedge\ \forall k \in I \cap J : t_k \leq T_k \\
 & & 5.\wedge\ \forall j \in J_V : \neg \mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v \Rightarrow (T_j = \infty)
\end{array}
$$

PROOF SKETCH: It follows immediately from the hypotheses that Π^t implies $\Box \mathit{Inv}.3$ and $\Box \mathit{Inv}.4$. The proofs for the other conjuncts of Inv are straightforward but somewhat tedious invariance proofs, which are omitted.

$\langle 2 \rangle 1.$ $\mathit{RT}_v \Rightarrow \Box \mathit{Inv}.2$

$\langle 2 \rangle 2.$ For all $j \in J$: $\Pi \wedge \Box \mathit{Inv}.2 \wedge \mathit{MaxTime}(T_j) \Rightarrow \Box \mathit{Inv}.1$

$\langle 2 \rangle 3.$ $\Pi^t \Rightarrow \Box \mathit{Inv}.3$

$\langle 2 \rangle 4.$ $\Pi^t \Rightarrow \Box \mathit{Inv}.4$

$\langle 2 \rangle 5.$ $\Pi^t \Rightarrow \Box \mathit{Inv}.5$

$\langle 2 \rangle 6.$ Q.E.D.

PROOF: $\langle 2 \rangle 1$–$\langle 2 \rangle 5$.

$\langle 1 \rangle 3.$ 1. $\Pi = \mathit{Init} \wedge \Box \mathcal{M}$

2. $\Pi^t = \mathit{Init}^t \wedge \Box \mathcal{N}^t$, for some predicate Init^t.

where

$$
\begin{array}{lcl}
\mathcal{M} & \triangleq & \wedge\ [\mathcal{N}]_w \\
 & & \wedge\ \forall j \in J_P : \mathit{PTact}(T_j, \mathcal{A}_j, \Delta_j, v) \\
 & & \wedge\ \forall j \in J_V : \mathit{VTact}(T_j, \mathcal{A}_j, \Delta_j, v) \\
\mathcal{N}^t & \triangleq & \wedge\ \mathcal{M} \\
 & & \wedge\ \mathit{RTact}_v \\
 & & \wedge\ \forall i \in I : \mathit{MinTact}(t_i, \mathcal{A}_i, v) \\
 & & \wedge\ \forall j \in J : \mathit{MaxTact}(T_j)
\end{array}
$$

PROOF: 1 follows from assumptions $\langle 0\rangle$.0a and $\langle 0\rangle$.3b, and 2 follows from the definition of Π^t, since $\Box$ distributes over $\wedge$. (A simple calculation shows that $\mathcal{M} = [\mathcal{M}]_f$ and $\mathcal{N}^t = [\mathcal{N}^t]_g$, for suitable tuples f and g, so $\Box\mathcal{M}$ and $\Box\mathcal{N}^t$ are TLA formulas.)

$\langle 1\rangle 4.$ $\langle \mathcal{C}\rangle_{(now,v)}$ is a subaction of Π^t.

$\langle 2\rangle 1.$ For all j in J : $\langle \mathcal{B}_j^t\rangle_v$ is a subaction of Π^t.

$\langle 2\rangle 2.$ $\langle NowT\rangle_{now}$ is a subaction of Π^t.

$\langle 2\rangle 3.$ Q.E.D.

PROOF: By $\langle 2\rangle 1$ and $\langle 2\rangle 2$, since $\langle NowT\rangle_{now} = \langle NowT\rangle_{(now,v)}$, $\langle \mathcal{B}_j^t\rangle_v = \langle \mathcal{B}_j^t\rangle_{(now,v)}$, Lemma 4.4 implies that the disjunction of subactions is a subaction, and Lemma 4.3 implies that if $\mathcal{D}$ is a subaction, then $P \wedge \mathcal{D}$ is also a subaction, for any predicate P.

$\langle 1\rangle 5.$ $\Pi^t \wedge \mathrm{WF}_{(now,v)}(\mathcal{C}) \Rightarrow NZ$

PROOF: By the Lattice Rule [13], it suffices to

ASSUME: $r \in \mathbf{R}$

PROVE: $\Pi^t \wedge \mathrm{WF}_{(now,v)}(\mathcal{C}) \;\Rightarrow\; (\,(now = r) \leadsto (now \in [r+\Delta, \infty))\,)$

PROOF SKETCH: The standard method of proving that $now = r$ leads to $now \in [r+\Delta, \infty)$ is to assume that $now = r$ and now is never in $[r+\Delta, \infty)$, and derive a contradiction. Step $\langle 2\rangle 2$ below proves that, if now equals r and it is never in $[r+\Delta, \infty)$, then it is always in $[r, r+\Delta)$. It therefore suffices to assume now is always in $[r, r+\Delta)$ and derive the contradiction.

The contradiction is obtained by showing that eventually there is no upper-bound timer preventing the advance of now past $r+\Delta$. The timers that could prevent the advance of now are the ones in the set U of timers that are less than $r+\Delta$. Step $\langle 2\rangle 4$ asserts that U eventually becomes empty, and $\langle 2\rangle 5$ asserts that time then advances past $r+\Delta$.

LET: $U \triangleq \{j \in J : T_j < r+\Delta\}$

$V \triangleq \{j \in J : now < T_j < r+\Delta\}$

$TimerAct \triangleq \forall j \in J : \vee\, VTact(T_j, \mathcal{A}_j, \Delta_j, v)$
$\qquad\qquad\qquad\qquad\;\; \vee\, PTact(T_j, \mathcal{A}_j, \Delta_j, v)$

$\langle 2\rangle 1.$ $Inv \Rightarrow \mathit{Enabled}\ \langle \mathcal{C}\rangle_{(now,v)}$

$\langle 2\rangle 2.$ $\Pi^t \;\Rightarrow$
$$\Box((now = r) \wedge \Box(now \in (-\infty, r+\Delta)) \;\Rightarrow\; \Box(now \in [r, r+\Delta)))$$

$\langle 2\rangle 3.$ ASSUME: $j \in J$

PROVE: 1. $\Pi^t \wedge \Box(now \in [r, r+\Delta)) \;\Rightarrow\; \Box((j \notin U) \Rightarrow \Box(j \notin U))$

2. $\Pi^t \wedge \Box(now \in [r, r+\Delta)) \;\Rightarrow\; \Box((j \notin V) \Rightarrow \Box(j \notin V))$

$\langle 2\rangle 4.$ $\Pi^t \wedge \Box(now \in [r, r+\Delta)) \wedge \mathrm{WF}_{(now,v)}(\mathcal{C}) \;\Rightarrow\; \Diamond\Box(U = \emptyset)$

⟨2⟩5. $\wedge \Box(now \in [r, r + \Delta))$
$\wedge \Box(U = \emptyset)$
$\wedge \Box Inv$
$\wedge \text{WF}_{(now,v)}(\mathcal{C})$
$\Rightarrow \mathsf{true} \leadsto (now \geq r + \Delta)$

⟨2⟩6. Q.E.D.

PROOF: ⟨2⟩2, ⟨2⟩4, ⟨2⟩5, ⟨1⟩2, and temporal logic.

⟨1⟩6. $(E \mathrel{-\!\!\triangleright} M^t,\ E \Rightarrow \text{WF}_{(now,v)}(\mathcal{C}))$ is μ-machine realizable.

⟨2⟩1. M^t constrains at most μ.

⟨2⟩2. Q.E.D.

PROOF: We apply Proposition 4, with E substituted for Φ, M^t substituted for Π, and the single formula $\text{WF}_{(now,v)}(\mathcal{C})$. Step ⟨2⟩1 asserts that M^t constrains at most μ. The three numbered hypotheses of the proposition are proved as follows:

1. ⟨1⟩4.
2. Assumptions ⟨0⟩.6a and ⟨0⟩.6b and the definition of $\mathcal{C}$.
3. Vacuous.

⟨1⟩7. (true, NZ) is μ-machine realizable.

PROOF: Assumption ⟨0⟩.6b.

⟨1⟩8. Q.E.D.

PROOF: Proposition 5, using ⟨0⟩.5, ⟨1⟩6, ⟨1⟩7, and ⟨1⟩5.

B.7.2 Detailed Proof of the Theorem

⟨1⟩1. Choose J_P and J_V such that:

1. $J = J_P \cup J_V$
2. $J_P \cap J_V = \emptyset$
3. $\forall j \in J_V : \Pi \Rightarrow VTimer(T_j, \mathcal{A}_j, \Delta_j, v)$
4. $\forall j \in J_P : \Pi \Rightarrow PTimer(T_j, \mathcal{A}_j, \Delta_j, v)$

⟨1⟩2. $\Pi^t \Rightarrow \Box Inv$, where

$$
\begin{aligned}
Inv \triangleq\ & 1.\wedge \forall j \in J : T_j \in [now, \infty] \\
& 2.\wedge now \in \mathbf{R} \\
& 3.\wedge \forall j \in J : (Enabled\ \langle\mathcal{A}_j\rangle_v = \\
& \qquad\qquad Enabled\ (\langle\mathcal{A}_j\rangle_v \wedge (now' = now)) \\
& 4.\wedge \forall k \in I \cap J : t_k \leq T_k \\
& 5.\wedge \forall j \in J_V : \neg Enabled\ \langle\mathcal{A}_j\rangle_v \Rightarrow (T_j = \infty)
\end{aligned}
$$

⟨2⟩1. $RT_v \Rightarrow \Box Inv.2$

PROOF: An invariance proof.

$\langle 2\rangle 2.$ For all $j \in J : \Pi \wedge \Box Inv.2 \wedge MaxTime(T_j) \Rightarrow \Box Inv.1$

PROOF: Assumption $\langle 0\rangle.0$c and an invariance proof.

$\langle 2\rangle 3.$ $\Pi^t \Rightarrow \Box Inv.3$

PROOF: Assumption $\langle 0\rangle.3$c.

$\langle 2\rangle 4.$ $\Pi^t \Rightarrow \Box Inv.4$

PROOF: Assumption $\langle 0\rangle.4$.

$\langle 2\rangle 5.$ $\Pi^t \Rightarrow \Box Inv.5$

PROOF: $\langle 1\rangle 1.3$, assumption $\langle 0\rangle.3$d, and an invariance proof.

$\langle 2\rangle 6.$ Q.E.D.

$\langle 1\rangle 3.$ 1. $\Pi = Init \wedge \Box \mathcal{M}$

2. $\Pi^t = Init^t \wedge \Box \mathcal{N}^t$, for some predicate $Init^t$.

where

$$\mathcal{M} \triangleq \wedge [\mathcal{N}]_w \\ \wedge \forall j \in J_P : PTact(T_j, \mathcal{A}_j, \Delta_j, v) \\ \wedge \forall j \in J_V : VTact(T_j, \mathcal{A}_j, \Delta_j, v)$$

$$\mathcal{N}^t \triangleq \wedge \mathcal{M} \\ \wedge RTact_v \\ \wedge \forall i \in I : MinTact(t_i, \mathcal{A}_i, v) \\ \wedge \forall j \in J : MaxTact(T_j)$$

$\langle 1\rangle 4.$ $\langle \mathcal{C}\rangle_{(now,v)}$ is a subaction of Π^t.

$\langle 2\rangle 1.$ For all j in J : $\langle \mathcal{B}^t_j\rangle_v$ is a subaction of Π^t.

ASSUME: $j \in J$

PROVE: $\Pi^t \Rightarrow \Box\,(\,Enabled\ \langle \mathcal{B}^t_j\rangle_v \ \Rightarrow\ Enabled\ (\langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{N}^t)\,)$

$\langle 3\rangle 1.$ $\Pi \ \Rightarrow\ \Box\,(Enabled\ (\langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{M}) \ \Rightarrow$
$\quad Enabled\ (\wedge \langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{M}$
$\quad\quad \wedge \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i\rangle_v \wedge \langle \mathcal{A}_j\rangle_v \wedge \mathcal{M})\,)\,)$

$\langle 4\rangle 1.$ $\Pi \ \Rightarrow\ \neg Enabled\ (\exists\, i \in I - \{j\} : \langle \mathcal{A}_i\rangle_v \wedge \langle \mathcal{A}_j\rangle_v \wedge \mathcal{M})$

PROOF: $\langle 1\rangle 3.1$, assumption $\langle 0\rangle.1$, and Lemma 4.4.

$\langle 4\rangle 2.$ Q.E.D.

PROOF: Lemma 4.2, substituting $\langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{M}$ for $\mathcal{A}$ and $\exists\, i \in I - \{j\}$: $\langle \mathcal{A}_i\rangle_v \wedge \langle \mathcal{A}_j\rangle_v \wedge \mathcal{M}$ for $\mathcal{B}$.

$\langle 3\rangle 2.$ $\wedge \langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{N}^t$
$\quad \wedge \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i\rangle_v \wedge \langle \mathcal{A}_j\rangle_v \wedge \mathcal{M})$
$\quad = \ \wedge \langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{M}$
$\quad\quad \wedge \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i\rangle_v \wedge \langle \mathcal{A}_j\rangle_v \wedge \mathcal{M})$

PROOF: $\wedge \langle \mathcal{B}^t_j\rangle_v \wedge \mathcal{N}^t$
$\quad \wedge \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i\rangle_v \wedge \langle \mathcal{A}_j\rangle_v \wedge \mathcal{M})$

$= $ [by definition of $\langle \mathcal{B}^t_j \rangle_v$]

$\quad \wedge\ \mathcal{A}^t_j \wedge (v' \neq v) \wedge (now' = now) \wedge \mathcal{N}^t$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i \rangle_v \wedge \langle \mathcal{A}_j \rangle_v \wedge \mathcal{M})$

$= $ [by definition of $\mathcal{N}^t$]

$\quad \wedge\ \mathcal{A}^t_j \wedge \mathcal{M} \wedge (v' \neq v) \wedge (now' = now)$

$\quad \wedge\ \forall i \in I : \mathcal{A}_i \Rightarrow (t_i \leq now)$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i \rangle_v \wedge \langle \mathcal{A}_j \rangle_v \wedge \mathcal{M})$

$= $ [by definition of $\mathcal{A}^t_j$]

$\quad \wedge\ P_j \wedge \mathcal{A}_j \wedge \mathcal{M} \wedge (v' \neq v) \wedge now' = now$

$\quad \wedge\ \forall i \in I : \mathcal{A}_i \Rightarrow (t_i \leq now)$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\mathcal{A}_i \wedge \mathcal{A}_j \wedge \mathcal{M} \wedge (v' \neq v))$

$= $ [by predicate logic]

$\quad \wedge\ P_j \wedge \mathcal{A}_j \wedge \mathcal{M} \wedge (v' \neq v) \wedge (now' = now)$

$\quad \wedge\ (j \in I) \Rightarrow (t_j \leq now)$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\mathcal{A}_i \wedge \mathcal{A}_j \wedge \mathcal{M} \wedge (v' \neq v))$

$= $ [by definition of $\langle \mathcal{A}_k \rangle_v$ and $\mathcal{M}$]

$\quad \wedge\ P_j \wedge \mathcal{A}_j \wedge \mathcal{M} \wedge (v' \neq v) \wedge (now' = now)$

$\quad \wedge\ (j \in I) \Rightarrow (t_j \leq now)$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i \rangle_v \wedge \langle \mathcal{A}_j \rangle_v \wedge \mathcal{M})$

$= $ [by definition of P_j]

$\quad \wedge\ P_j \wedge \mathcal{A}_j \wedge (v' \neq v) \wedge (now' = now) \wedge \mathcal{M}$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i \rangle_v \wedge \langle \mathcal{A}_j \rangle_v \wedge \mathcal{M})$

$= $ [by definition of $\langle \mathcal{B}^t_j \rangle_v$]

$\quad \wedge\ \langle \mathcal{B}^t_j \rangle_v \wedge \mathcal{M}$

$\quad \wedge\ \forall i \in I - \{j\} : \neg(\langle \mathcal{A}_i \rangle_v \wedge \langle \mathcal{A}_j \rangle_v \wedge \mathcal{M})$

⟨3⟩3. $\Pi \Rightarrow \Box\,(\, \mathit{Enabled}\ (\langle \mathcal{B}^t_j \rangle_v \wedge \mathcal{M}) \Rightarrow \mathit{Enabled}\ (\langle \mathcal{B}^t_j \rangle_v \wedge \mathcal{N}^t)\,)$

PROOF: ⟨3⟩1, ⟨3⟩2, and Lemma 4.5.

⟨3⟩4. $\Pi \Rightarrow \Box(P_j \wedge \mathit{Enabled}\ \langle \mathcal{B}_j \rangle_v \Rightarrow P_j \wedge \mathit{Enabled}\ (\langle \mathcal{B}_j \rangle_v \wedge \mathcal{M}))$

PROOF: ⟨1⟩3, Assumption ⟨0⟩.3a, and the definition of $\mathcal{B}_j$.

⟨3⟩5. $\Pi \Rightarrow \Box\,(\,(\mathit{Enabled}\ \langle \mathcal{B}^t_j \rangle_v) \Rightarrow \mathit{Enabled}\ (\langle \mathcal{B}^t_j \rangle_v \wedge \mathcal{M})\,)$

PROOF: ⟨3⟩4, the definition of $\mathcal{B}^t_j$, and Lemma 4.3.

⟨3⟩6. $\Pi \Rightarrow \Box((\mathit{Enabled}\ \langle \mathcal{B}^t_j \rangle_v) \Rightarrow \mathit{Enabled}\ (\langle \mathcal{B}^t_j \rangle_v \wedge \mathcal{N}^t))$

PROOF: ⟨3⟩5 and ⟨3⟩3.

⟨3⟩7. Q.E.D.

PROOF: ⟨3⟩6 and the definition of Π^t, which implies $\Pi^t \Rightarrow \Pi$.

⟨2⟩2. $\langle \mathit{NowT} \rangle_{now}$ is a subaction of Π^t.

⟨3⟩1. $\Pi^t \Rightarrow \Box\,(\,(\mathit{Enabled}\ \langle \mathit{NowT} \rangle_{now}) \Rightarrow (\mathit{Enabled}\ (\langle \mathit{NowT} \rangle_{now} \wedge \mathcal{M}))\,)$

$\langle 4 \rangle 1.$ $\Pi^t \Rightarrow \Box(\,(\mathit{Enabled}\ \mathit{NowT}) \Rightarrow (\mathit{Enabled}\ (\mathit{NowT} \wedge \mathcal{M}))\,)$

PROOF: $\langle 1 \rangle 3$, assumption $\langle 0 \rangle.2\text{b}$, and the definition of NowT, since Π^t implies Π.

$\langle 4 \rangle 2.$ Q.E.D.

PROOF: $\langle 4 \rangle 1$ and Lemma 4.3, substituting $\mathit{now} \neq T$ for P (since $\langle \mathit{NowT} \rangle_{\mathit{now}}$ equals $\mathit{NowT} \wedge (\mathit{now} \neq T)$).

$\langle 3 \rangle 2.$ $\Pi^t \Rightarrow \Box(\mathit{Enabled}\ (\langle \mathit{NowT} \rangle_{\mathit{now}} \wedge \mathcal{M}) \Rightarrow$
$\qquad\qquad \mathit{Enabled}\ (\langle \mathit{NowT} \rangle_{\mathit{now}} \wedge \mathcal{N}^t))$

$\langle 4 \rangle 1.$ $(\mathit{Inv} \wedge \langle \mathit{NowT} \rangle_{\mathit{now}} \wedge \mathcal{M}) \Rightarrow (\langle \mathit{NowT} \rangle_{\mathit{now}} \wedge \mathcal{N}^t)$

$\langle 5 \rangle 1.$ $\mathit{Inv} \wedge \langle \mathit{NowT} \rangle_{\mathit{now}} \Rightarrow \mathit{RTact}_v$

PROOF: $\langle 1 \rangle 3$ ($\mathit{Inv}.1$ and $\mathit{Inv}.2$).

$\langle 5 \rangle 2.$ $\forall i \in I : \langle \mathit{NowT} \rangle_{\mathit{now}} \Rightarrow \mathit{MinTact}(t_i, \mathcal{A}_i, v)$

PROOF: $\mathit{MinTact}(t_i, \mathcal{A}_i, v) = [\ldots]_v$, and $\langle \mathit{NowT} \rangle_{\mathit{now}}$ implies $v' = v$.

$\langle 5 \rangle 3.$ $\forall j \in J : \mathit{Inv} \wedge \langle \mathit{NowT} \rangle_{\mathit{now}} \wedge \mathcal{M} \Rightarrow \mathit{MaxTact}(T_j)$

ASSUME: 1. $j \in J$

2. $\mathit{Inv} \wedge \langle \mathit{NowT} \rangle_{\mathit{now}} \wedge \mathcal{M}$

PROVE: $\mathit{MaxTact}(T_j)$

$\langle 6 \rangle 1.$ CASE: $\mathit{Enabled}\ \langle \mathcal{A}_j \rangle_v$

$\langle 7 \rangle 1.$ $\mathit{now}' \leq T_j$

PROOF: $\mathit{Inv}.1$, $\mathit{Inv}.2$, $\langle \mathit{NowT} \rangle_{\mathit{now}}$, and the definitions of NowT, since case assumption $\langle 6 \rangle$ and the definition of T imply $T_j \geq T$.

$\langle 7 \rangle 2.$ CASE: $j \in J_P$

$\langle 8 \rangle 1.$ CASE: $T_j' = T_j$

PROOF: $\langle 7 \rangle 1$ and the definition of $\mathit{MaxTact}(T_j)$.

$\langle 8 \rangle 2.$ CASE: $T_j' = \mathit{now} + \Delta_j$

PROOF: $\mathit{Inv}.2$, $\langle \mathit{NowT} \rangle_{\mathit{now}}$, and definition of $\mathit{MaxTact}(T_j)$.

$\langle 8 \rangle 3.$ Q.E.D.

PROOF: $\langle 8 \rangle 1$, $\langle 8 \rangle 2$, and case assumption $\langle 7 \rangle$, since case assumption $\langle 6 \rangle$ and the definition of $\mathit{PTact}(T_j, \mathcal{A}_j, \Delta_j, v)$ imply that these are the only possibilities.

$\langle 7 \rangle 3.$ CASE: $j \in J_V$

$\langle 8 \rangle 1.$ $T_j = T_j'$

PROOF: Case assumption $\langle 7 \rangle$ and the definitions of $\mathcal{M}$ and $\mathit{VTact}(T_j, \mathcal{A}_j, \Delta_j, v)$, since $\langle \mathit{NowT} \rangle_v$ implies $v = v'$.

$\langle 8 \rangle 2.$ Q.E.D.

PROOF: $\langle 7 \rangle 1$, $\langle 8 \rangle 1$, and the definition of $\mathit{MaxTact}(T_j)$.

$\langle 7 \rangle 4.$ Q.E.D.

PROOF: $\langle 7 \rangle 2$, $\langle 7 \rangle 3$, $\langle 1 \rangle 1.1$, and assumption $\langle 5 \rangle.1$.

$\langle 6\rangle 2$. CASE: $\neg Enabled\ \langle \mathcal{A}_j \rangle_v$

$\langle 7\rangle 1$. $now, now' \in \mathbf{R}$ and $now' > now$.

PROOF: Inv.1, Inv.2, $\langle NowT \rangle_{now}$, and the definition of T.

$\langle 7\rangle 2$. CASE: $j \in J_P$

$\langle 8\rangle 1$. $T'_j \geq now'$

PROOF: $\langle 7\rangle 1$, case assumption $\langle 7\rangle$, Inv.1, and the definitions of $\mathcal{M}$ and $PTact(T_j, \mathcal{A}_j, \Delta_j, v)$.

$\langle 8\rangle 2$. Q.E.D.

PROOF: $\langle 8\rangle 1$ and the definition of $MaxTact(T_j)$.

$\langle 7\rangle 3$. CASE: $j \in J_V$

$\langle 8\rangle 1$. $T_j = \infty$

PROOF: By case assumption $\langle 7\rangle$ and Inv.5.

$\langle 8\rangle 2$. $T'_j = \infty$

PROOF: $\langle 8\rangle 1$, case assumption $\langle 7\rangle$, and the definitions of $\mathcal{M}$ and $VTact(T_j, \mathcal{A}_j, \Delta_j, v)$, since $\langle NowT \rangle_v$ implies $v = v'$.

$\langle 8\rangle 3$. Q.E.D.

PROOF: $\langle 7\rangle 1$, $\langle 8\rangle 2$, and the definition of $MaxTact(T_j)$.

$\langle 7\rangle 4$. Q.E.D.

PROOF: $\langle 7\rangle 2$, $\langle 7\rangle 3$, $\langle 1\rangle 1.1$, and assumption $\langle 5\rangle.1$.

$\langle 6\rangle 3$. Q.E.D.

PROOF: $\langle 6\rangle 1$ and $\langle 6\rangle 2$.

$\langle 5\rangle 4$. Q.E.D.

PROOF: $\langle 5\rangle 1$, $\langle 5\rangle 2$, $\langle 5\rangle 3$, and the definition of $\mathcal{N}^t$.

$\langle 4\rangle 2$. Q.E.D.

PROOF: $\langle 4\rangle 1$ and $\langle 1\rangle 2$, since by Lemma 4.3 and Lemma 4.5, $Inv \wedge \mathcal{D} \Rightarrow \mathcal{E}$ implies $\Box Inv \Rightarrow \Box((Enabled\ \mathcal{D}) \Rightarrow (Enabled\ \mathcal{E}))$, for any actions $\mathcal{D}$ and $\mathcal{E}$.

$\langle 3\rangle 3$. Q.E.D.

PROOF: $\langle 3\rangle 1$ and $\langle 3\rangle 2$.

$\langle 2\rangle 3$. Q.E.D.

$\langle 1\rangle 5$. $\Pi^t \wedge \mathrm{WF}_{\langle now, v\rangle}(\mathcal{C}) \Rightarrow NZ$

ASSUME: $r \in \mathbf{R}$

PROVE: $\Pi^t \wedge \mathrm{WF}_{\langle now, v\rangle}(\mathcal{C}) \Rightarrow ((now = r) \leadsto (now \in [r + \Delta, \infty)))$

LET: $U \triangleq \{j \in J : T_j < r + \Delta\}$

$V \triangleq \{j \in J : now < T_j < r + \Delta\}$

$TimerAct \triangleq \forall j \in J : \vee\ VTact(T_j, \mathcal{A}_j, \Delta_j, v)$
$\qquad\qquad\qquad\qquad \vee\ PTact(T_j, \mathcal{A}_j, \Delta_j, v)$

$\langle 2\rangle 1.$ $Inv \Rightarrow Enabled\ \langle \mathcal{C}\rangle_{(now,v)}$

$\langle 3\rangle 1.$ CASE: $T \neq now$

PROOF: By the definition of $\mathcal{C}$, since case assumption $\langle 3\rangle$ implies $Enabled\ \langle NowT\rangle_{now}$.

$\langle 3\rangle 2.$ CASE: $T = now$

$\langle 4\rangle 1.$ Choose $j \in J$ such that $(T_j = T) \wedge Enabled\ \langle \mathcal{A}_j\rangle_v$.

PROOF: Inv.2, case assumption $\langle 3\rangle$, and the definition of T.

$\langle 4\rangle 2.$ $Enabled\ \langle \mathcal{B}_j\rangle_v$

PROOF: $\langle 4\rangle 1$, Inv.3, and the definition of $\langle \mathcal{B}_j\rangle_v$.

$\langle 4\rangle 3.$ $Enabled\ \langle \mathcal{B}^t_j\rangle_v$

PROOF: $\langle 4\rangle 2$, Inv.4, case assumption $\langle 3\rangle$, and the definition of $\langle \mathcal{B}^t_j\rangle_v$.

$\langle 4\rangle 4.$ Q.E.D.

PROOF: Case assumption $\langle 3\rangle$, $\langle 4\rangle 3$, and the definition of $\mathcal{C}$.

$\langle 3\rangle 3.$ Q.E.D.

PROOF: $\langle 3\rangle 1$ and $\langle 3\rangle 2$.

$\langle 2\rangle 2.$ $\Pi^t \Rightarrow$

$$\Box((now = r) \wedge \Box(now \in (-\infty, r+\Delta)) \Rightarrow \Box(now \in [r, r+\Delta)))$$

$\langle 3\rangle 1.$ $\Box[RTact_v]_{now} \Rightarrow ((now = r) \Rightarrow \Box(now \in [r, \infty)))$

PROOF: A standard invariance argument.

$\langle 3\rangle 2.$ $\Box[RTact_v]_{now} \Rightarrow \Box((now = r) \Rightarrow \Box(now \in [r, \infty)))$

PROOF: $\langle 3\rangle 1$ and simple temporal logic.

$\langle 3\rangle 3.$ $\Pi^t \Rightarrow \Box((now = r) \Rightarrow \Box(r \leq now))$

PROOF: $\langle 3\rangle 2$, since $\Pi^t \Rightarrow RT_v$ and $RT_v \Rightarrow \Box[RTact_v]_{now}$.

$\langle 3\rangle 4.$ Q.E.D.

PROOF: $\langle 3\rangle 3$, using the temporal logic tautology

$$(F \leadsto G) \Rightarrow ((F \wedge \Box H) \leadsto (G \wedge \Box H))$$

$\langle 2\rangle 3.$ ASSUME: $j \in J$

PROVE: 1. $\Pi^t \wedge \Box(now \in [r, r+\Delta)) \Rightarrow \Box((j \notin U) \Rightarrow \Box(j \notin U))$

2. $\Pi^t \wedge \Box(now \in [r, r+\Delta)) \Rightarrow \Box((j \notin V) \Rightarrow \Box(j \notin V))$

PROOF: A standard invariance proof, using assumption $\langle 0\rangle$.3b.

$\langle 2\rangle 4.$ $\Pi^t \wedge \Box(now \in [r, r+\Delta)) \wedge \mathrm{WF}_{(now,v)}(\mathcal{C}) \Rightarrow \Diamond\Box(U = \emptyset)$

PROOF SKETCH: The set V consists of those timers in U that are not equal to now. To prove that U is eventually empty, we show that, whenever U is nonempty, eventually U or V gets smaller. Since U and V are finite, U must eventually become empty.

$\langle 3\rangle 1.$ ASSUME: U_0 and V_0 sets, with $U_0 \neq \emptyset$.

PROVE: $\wedge\ \Box((U \subseteq U_0) \wedge (V \subseteq V_0))$
$\wedge\ \Box(now \in [r, r + \Delta))$
$\wedge\ \Box Inv$
$\wedge\ \Box[TimerAct]_{(now,v)}$
$\wedge\ \mathrm{WF}_{(now,v)}(\mathcal{C})$
$\Rightarrow ((U = U_0) \wedge (V = V_0)) \leadsto$
$(U \subset U_0) \vee ((U \subseteq U_0) \wedge (V \subset V_0))$

PROOF SKETCH: This is a straightforward application of rule WF1 (Lemma 3), with the following substitutions.

LET: $I \triangleq$ 1.$\wedge$ $now \in [r, r + \Delta)$
2.$\wedge$ $(U \subseteq U_0) \wedge (V \subseteq V_0)$
3.$\wedge$ Inv

$P \triangleq (U = U_0) \wedge (V = V_0)$

$Q \triangleq (U \subset U_0) \vee ((U \subseteq U_0) \wedge (V \subset V_0))$

$\mathcal{N} \triangleq TimerAct$

$\mathcal{A} \triangleq \mathcal{C}$

$f \triangleq (now, v)$

$\langle 4\rangle 1.$ $I' \Rightarrow (P' \vee Q')$
PROOF: Obvious.

$\langle 4\rangle 2.$ ASSUME: $P \wedge I \wedge I' \wedge \langle \mathcal{N} \wedge \mathcal{A}\rangle_f$
PROVE: Q'

$\langle 5\rangle 1.$ CASE: $T = now$
$\langle 6\rangle 1.$ Choose j in J such that $(T_j = T) \wedge \langle \mathcal{B}^t_j\rangle_v$.
PROOF: $\mathcal{A}$ and case assumption $\langle 5\rangle$.
$\langle 6\rangle 2.$ $T'_j \geq r + \Delta$
PROOF: $\langle 6\rangle 1$, $I.1$, the definition of $\mathcal{B}^t_j$, and $TimerAct$.
$\langle 6\rangle 3.$ $j \in U \wedge j \notin U'$
PROOF: $\langle 6\rangle 1$, case assumption $\langle 5\rangle$, and $I.1$.
$\langle 6\rangle 4.$ $j \notin U'$
PROOF: $\langle 6\rangle 2$.
$\langle 6\rangle 5.$ Q.E.D.
PROOF: $\langle 6\rangle 3$, $\langle 6\rangle 4$, and $I'.2$.

$\langle 5\rangle 2.$ CASE: $T \neq now$
$\langle 6\rangle 1.$ $(now' = T) \wedge (v' = v)$
PROOF: $\mathcal{A}$ and case assumption $\langle 5\rangle$.
$\langle 6\rangle 2.$ CASE: $T \in (now, r + \Delta)$
$\langle 7\rangle 1.$ Choose j in J such that $(T_j = T) \wedge Enabled\ \langle \mathcal{A}\rangle_v$.
PROOF: Case assumption $\langle 6\rangle$ and the definition of T.

$\langle 7\rangle 2.$ $\vee\ T_j' = \mathit{now}'$
$\vee\ T_j' \in [r + \Delta, \infty]$

PROOF: $\langle 6\rangle 1$, $\langle 7\rangle 1$, $I.1$, and *TimerAct*.

$\langle 7\rangle 3.$ $j \in V$

PROOF: $\langle 7\rangle 1$, case assumption $\langle 6\rangle$, and the definition of V.

$\langle 7\rangle 4.$ $j \notin V'$

PROOF: $\langle 7\rangle 2$ and the definition of V.

$\langle 7\rangle 5.$ Q.E.D.

$\langle 7\rangle 3$, $\langle 7\rangle 4$, and $I'.2$.

$\langle 6\rangle 3.$ CASE: $T \in [r + \Delta, \infty]$

PROOF: Impossible by $\langle 6\rangle 1$ and $I'.1$.

$\langle 6\rangle 4.$ Q.E.D.

PROOF: $\langle 6\rangle 2$, $\langle 6\rangle 3$, $I.3.1$, and case assumption $\langle 5\rangle$.

$\langle 5\rangle 3.$ Q.E.D.

PROOF: $\langle 5\rangle 1$ and $\langle 5\rangle 2$.

$\langle 4\rangle 3.$ $P \wedge I \Rightarrow \mathit{Enabled}\ \langle \mathcal{A}\rangle_f$

PROOF: $\langle 2\rangle 1$.

$\langle 4\rangle 4.$ Q.E.D.

PROOF: $\langle 4\rangle 1$, $\langle 4\rangle 2$, $\langle 4\rangle 3$, and Lemma 3.

$\langle 3\rangle 2.$ ASSUME: U_0 and V_0 sets, with $U_0 \neq \emptyset$.

PROVE: $\Pi^t \wedge \Box(\mathit{now} \in [r, r + \Delta)) \wedge \mathrm{WF}_{(\mathit{now},v)}(\mathcal{C}) \Rightarrow$
$((U = U_0) \wedge (V = V_0)) \leadsto$
$((U \subset U_0) \vee ((U \subseteq U_0) \wedge (V \subset V_0)))$

$\langle 4\rangle 1.$ $\Pi^t \Rightarrow \Box((U \subseteq U_0) \wedge (V \subseteq V_0) \Rightarrow \Box((U \subseteq U_0) \wedge (V \subseteq V_0)))$

PROOF: Follows from $\langle 2\rangle 3$.

$\langle 4\rangle 2.$ Q.E.D.

PROOF: $\langle 3\rangle 1$, $\langle 4\rangle 1$, and $\langle 1\rangle 2$, since $\Pi^t \Rightarrow \Box[\mathit{TimerAct}]_{(\mathit{now},v)}$ by assumption $\langle 0\rangle.3\mathrm{b}$.

$\langle 3\rangle 3.$ Q.E.D.

PROOF: Since U and V are finite by assumption $\langle 0\rangle.0\mathrm{d}$, it follows from $\langle 3\rangle 2$ and the Lattice Rule [13] that $\Pi^t \wedge \Box(\mathit{now} \in [r, r + \Delta)) \wedge \mathrm{WF}_{(\mathit{now},v)}(\mathcal{C})$ implies $\Diamond(U = \emptyset)$. By $\langle 2\rangle 3$, $\Pi^t \wedge \Box(\mathit{now} \in [r, r + \Delta))$ implies $\Diamond(U = \emptyset) \Rightarrow \Diamond\Box(U = \emptyset)$.

$\langle 2\rangle 5.$ $\wedge\ \Box(\mathit{now} \in [r, r + \Delta))$
$\wedge\ \Box(U = \emptyset)$
$\wedge\ \Box \mathit{Inv}$
$\wedge\ \mathrm{WF}_{(\mathit{now},v)}(\mathcal{C})$
$\Rightarrow \mathsf{true} \leadsto (\mathit{now} \geq r + \Delta)$

LET: $I \triangleq$ 1.$\land$ $U = \emptyset$
2.$\land$ $now \in [r, r + \Delta)$
3.$\land$ Inv

$P \triangleq \mathsf{true}$

$Q \triangleq now \geq r + \Delta$

$\mathcal{N} \triangleq \mathsf{true}$

$\mathcal{A} \triangleq \mathcal{C}$

$f \triangleq (now, v)$

$\langle 3 \rangle 1.$ $I \land I' \land [\mathcal{N}]_f \Rightarrow (P' \lor Q')$

PROOF: Immediate, since $P' = \mathsf{true}$.

$\langle 3 \rangle 2.$ ASSUME: $P \land I \land I' \land \langle \mathcal{N} \land \mathcal{A} \rangle_f$
PROVE: Q'

$\langle 4 \rangle 1.$ $T \in [r + \Delta, \infty]$

PROOF: By I.

$\langle 4 \rangle 2.$ $\langle NowT \rangle_{now}$

PROOF: By $\langle 4 \rangle 1$, I.2, and $\mathcal{A}$.

$\langle 4 \rangle 3.$ Q.E.D.

PROOF: $\langle 4 \rangle 1$, $\langle 4 \rangle 2$, and the definition of $NowT$.

$\langle 3 \rangle 3.$ $P \land I \Rightarrow Enabled\ \langle \mathcal{A} \rangle_f$

PROOF: $\langle 3 \rangle 1$.

$\langle 3 \rangle 4.$ Q.E.D.

PROOF: $\langle 3 \rangle 1$, $\langle 3 \rangle 2$, $\langle 3 \rangle 3$, and Lemma 3.

$\langle 2 \rangle 6.$ Q.E.D.

$\langle 1 \rangle 6.$ $(E \mathrel{-\!\!\triangleright} M^t, E \Rightarrow \mathrm{WF}_{(now,v)}(\mathcal{C}))$ is μ-machine realizable.

$\langle 2 \rangle 1.$ M^t constrains at most μ.

$\langle 3 \rangle 1.$ M constrains at most μ.

PROOF: Assumption $\langle 0 \rangle .5$.

$\langle 3 \rangle 2.$ RT_v constrains at most μ.

PROOF: Assumption $\langle 0 \rangle .6$b.

$\langle 3 \rangle 3.$ For all i in I, $MinTime(t_i, \mathcal{A}_i, v)$ constrains at most μ.

PROOF: By definition of $MinTime$, a step violates $MinTime(t_i, \mathcal{A}_i, v)$ only if it is an $\langle \mathcal{A}_i \rangle_v$ step, so this follows from Assumption $\langle 0 \rangle .6$a.

$\langle 3 \rangle 4.$ For all j in J, $MaxTime(T_j)$ constrains at most μ.

PROOF: Assumption $\langle 0 \rangle .6$b.

$\langle 3 \rangle 5.$ Q.E.D.

PROOF: $\langle 3 \rangle 1$–$\langle 3 \rangle 4$ and the definition of M^t.

$\langle 2 \rangle 2.$ Q.E.D.

$\langle 1 \rangle 7$. (true, NZ) is μ-machine realizable.
$\langle 1 \rangle 8$. Q.E.D.

References

[1] Martín Abadi and Leslie Lamport. Composing specifications. In J. W. de Bakker, W.-P. de Roever, and G. Rozenberg, editors, *Stepwise Refinement of Distributed Systems*, volume 430 of *Lecture Notes in Computer Science*, pages 1–41. Springer-Verlag, May/June 1989.

[2] Martín Abadi and Leslie Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, May 1991.

[3] Martín Abadi and Gordon Plotkin. A logical view of composition. Research Report 86, Digital Equipment Corporation, Systems Research Center, May 1992.

[4] Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, October 1985.

[5] Krzysztof R. Apt, Nissim Francez, and Shmuel Katz. Appraising fairness in languages for distributed programming. *Distributed Computing*, 2:226–241, 1988.

[6] Arthur Bernstein and Paul K. Harter, Jr. Proving real time properties of programs with temporal logic. In *Proceedings of the Eighth Symposium on Operating Systems Principles*, pages 1–11, New York, 1981. ACM. *Operating Systems Review 15*, 5.

[7] K. Mani Chandy and Jayadev Misra. *Parallel Program Design*. Addison-Wesley, Reading, Massachusetts, 1988.

[8] David L. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits*. PhD thesis, Carnegie Mellon University, February 1988.

[9] Michael Fischer. Re: Where are you? E-mail message to Leslie Lamport. Arpanet message number 8506252257.AA07636@YALE-BULLDOG.YALE.ARPA (47 lines), June 25, 1985 18:56:29 EDT.

[10] Cliff B. Jones. Specification and design of (parallel) programs. In R. E. A. Mason, editor, *Information Processing 83: Proceedings of the IFIP 9th World Congress*, pages 321–332. IFIP, North-Holland, September 1983.

[11] Leslie Lamport. An assertional correctness proof of a distributed algorithm. *Science of Computer Programming*, 2(3):175–206, December 1982.

[12] Leslie Lamport. A fast mutual exclusion algorithm. *ACM Transactions on Computer Systems*, 5(1):1–11, February 1987.

[13] Leslie Lamport. The temporal logic of actions. Research Report 79, Digital Equipment Corporation, Systems Research Center, December 1991.

[14] Keith Marzullo, Fred B. Schneider, and Navin Budhiraja. Derivation of sequential, real-time process-control programs. In André M. van Tilborg and Gary M. Koob, editors, *Foundations of Real-Time Computing: Formal Specifications and Methods*, chapter 2, pages 39–54. Kluwer Academic Publishers, Boston, Dordrecht, and London, 1991.

[15] Jayadev Misra and K. Mani Chandy. Proofs of networks of processes. *IEEE Transactions on Software Engineering*, SE-7(4):417–426, July 1981.

[16] Peter G. Neumann and Leslie Lamport. Highly dependable distributed systems. Technical report, SRI International, June 1983. Contract Number DAEA18-81-G-0062, SRI Project 4180.

[17] Amir Pnueli. In transition from global to modular temporal reasoning about programs. In Krzysztof R. Apt, editor, *Logics and Models of Concurrent Systems*, NATO ASI Series, pages 123–144. Springer-Verlag, October 1984.

[18] Fred B. Schneider, Bard Bloom, and Keith Marzullo. Putting time into proof outlines. In J. W. de Bakker, C. Huizing, W.-P. de Roever, and G. Rozenberg, editors, *Real-Time: Theory in Practice*, volume 600 of *Lecture Notes in Computer Science*, pages 618–639, Berlin, Heidelberg, New York, 1992. Springer-Verlag.

Index