Modeling Time in Computing:
A Taxonomy and a Comparative Survey

Carlo A. Furia, Dino Mandrioli,
Angelo Morzenti, and Matteo Rossi

October 9, 2018

Abstract

The increasing relevance of areas such as real-time and embedded systems, pervasive computing, hybrid systems control, and biological and social systems modeling is bringing a growing attention to the temporal aspects of computing, not only in the computer science domain, but also in more traditional fields of engineering.

This article surveys various approaches to the formal modeling and analysis of the temporal features of computer-based systems, with a level of detail that is also suitable for nonspecialists. In doing so, it provides a unifying framework, rather than just a comprehensive list of formalisms.

The article first lays out some key dimensions along which the various formalisms can be evaluated and compared. Then, a significant sample of formalisms for time modeling in computing are presented and discussed according to these dimensions. The adopted perspective is, to some extent, historical, going from "traditional" models and formalisms to more modern ones.

Contents

1 Introduction

In many fields of science and engineering, the term *dynamics* is intrinsically bound to a notion of time. In fact, in classical physics a mathematical model of a dynamical system most often consists of a set of equations that state a relation between a *time variable* and other quantities characterizing the system, often referred to as system *state*.

In the theory of computation, conversely, the notion of time does not always play a major role. At the root of the theory, a problem is formalized as a *function* from some input domain to an output range. An algorithm is a process aimed at computing the value of the function; in this process, dynamic aspects are usually abstracted away, since the only concern is the result produced.

Timing aspects, however, are quite relevant in computing too, for many reasons; let us recall some of them by adopting a somewhat historical perspective.

- First, *hardware* design leads down to electronic devices where the physical world of circuits comes back into play, for instance when the designer must verify that the sequence of logical gate switches that is necessary to execute an instruction can be completed within a clock's tick. The time models adopted here are borrowed from physics and electronics, and range from differential equations on continuous time for modeling devices and circuits, to discrete time (coupled with discrete mathematics) for describing logical gates and digital circuits.

- When the level of description changes from hardware to software, physical time is progressively disregarded in favor of more "coarse-grained" views of time, where a time unit represents a computational step, possibly in a high-level programming language; or it is even completely abstracted away when adopting a purely functional view of software, as a mapping from some input to the computed output. In this framework, *computational complexity* theory was developed as a natural complement of computability theory: it was soon apparent that knowing an algorithm to solve a problem is not enough if the execution of such an algorithm takes an unaffordable amount of time. As a consequence, models of abstract machines have been developed or refined so as to measure the time needed for their operations. Then, such an abstract notion of time measure (typically the number of elementary computation steps) could be mapped easily to physical time.

- The advent of *parallel* processing mandated a further investigation of timing issues in the theory of computing. To coordinate appropriately the various concurrent activities, in fact, it is necessary to take into account their temporal evolution. Not by chance the term synchronization derives from the two Greek words $\sigma\upsilon\nu$ (meaning "together") and $\chi\rho o\nu o\sigma$ (meaning "time").

- In relatively recent times the advent of novel methods for the design and verification of *real-time* systems also requires the inclusion of the envi-

ronment with which the computer interacts in the models under analysis. Therefore the various activities are, in general, not fully synchronized, that is, it is impossible to delay indefinitely one activity while waiting for another one to come alive. Significant classes of systems that possess real-time features are, among others, social organizations (in a broad sense), and distributed and embedded systems. For instance, in a plant control system, the control apparatus must react to the stimuli coming from the plant at a pace that is mandated by the dynamics of the plant. Hence physical time, which was progressively abstracted away, once again plays a prominent role.

As a consequence, some type of time modeling is necessary in the theory of computing as well as in any discipline that involves dynamics. Unlike other fields of science and engineering, however, time modeling in computing is far from exhibiting a unitary and comprehensive framework that would be suitable in a general way for most needs of system analysis: this is probably due to the fact that the issue of time modeling arose in different fields, in different circumstances, and was often attacked in a fairly *ad hoc* manner.

In this article we survey various approaches that have been proposed to tackle the issue of time modeling in computing. Rather than pursuing an exhaustive list of formalisms, our main goal is to provide a unifying framework so that the various models can be put in perspective, compared, evaluated, and possibly adapted to the peculiar needs of specific application fields. In this respect, we selected the notations among those that are most prominent in the scientific literature, both as basic research targets and as useful modeling tools in applications. We also aimed at providing suitable "coverage" of the most important features that arise in time modeling. We tried to keep our exposition at a level palatable for the nonspecialist who wishes to gain an overall but not superficial understanding of the issue. Also, although the main goal of time modeling is certainly to use it in the practice of system design, we focus on the conceptual aspects of the problem (what can and cannot be done with a given model; how easy it is to derive properties, etc.) rather than on practical "recipes" of how to apply a formal language in specific projects. The presentation is accompanied by many examples from different domains; most of them are inspired by embedded systems concepts, others, however, show that the same concepts apply as well to a wider class of systems such as biological and social ones.

We deliberately excluded from our survey time modeling approaches based on stochastic formalisms. This sector is certainly important and very relevant for several applications, and it has recently received increasing attention from the research community (e.g., [RKNP04, DK05]). In fact, most of the formal notations presented in this survey have some variants that include stochastic or probabilistic features. However, including such variants in our presentation would have also required us to present the additional mathematical notions and tools needed to tackle stochastic processes. These are largely different from the notions discussed in the article, which aim at gaining "certainty" (e.g., "the system will not crash under any circumstances") rather than a "measure of

uncertainty" (e.g., "the system will crash with probability 10^{-3}") as happens with probabilistic approaches. Thus, including stochastic formalisms would have weakened the focus of the article and made it excessively long.

The first part of this article introduces an informal reference framework within which the various formalisms can be explained and evaluated. First, Section 2 presents the notion of language, and gives a coarse categorization of formalisms; then, Section 3 proposes a collection of "dimensions" along which the various modeling approaches can be classified.

The second part of the article is the actual survey of time modeling formalisms. We do not aim at exhaustiveness; rather, we focus on several relevant formalisms, those that better exemplify the various approaches found in the literature, and analyze them through the dimensions introduced in Section 3. We complement the exposition, however, with an extensive set of bibliographic references. In the survey, we follow a rather historical ordering: Section 4 summarizes the most traditional ways of taking care of timing aspects in computing, whereas Section 5 is devoted to the more recent proposals, often motivated by the needs of new, critical, real-time applications. Finally, Section 6 contains some concluding remarks.

2 Languages and Interpretations

When studying the different ways in which time has been represented in the literature, and the associated properties, two aspects must be considered: the language used to describe time and the way in which the language is interpreted.[1] Let us illustrate this point in some detail.

A *language* (in the broad sense of the term) is the device that we employ to describe anything of interest (an object, a function, a system, a property, a feature, etc.). Whenever we write a "sentence" in a language (any language), a *meaning* is also attached to that sentence. Depending on the extent to which mathematical concepts are used to associate a sentence with its meaning, a language can be informal (no elements of the language are associated with mathematical concepts), semiformal (some are, but not all), or formal (everything is).

More precisely, given a sentence ϕ written in some language $\mathcal{L}$, we *can* assign it a variety of *interpretations*; we then define the meaning of ϕ by establishing which ones, among all possible interpretations, are those that are *actually associated* with it (in other words, by deciding which interpretations are "meaningful", and which ones are not); we say that an interpretation *satisfies* a sentence with which it is associated, or, dually, that the sentence *expresses* its associated interpretations. In the rest of this article, we will sometimes refer to the language as the *syntax* used to describe a sentence, as opposed to the interpretations that the latter expresses, which constitute its *semantics*.

---

[1]Such interpretations are referred to in mathematical logic as the *models* of ϕ; in this article we will in general avoid this terminology, as it might generate some confusion with the different notion of a model as "description of a system".

In this survey we mainly deal with languages that have the distinguishing feature of including a notion of time. Then, the interpretations associated with sentences in these languages include a notion of *temporal evolution* of elements; that is, they define what value is associated with an element at a certain time instant. As a consequence, we refer to the possible interpretations of sentences in timed languages as *behaviors*. In fact, the semantics of every formal language that has a notion of time is defined through some idea of "behavior" (or trace): infinite words for linear temporal logic [Eme90], timed words for timed automata [AD94], sequences of markings for Petri nets [Pet81], and so on.

For example, a behavior of a system is a mapping $b : \mathbb{T} \to S$, where $\mathbb{T}$ is a temporal domain and S is a state space; the behavior represents the system's state (i.e., the value of its elements) in the various time instants of $\mathbb{T}$.

Let us consider a language $\mathcal{L}$ and a sentence ϕ written in $\mathcal{L}$. The nature of ϕ depends on $\mathcal{L}$; for example it could be a particular kind of graph if $\mathcal{L}$ is some type of automaton (a Statechart, a Petri net, etc.), a logic formula if $\mathcal{L}$ is some sort of logic, and so on. Given a behavior b in the system model, we write $b \models \phi$ to indicate that b satisfies ϕ, that is, it is one of the behaviors expressed by the sentence. The satisfaction relation $\models$ is not general, that is, it is language-dependent (it is, in fact, $\models_{\mathcal{L}}$, but we omit the subscript for conciseness), and is part of the definition of the language.

Figure 1 depicts informally the relations among behaviors, language, system descriptions, real world, and semantics. Solid arrows denote that the entities they point to are obtained by combining elements of entities they originate from; for instance, a system description consists of formalizations of (parts of) the real world through sentences in some language. Dashed arrows, on the other hand, denote indirect influences; for example, the features of a language can suggest the adoption of certain behavioral structures. Finally, the semantics of a system is given by the set of all behaviors b satisfying system description Φ. These relations will become clearer in the following examples.

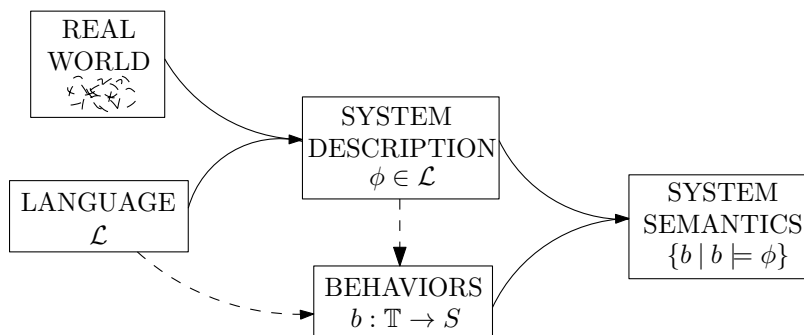

Figure 1: Behaviors, language, system descriptions, world.

Example 1 (Continuous, Scalar Linear Dynamic System)**.** Suppose $\mathcal{L}$ is the language of differential equations used to describe traditional linear dynamic

systems. With such a language we might model, for example, the simple RC circuit of Figure 2. In this case, the sentence ϕ that describes the system could be $\dot{q} = -\frac{1}{RC}q$ (where q is the charge of the capacitor); then, a behavior b that satisfies ϕ (i.e., such that $b \models \phi$) is $b(t) = C_0 e^{-t/RC}$, where C_0 is the initial charge of the capacitor, at the time when the circuit is closed (conventionally assumed to be 0).


Figure 2: An example of sentence in graphical language describing electric circuits.

To conclude this section, let us present a widely used categorization of languages that, while neither sharp nor precise, nevertheless quickly conveys some important features of a language.

Languages are often separated into two broad classes: *operational* languages and *descriptive* languages [GJM02].

Operational languages are well-suited to describe the *evolution* of a system starting from some initial state. Common examples of operational languages are the differential equations used to describe dynamic systems in control theory (see Section 4.1), automata-based formalisms (finite-state automata, Turing machines, timed automata, which are described in Sections 4.3 and 5.1.1) and Petri nets (which are presented in Section 5.1.2). Operational languages are usually based on the key concepts of *state* and *transition* (or *event*), so that a system is modeled as evolving from a state to the next one when a certain transition/event occurs. For example, an operational description of the dynamics of a digital safe could be the following:

Example 2 (Safe, operational formulation)**.** "When the last digit of the correct security code is entered, the safe opens; if the safe remains open for three minutes, it automatically closes."

Descriptive languages, instead, are better suited to describing the *properties* (static or dynamic) that the system must satisfy. Classic examples of descriptive languages are logic-based and algebra-based formalisms, such as those presented in Section 5.2. An example of descriptive formulation of the properties of a safe is the following:

Example 3 (Safe, descriptive formulation)**.** "The safe is open if and only if the correct security code has been entered no more than three minutes ago."

As mentioned above, the distinction between operational and descriptive languages is not as sharp as it sounds, for the following reasons. First, it is

possible to use languages that are operational to describe system properties (e.g., [AD94] used timed automata to represent both the system and its properties to be verified through model checking), and languages that are descriptive to represent the system evolution with state/event concepts [GM01] (in fact, the dynamics of Example 2 can be represented using a logic language, while the property of Example 3 can be formalized through an automata-based language). In addition, it is common to use a combination of operational and descriptive formalisms to model and analyze systems in a so-called *dual-language approach*. In this dual approach, an operational language is used to represent the dynamics of the system (i.e., its evolution), while its requirements (i.e., the properties that it must satisfy, and which one would like to verify in a formal manner) are expressed in a descriptive language. Model checking techniques [CGP00, HNSY94] and the combination of Petri nets with the TRIO temporal logic [FMM94] are examples of the dual language approach.

3 Dimensions of the Time Modeling Problem

When describing the modeling of time several distinctive issues need to be considered. These constitute the "dimensions" of the problem from the perspective of this article. They will help the analysis of how time is modeled in the literature, which is carried out in Sections 4 and 5.

Some of the dimensions proposed here are indicative of issues that are pervasive in the modeling of time in the literature (e.g., using discrete vs. continuous time domains); others shed more light on subtle aspects of some formalisms. We believe that the systematic, though not exhaustive, analysis of the formalisms surveyed in Sections 4 and 5 against the dimensions proposed below should not only provide the reader with an overall comparative assessment of the formalisms described in this article, but also help her build her own evaluation of other present and future formalisms in the literature.

3.1 Discrete vs. Dense Time Domains

A first natural categorization of the formalisms dealing with time-dependent systems and the adopted time model is whether such a model is a discrete or dense set.

A discrete set consists of isolated points, whereas a dense set (ordered by "<") is such that for every two points t_1, t_2, with $t_1 < t_2$, there is always another point t_3 in between, that is, $t_1 < t_3 < t_2$. In the scientific literature and applications, the most widely adopted discrete time models are natural and integer numbers — herewith denoted as $\mathbb{N}$ and $\mathbb{Z}$, respectively — whereas the typical dense models are rational and real numbers — herewith denoted as $\mathbb{Q}$ and $\mathbb{R}$, respectively. For instance, differential equations are normally stated with respect to real variable domains, whereas difference equations are defined on integers. Computing devices are formalized through discrete models when their behavior is paced by a clock, so that it is natural to measure time by

counting clock ticks, or when they deal with (i.e., measure, compute, or display) values in discrete domains.

Besides the above well-known premises, however, a few more accurate distinctions are useful to better evaluate and compare the many formalisms available in the literature and those that will be proposed in the future.

Continuous vs. Noncontinuous Time Models

Normally in mathematics, continuous time models (i.e., those in which the temporal domain is a dense set such that every nonempty set with an upper bound has a least upper bound) such as real numbers are preferred to other dense domains such as the rationals, thanks to their completeness/closure with respect to all common operations (otherwise, referring to $\sqrt{2}$ or π would be cumbersome). Instead, normal numerical algorithms deal with rational numbers since they can approximate real numbers — which cannot be represented by a finite sequence of digits — up to any predefined error. There are cases, however, where the two sets exhibit a substantial difference. For instance, assume that a system is composed of two devices whose clocks c_1 and c_2 are incommensurable (i.e., such that there are no integer numbers n, m such that $nc_1 = mc_2$). In such a case, if one wants to "unify" the system model, $\mathbb{Q}$ is not a suitable temporal domain. Also, there are some sophisticated time analysis algorithms that impose the restriction that the time domain is $\mathbb{Q}$ but not $\mathbb{R}$. We refer to one such algorithm when discussing Petri nets in Section 5.1.2.

Finite or Bounded Time Models

Normal system modeling assumes behaviors that may proceed indefinitely in the future (and maybe in the past), so that it is natural to model time as an unbounded set. There are significant cases, however, where all relevant system behaviors can be *a priori* enclosed within a bounded "time window". For instance, braking a car to a full stop requires at most a few seconds; thus, if we want to model and analyze the behavior of an anti-lock braking system there is no loss of generality if we assume as a temporal domain, say, the real range $[0 \ldots 60\text{secs}]$. In many cases this restriction highly simplifies several analyses and/or simulation algorithms. In other cases the system under consideration is periodic; thus, knowing its behaviors during a full period provides enough information to determine its relevant properties over the whole time axis.

Hybrid Systems

In this article, by hybrid system model we mean a model that uses both discrete and dense domains. There are several circumstances when this may occur, mainly but not exclusively related to the problem of integrating heterogeneous components: for instance, monitoring and controlling a continuous process by means of a digital device.

- A system (component) with a discrete — possibly finite — set of states is modeled as evolving in a dense time domain. In such a case its behavior is graphically described as a *square wave* form (see Figure 3) and its state can be formalized as a piecewise constant function of time, as shown in Figure 3.


Figure 3: A square-wave form over dense time.

- In a fairly symmetric way, a continuous behavior can be sampled at regular intervals, as exemplified in Figure 4.

Figure 4: A sampled continuous behavior.

- A more sophisticated, but fairly common, case of hybridness may arise when a model evolves through a discrete sequence of "steps" while other, independent, variables evolve taking their values in nondiscrete domains, for instance, finite state automata augmented with dense-timed clock variables. We see examples of such models in Section 5.1, in which timed and hybrid automata are discussed.

Time Granularity

In some sense, time granularity can be seen as a special case of hybridness. We say that two system components have different time granularities when their "natural time scales" differ, possibly by orders of magnitude. This, again, is quite frequent when we pair a process that evolves in the order of seconds or minutes, or even days or months (such as a chemical plant, or a hydroelectric power plant) with a controller based on digital electronic devices. In principle, if we assume a unique continuous time model, say the reals, the problem is reduced to a, possibly cumbersome, change of time unit.[2]

---

[2]Notice, however, that in very special cases the different time units could be incommensurable. In fact, even if in practice this circumstance may seldom arise, after all the two main

However, if discrete domains are adopted, subtle semantic issues related to the approximation of the coarser time unit may arise. Consider, for instance, the sentences "every month, if an employee works, then she gets her salary" and "whenever an employee is assigned a job, this job should be completed within three days". They may both be part of the same specification of an office system, so an admissible behavior for the office system must satisfy both sentences. It would be natural to assume a discrete temporal domain in which the time unit is the day, which is of finer granularity than the month. However, it is clear that stating that "every month, if an employee works, then she gets her salary" is not the same as "every day, if an employee works, then she gets her salary". In fact, in general, working for one month means that one works for 22 days of the month, whereas getting a monthly salary means that there is one day when one gets the salary for the month. Hence, a simple change in the time unit (from months to days) in this case does not achieve the desired effect.

As a further example, suppose that the sentence "this job has to be finished within 3 days from now" is stated at 4 PM of a given day: should this be interpreted as "this job has to be finished within $3 \times 24 \times 60 \times 60$ seconds from now", or "this job has to be finished before midnight of the third day after today"? Both interpretations may be adopted depending on the context of the claim.

An approach that deals rigorously with different time granularities is presented in Section 5.2.1 when discussing temporal logics.

3.2 Ordering vs. Metric

Another central issue is whether a formalism permits the expression of metric constraints on time, or, equivalently, of constraints that exploit the metric structure of the underlying time model (if it has any).

A domain (a time domain, in our case) has a *metric* when a notion of *distance* is defined on it (that is, when a nonnegative *measure* $\mathrm{d}(t_1, t_2) \geq 0$ is associated with any two points t_1, t_2 of the domain).

As mentioned above, typical choices for the time domain are the usual discrete and dense numeric sets, that is $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}$. All these domains have a "natural" metric defined on them, which corresponds to simply taking the distance[3] between two points: $\mathrm{d}(t_1, t_2) = |t_1 - t_2|$.[4]

Notice, however, that although all common choices for the time domains possess a metric, we focus on whether the *language* in which the system is described permits descriptions using the same form of metric information as that embedded in the underlying time domain. For instance, some languages allow the user to state that an event p (e.g., "push button") must temporally precede another event q (e.g., "take picture"), but do not include constructs to

---

units offered by nature, the day and the year, are incommensurable.

[3]Technically, this is called the Euclidean distance.

[4]We focus our attention here on temporal domains $\mathbb{T}$ that are totally ordered; although partially-ordered sets may be considered as time domains (see Section 3.6), they have not attracted as much research activity as totally ordered domains.

specify how long it takes between the occurrence of p and that of q; hence, they cannot distinguish the case in which the delay between p and q is 1 time unit from the case in which the delay is 100 time units. Thus, whenever the language does not allow users to state metric constraints, it is possible to express only information about the relative *ordering* of phenomena ("q occurs after p"), but not about their distance ("q occurs 100 time units after p"). In this case, we say that the language has a purely *qualitative* notion of time, as opposed to allowing *quantitative* constraints, which are expressible with metric languages.

Parallel systems have been defined [Wir77] as those where the correctness of the computation depends only on the *relative ordering* of computational steps, irrespective of the absolute distance between them. Reactive systems can often be modeled as parallel systems, where the system evolves concurrently with the environment. Therefore, for the formal description of such systems a purely qualitative language is sufficient. On the contrary, real-time systems are those whose correctness depends on the *time distance* among events as well. Hence, a complete description of such systems requires a language in which metric constraints can be expressed. In this vein, the research in the field of formal languages for system description has evolved from dealing with purely qualitative models to the more difficult task of providing the user with the possibility of expressing and reasoning about metric constraints.

For instance, consider the two sequences b_1, b_2 of events p, q, where exactly one event per time step occurs:

- $b_1 = \mathsf{pqpqpq} \cdots$

- $b_2 = \mathsf{ppqqppqq} \cdots$

b_1 and b_2 share all the qualitative properties expressible without using any metric operator. For instance "every occurrence of p is eventually followed by an occurrence of q" is an example of qualitative property that holds for both behaviors, whereas "p occurs in every instant" is another qualitative property, that instead is false for both behaviors. If referring to metric properties is allowed, one can instead discriminate between b_1 and b_2, for example through the property "every occurrence of q is followed by another occurrence of q after two time steps", which holds for b_1 but not for b_2.

Some authors have introduced a notion of equivalence between behaviors that captures the properties expressed by qualitative formulas. In particular Lamport [Lam83] first proposed the notion of *invariance under stuttering*. Whenever a (discrete time) behavior b_3 can be obtained from another behavior b_4 by adding and removing "stuttering steps" (i.e., pairs of identical states on adjacent time steps), we say that b_3 and b_4 are stutter-equivalent. For instance, behaviors b_1 and b_2 outlined above are stutter-equivalent. Then, the equivalence classes induced by this equivalence relation precisely identify classes of properties that share identical qualitative properties. Note that stutter invariance is defined for discrete time models only.

3.3 Linear vs. Branching Time Models

The terms *linear* and *branching* refer to the structures on which a formal language is interpreted: *linear*-time formalisms are interpreted over *linear* sequences of states, whereas *branching*-time formalisms are interpreted over *trees* of states. In other words, each description of a system adopting a linear notion of time refers to (a set of) linear behaviors, where the future evolution from a given state at a given time is always unique. Conversely, a branching-time interpretation refers to behaviors that are structured in trees, where each "present state" may evolve into different "possible futures". For instance, assuming discrete time, Figure 5 pictures a linear sequence of states and a tree of states over six time instants.


Figure 5: A linear and a branching time model.

A linear behavior is a special case of a tree. Conversely, a tree might be thought of as a set of linear behaviors that share common prefixes (i.e., that are prefix-closed); this notion is captured formally by the notion of *fusion closure* [AH92b]. Thus, linear and branching models can be put on a common ground and compared. This has been done extensively in the literature.

Linear or branching semantic structures are then matched, in the formal languages, by corresponding syntactic elements that allow us to express properties about the specific features of the interpretation. This applies in principle to all formal languages, but it has been the focus of logic languages especially, and temporal logics in particular. Thus, a linear-time temporal logic is interpreted over linear structures, and is capable of expressing properties of behaviors with unique futures, such as "if event p happens, then event q will happen eventually in the future". On the other hand, branching-time temporal logics are interpreted over tree structures and allow users to state properties of branching futures, such as "if event p happens at some time t, then there is some possible future where event q holds". We discuss this in greater depth in our consideration of temporal logics (Section 5.2.1); for general reference we cite the classic works by Lamport [Lam80], Emerson and Halpern [EH86], Emerson [Eme90], and Alur and Henzinger [AH92b] — the last focusing on real-time models.

13

Finally, we mention that it is possible to have semantic structures that are also *branching in the past* [Koy92], which is where different pasts merge into a single present. However, in practice, branching-in-the-past models are relatively rare, so we will not deal with them in the remainder of this article.

Determinism vs. Nondeterminism

Linear and branching time are features of the languages and structures on which they are interpreted, whereas the notions of determinism and nondeterminism are attributes of the *systems* being modeled or analyzed. More precisely, let us consider systems where a notion of *input* is defined: one such system evolves over time by reading its input and changing the current state accordingly. Whenever the future state of the system is *uniquely* determined by its current state and input values, then we call the system *deterministic*. For instance, a light button is a deterministic system where pressing the button (input) when the light is in state *off* yields the unique possible future state of light *on*. Notice that, for a given input sequence, the behavior of a deterministic system is uniquely determined by its initial state. Conversely, systems that can evolve to different future states from the same present state and the same input by making arbitrary "choices" are called *nondeterministic*. For example, a resource arbiter might be a nondeterministic system that responds to two requests happening at the same time by "choosing" arbitrarily to whom to grant the resource first. The Ada programming language [BB94] embeds such a nondeterminism in its syntax and semantics.

In linear-time models the future of any instant is unique, whereas in branching-time models each instant branches into different possible futures; then, there is a natural coupling between deterministic systems and linear models on one side, and on the other side nondeterministic systems and branching models, where all possible "choices" are mapped at some time to branches in the tree. Often, however, linear-time models are still preferred even for nondeterministic systems for their intuitiveness and simplicity. In the discussion of Petri nets (Section 5.1.2) we see an example of linear time domains expressing the semantics of nondeterministic formalisms.

3.4 Implicit vs. Explicit Time Reference

Some languages allow, or impose on, the user to make explicit reference to temporal items (attributes or entities of "type time"), whereas other formalisms, though enabling reasoning about temporal system properties, leave all or some references to time-related properties (occurrences of events, durations of states or actions) implicit in the adopted notation.

To illustrate, at one extreme consider pure first-order predicate calculus to specify system behavior and its properties. In such a case we could use explicit terms ranging over the time domain and build any formula, possibly with suitable quantifiers, involving such terms. We could then could express properties such as "if event p occurs at instant t, then q occurs at some instant

t' no later than k time units after t". At the other extreme, classic temporal logic [Kam68], despite its name, does not offer users the possibility to explicitly mention any temporal quantity in its formulas, but aims at expressing temporal properties by referring to an implicit "current time" and to the ordering of events with respect to it; for example, it has operators through which it is possible to represent properties such as "if event p occurs now, then sometime in the future event q will follow".

Several formalisms adopt a kind of intermediate approach. For instance, many types of abstract machines allow the user to specify explicitly, say, the duration of an activity with implicit reference to its starting time (e.g., Statecharts, discussed in Section 5.1.1, and Petri nets, discussed in Section 5.1.2). Similarly, some languages inspired by temporal logic (e.g., MTL, presented in Section 5.2.1) keep its basic approach of referring any formula to an implicit current instant (the *now* time), but allow the user to explicitly express a time distance with respect to it. Typical examples of such formulas may express properties such as "if event p occurs now then event q will follow in the future within t time units".

In general, using implicit reference to time instants — in particular the use of an implicit *now* — is quite natural and allows for compact formalizations when modeling and expressing properties of so-called "time-invariant systems", which are the majority of real-life systems. In most cases, in fact, the system behavior is the same if the initial state of and the input supplied to the system are the same, even if the same computation occurs at different times. Hence the resulting behaviors are simply a temporal translation of one another, and in such cases, explicitly expressing where the *now* is located in the time axis is superfluous.

3.5 The Time Advancement Problem

The problem of time advancement arises whenever the model of a timed system exhibits behaviors that do not progress past some instant. Such behaviors do not correspond to any physical "real" phenomena; they may be the consequence of some incompleteness and inconsistency in the formalization of the system, and thus must be ruled out.

The simplest manifestation of the time advancement problem arises with models that allow transitions to occur in a null time. For instance, several automata-based formalisms such as Statecharts and timed versions of Petri nets adopt this abstract notion (see Section 5.1.1). Although a truly instantaneous action is physically unfeasible, it is nonetheless a very useful abstraction for events that take an amount of time which is negligible with respect to the overall dynamics of the system [BB06]. For example, pushing a button is an action whose actual duration can usually be ignored and thus can be represented abstractly as a zero-time event. When zero-time transitions are allowed, an infinite number of such transitions may accumulate in an arbitrarily small interval to the left of a given time instant, thus modeling a fictitious infinite computation where time does not advance at all. Behaviors where time does not advance

are usually called *Zeno* behaviors, from the ancient philosopher Zeno of Elea[5] and his paradoxes on time advancement (the term was coined by Abadi and Lamport [AL94]). Notice that, from a rigorous point of view, even the notion of behavior as a function — whose domain is time and whose range is system state — is ill-defined if zero-time transitions are allowed, since the consequences of a transition that takes zero time to occur is that the system is both at the source state and at the target state of the transition in the same instant.

Even if actions (i.e., state transformations) are noninstantaneous, it is still possible for Zeno behaviors to occur if time advances only by *arbitrarily small* amounts. Consider, for instance, a system that delivers an unbounded sequence of events p_k, for $k \in \mathbb{N}$; each event p_k happens exactly t_k time units after the previous one (i.e., p_{k-1}). If the sum of the relative times (that is, the sum $\Sigma_k t_k$ of the time distances between consecutive events) converges to a finite limit t, then the absolute time never surpasses t; in other words, *time stops* at t, while an infinite number of events occur between any t_k and t. Such behaviors allow an infinite number of events to occur within a finite time.

Even when we consider continuous-valued time-dependent functions of time that vary smoothly, we may encounter Zeno behaviors. Take, for instance, the real-valued function of time $b(t) = \exp\left(-1/t^2\right)\sin(1/t)$; $b(t)$ is very smooth, as it possesses continuous derivatives of all orders. Nonetheless, its sign changes an infinite number of times in any interval containing the origin; therefore a natural notion such as "the next instant at which the sign of b changes" is not defined at time 0, and, consequently, we cannot describe the system by relating its behavior to such — otherwise well-defined — notions. Indeed, as explained precisely in Section 5.2 when discussing temporal logics, non-Zenoness can be mathematically characterized by the notion of *analyticity*, which is even stronger than infinite derivability.

The following remarks are to some extent related to the problem of time advancement, and might help a deeper understanding thereof.

- Some formal systems possess "Zeno-like" behaviors, where the distance between consecutive events gets indefinitely smaller, even if time progresses (these behaviors have been called "Berkeley" in [FPR08a], from the philosopher George Berkeley[6] and his investigations arguing against the notion of infinitesimal). These systems cannot be controlled by digital controllers operating with a fixed sampling rate such as in [CHR02], since in this case their behaviors cannot be suitably discretized [FR06, FPR08a].

- Some well-known problems of — possibly — concurrent computation such as *termination*, *deadlocks*, and *fairness* [Fra86] can be considered as *dual* problems to time advancement. In fact, they concern situations where some processes fail to advance their *states*, while time keeps on flowing. Examples of these problems and their solutions are discussed with reference to a variety of formalisms introduced in Section 5.

---

[5]Circa 490–425 BC.
[6]Kilkenny, 1685–Oxford, 1753.

We can classify solutions to the time advancement problem into two categories: *a priori* and *a posteriori* methods. In *a priori* methods, the syntax or the semantics of the formal notation is restricted beforehand, in order to guarantee that the model of any system described with it will be exempt from time advancement problems. For instance, in some notations zero-time events are simply forbidden, or only finite sequences of them are allowed.

On the contrary, *a posteriori* methods do not deal with time advancement issues until *after* the system specification has been built; then, it is analyzed against a formal definition of time advancement in order to check that all of its actual behaviors do not incur into the time advancement problem. An *a posteriori* method may be particularly useful to spot possible risks in the behavior of the real system. For instance, in some cases oscillations exhibited by the mathematical model with a frequency that goes to infinity within a finite time interval, such as in the example above, may be the symptom of some instability in the modeled physical system, just in the same way as a physical quantity — say, a temperature or a pressure — that, in the mathematical model, tends to infinity within a finite time is the symptom of the risk of serious failure in the real system.

3.6 Concurrency and Composition

Most real *systems* — as the term itself suggests — are complex enough that it is useful, if not outright unavoidable, to model, analyze, and synthesize them as the composition of several subsystems. Such a composition/decomposition process may be iterated until each component is simple enough so that it can be analyzed in isolation.

Composition/decomposition, also referred to as *modularization*, is one of the most general and powerful design principles in any field of engineering. In particular, in the case of — sequential — software design it produced a rich collection of techniques and language constructs, from subprograms to abstract data types.

The state of the art is definitely less mature when we come to the composition of concurrent activities. In fact, it is not surprising that very few programming languages deal explicitly with concurrency. It is well-known that the main issue with the modeling of concurrency is the *synchronization* of activities (for which a plethora of more or less equivalent constructs are used in the literature: processes, tasks, threads, etc.) when they have to access shared resources or exchange messages.

The problem becomes even more intricate when the various activities are heterogeneous in nature. For instance, they may involve "environment activities" such as a plant or a vehicle to be controlled, and monitoring and control activities implemented through some hardware and software components. In such cases the time reference can be implicit for some activities, explicit for others; also, the system model might include parts in which time is represented simply as an ordering of events and parts that are described through a metric notion of time; finally, it might even be the case that different components

17

are modeled by means of different time domains (discrete or continuous), thus producing hybrid systems.

Next, a basic classification of the approaches dealing with the composition of concurrent units is provided.

Synchronous vs. Asynchronous Composition

When composing concurrent modules there are two foremost ways of relating their temporal evolution: these are called synchronous and asynchronous composition.

Synchronous composition constraints state changes of the various units to occur at the very same time or at time instants that are strictly and rigidly related. Notice that synchronous composition is naturally paired with a discrete time domain, but meaningful exceptions may occur where the global system is synchronized over a continuous time domain.

Conversely, in an *asynchronous* composition of parallel units, each activity can progress at a speed relatively unrelated with others; in principle there is no need to know in which state each unit is at every instant; in some cases this is even impossible: for instance, if we are dealing with a system that is geographically distributed over a wide area and the dynamics of some component evolves at a speed that is of the same order of magnitude as the light speed (more precisely, the state of a given component changes in a time that is shorter than the time needed to send information about the component's state to other components).

A similar situation occurs in totally different realms, such as the world-wide stock market. There, the differences in local times between locations all over the world make it impossible to define certain states about the global market, such as when it is "closed".

For asynchronous systems, interaction between different components occurs only at a few "meeting points" according to suitably specified rules. For instance, the Ada programming language [BB94] introduces the notion of *rendez vous* between asynchronous tasks: a task owning a resource waits to grant it until it receives a request thereof; symmetrically a task that needs to access the resources raises a request (an *entry call*) and waits until the owner is ready to accept it. When both conditions are verified (an entry call is issued and the owner is ready to accept it), the rendez vous occurs, that is, the two tasks are synchronized. At the end of the entry execution by the owner, the tasks split again and continue their asynchronous execution.

Many formalisms exist in the literature that aim at modeling some kind of asynchronous composition. Among these, Petri nets exhibit similarities with the above informal description of Ada's task system.

Not surprisingly, however, precisely formalizing the semantics of asynchronous composition is somewhat more complex than the synchronous one, and several approaches have been proposed in the literature. We examine some of them in Section 5.

3.7 Analysis and Verification Issues

A fundamental feature of a formal model is its amenability to analysis; namely, we can probe the model of a system to be sure that it ensures certain desired features. In a widespread paradigm [GJM02, Som04], we call *specification* the model under analysis, and *requirements* the properties that the specification model must exhibit. The task of ensuring that a given specification satisfies a set of requirements is called *verification*. Although this survey does not focus on verification aspects, we will occasionally deal with some related notions.

Expressiveness

A fundamental criterion according to which formal languages can be classified is their *expressiveness*, that is, the possibility of characterizing extensive classes of properties. Informally, a language is more expressive than another one if it allows the designer to write sentences that can more finely and accurately partition the set of behaviors into those that satisfy or fail to satisfy the property expressed by the sentence itself. Note that the expressiveness relation between languages is a partial order, as there are pairs of formal languages whose expressive power is incomparable: for each language there exist properties that can be expressed only with the other language. Conversely, there exist formalisms whose expressive powers coincide; in such cases they are equivalent in that they can express the very same properties. Expressiveness deals only with the logical possibility of expressing properties; this feature is totally different from other — somewhat subjective, but nonetheless very relevant — characterizations such as conciseness, readability, naturalness, and ease of use.

Decidability and Complexity

Although in principle we might prefer the "most expressive" formalism, in order not to be restrained in what it be expressed, there is a fundamental trade-off between expressiveness and another important characteristic of a formal notation, namely, its *decidability*. A certain property is *decidable* for a formal language if there exists an algorithmic procedure that is capable of determining, for any specification written in that language, whether the property holds or not in the model. Therefore, the verification of decidable properties can be — at least in principle — a totally automated process. The trade-off between expressiveness and decidability arises because, when we increase the expressiveness of the language, we may lose decidability, and thus have to resort to semi-automated or manual methods for verification, or adopt *partial* verification techniques such as testing and simulation. Here the term *partial* refers to the fact that the analysis conducted with these techniques provides results that concern only a subset of all possible behaviors of the model under analysis.

While decidability is just a yes/no property, *complexity* analysis provides, in the case when a given property is decidable, a measure of the computational effort that is required by an algorithm that decides whether the property holds or not for a model. The computational effort is typically measured in terms

of the amount of memory or time required to perform the computation, as a function of the length of the input (that is, the size of the sentence that states it in the chosen formal language; see also Section 4.3).

Analysis and Verification Techniques

There exist two large families of verification techniques: those based on *exhaustive enumeration* procedures and those based on *syntactic transformations* like deduction or rewriting, typically in the context of some axiomatic description. Although broad, these two classes do not cover, by any means, all the spectrum of verification algorithms, which comprises very different techniques and methods; here, however, we limit ourselves to sketching a minimal definition of these two basic techniques.

Exhaustive enumeration techniques are mostly automated, and are based on exploration of graphs or other structures representing an operational model of the system, or the space of all possible interpretations for the sentence expressing the required property.

Techniques based on *syntactic transformations* typically address the verification problem by means of logic deduction [Men97]. Therefore, usually both the specification and the requirements are in descriptive form, and the verification consists of successive applications of some deduction schemes until the requirements are shown to be a logical consequence of the system specification.

4   Historical Overview

In the rest of this article, in the light of the categories outlined in Section 3, we survey and compare a wide range of time models that have been used to describe computational aspects of systems.

This section presents an overview of the "traditional" models that first tackled the problem of time modeling, whereas Section 5 discusses some more "modern" formalisms. As stated in Section 2, we start from the description of *formalisms*, but we will ultimately focus on their *semantics* and, therefore, on what kind of *temporal modeling* they allow.

Any model that aims at describing the "dynamics" of phenomena, or a "computation" will, in most cases, have some notion of time. The modeling languages that have been used from the outset to describe "systems", be they physical (e.g., moving objects, fluids, electric circuits), logical (e.g., algorithms), or even social or economic ones (e.g., administrations) are no exception and incorporate a more or less abstract idea of time.

This section presents the relevant features of the notion of time as traditionally used in three major areas of science and engineering: control theory (Section 4.1), electronics (Section 4.2) and computer science (Section 4.3). As the traditional modeling languages used in these disciplines have been widely studied and are well understood, we will only sketch their (well-known) main features; we will nonetheless pay particular attention to the defining character-

istics of the notion of time employed in these languages, and highlight its salient points.

4.1 Traditional Dynamical Systems

A common way used to describe systems for control purposes in various engineering disciplines (mechanical, aerospace, chemical, electrical, etc.) is through the so-called state-space representation [Kha95, SP05].

The state-space representation is based on three key elements: a vector $\mathbf{x}$ of *state variables*, a vector $\mathbf{u}$ of *input variables*, and a vector $\mathbf{y}$ of *output variables*. $\mathbf{x}$, $\mathbf{u}$, and $\mathbf{y}$ are all *explicit functions of time*, hence their values depend on the time at which they are evaluated and they are usually represented as $\mathbf{x}(t)$, $\mathbf{u}(t)$, and $\mathbf{y}(t)$.[7]

In the state-space representation the temporal domain is usually either continuous (e.g., $\mathbb{R}$), or discrete (e.g., $\mathbb{Z}$). Depending on whether the temporal domain is $\mathbb{R}$ or $\mathbb{Z}$, the relationship between $\mathbf{x}$ and $\mathbf{u}$ is often expressed through differential or difference equations, respectively, e.g., in the following form:

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t), t)$$
$$\mathbf{x}(k+1) = f(\mathbf{x}(k), \mathbf{u}(k), k) \tag{1}$$

where $t \in \mathbb{R}$ and $k \in \mathbb{Z}$ (the relationship between $\mathbf{y}$ and the state and input variables is instead purely algebraic in the form $\mathbf{y}(t) = g(\mathbf{x}(t), \mathbf{u}(t), t)$.

Given an initial condition $\mathbf{x}(0)$, and fixed a function $\mathbf{u}(t)$, all functions $\mathbf{x}(t)$ (or $\mathbf{x}(k)$) if time is discrete) that are *solutions* to the equations (1) represent the possible system behaviors. Notice that suitable constraints on the formalization of the system's dynamics are defined so that the derived behaviors satisfy some natural causality principles. For instance, the form of equations (1) must ensure that the state at time t depends only on the initial state and on the value of the input in the interval $[0, t]$ (the future cannot modify the past).

Also, systems described through state-space equations are usually *deterministic* (see Section 3.3), since the evolution of the state $\mathbf{x}(t)$ is unique for a fixed input signal $\mathbf{u}(t)$ (and initial condition $\mathbf{x}(0)$).[8] Therefore, dynamical system models typically assume a linear time model (see also the discussion in Section 3.3).

Moreover, time is typically treated quantitatively in these models, as the metric structure of the time domains $\mathbb{R}$ or $\mathbb{Z}$ is exploited.

Notice also that often the first equation of (1) takes a simplified form:

$$\dot{\mathbf{x}} \quad = \quad f(\mathbf{x}, \mathbf{u})$$

---

[7]Another classic way of representing a dynamical system is through its *transfer function*, which describes the input/output relationship of the system; unlike the state-space representation, the transfer function uses an *implicit*, rather than explicit, notion of time. Despite its popularity and extensive use in the field of control theory, the transfer function has little interest in the modeling of computation, so we do not delve any further in its analysis.

[8]Notice that for a dynamical system described by equations such as (1) to be nondeterministic, the solution of the equation should be non-unique; this is usually ruled out by suitable hypotheses on the f function [Kha95].

where the *time* variable does not occur explicitly but is *implicit* in the fact that **x** and **u** are functions thereof. The time variable, of course, occurs explicitly in the solution of the equation. This is typical of time-invariant systems i.e., those systems that behave identically if the same "experiment" is translated along the time axis by a finite constant.

A typical example of continuous-time system is the electric circuit of Figure 2. A less common instance of discrete-time system is provided in the next example.

Example 4 (Monodimensional Cellular Automata). Let us consider a discrete-time family of dynamical systems called *cellular automata*, where $\mathbb{T} = \mathbb{N}$. More precisely, we consider the following instance, named *rule 110* by Wolfram [Wol94]. The state domain is a bi-infinite string of binary values $\mathbf{s}(k) = \ldots s_{i-2}(k)s_{i-1}(k)s_i(k)s_{i+1}(k)s_{i+2}(k) \ldots \in \{0,1\}^{2\omega}$, and the output coincides with the whole state. The system is closed, since it has no input, and its evolution is entirely determined by its initial state $s_i(0)$ ($i \in \mathbb{Z}$).

The dynamics is defined by the following equation, which determines the update of the state according to its value at the previous instant (starting from instant 1).

$$s_i(k+1) = \begin{cases} 1 & \text{if } s_{i-1}(k)s_i(k)s_{i+1}(k) \in \{110, 101, 011, 010, 001\} \\ 0 & \text{otherwise} \end{cases}$$

Despite the simplicity of the update rule, the dynamics of such a system is highly complex; in particular it has been shown to be capable of universal computation [Coo04].

Let us now discuss the main advantages in modeling an evolving process by means of a dynamical system. In doing so, perhaps we shift the point of view of system analysis from a "control attitude" towards a "computing attitude".

The foremost advantage of dynamical system models is probably that they borrow directly from the models used in physics. Therefore, they are capable of describing very general and diverse dynamic behaviors, with the powerful long-standing mathematical tools of differential (or difference) calculus. In particular, very different heterogeneous time models can be described within the same framework.[9] In a sense, many other formalisms for time modeling can be seen as a specialization of dynamical systems and can be reformulated in terms of state-space equations, including more computationally-oriented formalisms such as finite state automata and Turing machines.

The main limitations of the dynamical system models in describing timed systems lie in their being "too detailed" for some purposes. Being intrinsically operational and deterministic in most cases, such models provide complete descriptions of a system behavior, but are unsuitable for partial specifications or

---

[9]The recent literature of control theory also deals with hybrid systems where discrete and continuous time domains are integrated in the same system formalization [vS00, Ant00, BBM98].

very high-level descriptions, which are instead a necessary feature in the early phases of development of a system (e.g., the requirements engineering phase in the development of software).

Moreover, since the time domain is usually a totally ordered metric set, dynamical systems are unsuitable for describing distributed systems where a notion of "global time" cannot be defined, and systems where the exact timing requirements are unspecified or unimportant. Some models that we will describe in the next sections try to overcome these limits by introducing suitable abstractions.

4.2   The Hardware View

One field in which the modeling of time has always been a crucial issue is (digital) electronic circuits design.

The key modeling issue that must be addressed in describing digital devices is the need to have *different abstraction levels* in the description of the same system. More exactly, we typically have two "views" of a digital component. One is the *micro view*, which is nearest to a physical description of the component. The other is the *macro view*, where most lower-level details are abstracted away.

The *micro view* is a low-level description of a digital component, where the basic physical quantities are modeled explicitly. System description usually partitions the relevant items into input, output, and state values. All of them represent physical quantities that vary continuously over time. Thus, the time domain is *continuous*, and so is the state domain. More precisely, since we usually define an initialization condition, the temporal domain is usually *bounded* on the left (i.e., $\mathbb{R}_{\geq 0}$). Conversely, the state domain is often, but not always, restricted to a bounded subset $[L, U]$ of the whole domain $\mathbb{R}$ (in many electronic circuits, for example, voltages vary from a lower bound of approximately 0V to an upper bound of approximately 5V).

Similarly to the case of time-invariant dynamical systems, time is generally *implicit* in formalisms adopting the micro view. It is also *metric* — as it is always the case in describing directly physical quantities — and fully *asynchronous*, so that inputs may change at any instant of time, and outputs and states react to the changes in the inputs at any instant of time.

A simple operational formalism used to describe systems at the micro view is that of *logic gates* [KB04], which can then be used to represent more complex digital components, with memory capabilities, such as *flip-flops* and *sequential machines*.

Figure 6 shows an example of behavior of a sequential machine with two inputs i_0 and i_1, one output o, and two state values m_0 and m_1.

The figure highlights the salient features of the modeling of time at the micro (physical) level: continuity of both time and state, and asynchrony (for example memory signals m_0 and m_1 can change their values at different time instants).

More precisely, Figure 6 pictures a possible evolution of the state (i.e., the pair $\langle m_0, m_1 \rangle$) and of the output (i.e., signal o) of the sequential machine with

Figure 6: A behavior of a sequential machine.

respect to its input (i.e., the pair $\langle i_0, i_1 \rangle$). For example, it shows that if all four signals i_0, i_1, m_0, m_1 are "low" (e.g., at time t_0), then the pair $\langle m_0, m_1 \rangle$ remains "low"; however, if both input signals are "high" and the state is "low" (e.g., at time t_1), m_1 becomes "high" after a certain delay. The output is also related to the state, in that o is "high" when both m_0 and m_1 are "high" (in fact, o becomes "high" a little after m_0 and m_1 both become "high", as shown at time t_3 in the figure). Notice how the reaction to a change in the values of the input signals is not instantaneous, but takes a non-null amount of time (a *propagation delay*), which depends on the propagation delays of the logic gates composing the sequential machine.

As the description above suggests, the micro view of digital circuits, being close to the "physical" representation, is very detailed (e.g., it takes into account the transient state that occurs after variation in the inputs). However, if one is able to guarantee that the circuit will eventually reach a stable state after a variation of the inputs, and that the duration of the transient state is short with respect to the rate with which input variations occur, it is possible to abstract away the inner workings of the digital circuits, and focus on the effects of a change in the inputs on the machine state, instead. In addition, it is common practice to represent the "high" and "low" values of signals in an abstract way, usually as the binary values 1 (for "high") and 0 (for "low"). Then, we can associate a sequential machine with a logic function that describes the evolution of only the stable states. Table 1 represents such a logic function where we associate a letter to every possible stable configuration of the inputs (column header) and the memory (row header), while the output is instead simply defined to be 1 if and only if the memory has the stable value 11. A blank cell in the table denotes an undefined ("don't care") behavior for the corresponding pair of current state and current input. Then, the evolution in Figure 6 is compatible with the system specification introduced by Table 1.

24

	$a(00)$	$b(01)$	$c(11)$	$d(10)$
$A(00)$	00	00	01	10
$B(01)$		10	11	01
$C(11)$	00		10	11
$D(10)$	00	10	10	11

Table 1: A tabular description of the behavior of a sequential machine.

Notice that by applying the above abstraction we discretized the state domain and assumed zero-time transitions. However, in the behavior of Figure 6 the inputs i_0 and i_1 vary too quickly to guarantee that the component realizes the input/output function described by the table above. For example, when at instant t_2 both m_0 and m_1 become 1, memory signal m_1 does not have the time to become 0 (as stated in Table 1) before input i_1 changes anew. In addition, the output does not reach a stable state (and become 1) before state m_1 switches to 0. Thus, the abstraction of zero-time transition was not totally correct.

As the example suggests, full asynchrony in sequential machines poses several problems both at the modeling and at the implementation level. A very common way to avoid these problems, thus simplifying the design and implementation of digital components, is to synchronize the evolution of the components through a *clock* (i.e., a physical signal that forces variations of other signals to occur only at its edges).

The benefits of the introduction of a clock are twofold: on the one hand, a clock rules out "degenerate behaviors", in which signal stability is never reached [KB04]; on the other hand, it permits a higher-level view of the digital circuit, which we call the *macro view*.

In the macro view, not only physical quantities are represented symbolically, as a combination of binary values; such values, in turn, are observed only when they have reached stable values. The time domain becomes discrete, too. In fact inputs are read only at periodic instants of time, while the state and outputs are simultaneously (and instantaneously) updated. Thus, their observation is *synchronized* with a clock that beats the time periodically. Since we disregard any transient state, time is now a *discrete* domain. In practice, we adopt the naturals $\mathbb{N}$ as time domain, whose origin matches the initialization instant of the system.

Typical languages that adopt "macro" time models are those that belong to the large family of abstract state machines [Sip05, HMU00, MG87]. More precisely, the well-known Moore machines [Moo56] and Mealy machines [Mea55] have been used for decades to model the dynamics of digital components. For example, the Moore machine of Figure 7 represents the dynamics of the sequential machine implementing the logic function defined by Table 1. Every transition in the Moore machine corresponds to the elapsing of a clock interval; thus, the model abstracts away from all physical details, and focuses on the evolution of the system at clock ticks.

Figure 7: A Moore machine.

We will discuss abstract state machines and their notion of time in more detail in Section 5.1.1.

4.3   The Software View

As mentioned above, abstract state machines such as the Moore machine of Figure 7 give a representation of digital components that is more "computational-oriented" and abstract than the "physics-oriented" one of logic gates.

Traditionally, the software community has adopted a view of the evolution of programs over time that is yet more abstract.

In the most basic view of computation, time is not modeled at all. In fact, a software application implements a function of the inputs to the outputs. Therefore, the whole computational process is considered atomic, and time is absent from functional formalization. In other words, behaviors have no temporal characteristics in this basic view, but they simply represent input/output pairs for some computable function. An example of a formal language adopting such black-box view of computation is that of *recursive functions*, at the roots of the *theory of computation* [Odi99, Rog87, BL74].

A refinement of this very abstract way of modeling software would keep track not only of the inputs and outputs of some computation, but also of the whole sequence of discrete steps the computing device undergoes during the computation (i.e., of the *algorithm* realized by the computation). More precisely, the actual time length of each step is disregarded, assigning uniformly a unit length to each of them; this corresponds to choosing the naturals ℕ as time domain. Therefore, time is discrete and bounded on the left: the initial time 0 represents the time at which the computation starts. The time measure represents the number of elementary computational steps that are performed through the computation. Notice that no form of concurrency is allowed in these computational models, which are strictly *sequential*, that is each step is followed uniquely by its successor (if any).

Turing Machines [Pap94, Sip05, HMU00, MG87] are a classic formalism to describe computations (i.e., algorithms). For example, the Turing machine of

Figure 8 describes an algorithm to compute the successor function for a binary input (stored with the least significant bit on the left, that is in a "little-endian" fashion).


Figure 8: A Turing machine computing the successor function. $\triangleright$ denotes the origin of the Turing machine tape, $\square$ denotes the blank symbol; a double circle marks a halting state; in every transition, $I/O, M$ denotes the symbol read on the tape upon taking the transition (I), the symbol (O) written on the tape in place of I, and the way (M) in which the tape head is moved (l for "left", s for "stay", and r for "right").

For a given Turing machine M computing a function f, or any other abstract machine for which it is assumed that an elementary transition takes a time unit to execute, by associating the number of steps from time 0 until the reaching of a halting state — if ever — we may define a *time complexity* function $T_M(n)$, whose argument n represents the *size* of the data input to f, and whose value is the maximum number of steps required to complete the computation of f when input data has size n.[10] For example, the computational complexity $T_{\text{succ}}(n)$ of the Turing machine of Figure 8 is proportional to the length n of the input string.

In the software view the functional behavior of a computation is normally separated from its timed behavior. Indeed, while the functional behavior is studied without taking time into account, the modeling of the timed behavior focuses only on the number of steps required to complete the computation. In other words, functional correctness and time complexity analysis are usually separated and adopt different techniques.

In some sense the software view of time models constitutes a further abstraction of the macro hardware view. In particular, the adoption of a discrete time domain reflects the fact that the hardware is what actually performs the computations formalized by means of a software model. Therefore, all the abstract automata that are used for the macro modeling of hardware devices are also commonly considered models of computation.

---

[10] As a particular case, if M's computation never reaches a halting state we conventionally define $T_M(n) = \infty$.

5 Temporal Models in Modern Theory and Practice

The growing importance and pervasiveness of computer systems has required the introduction of new, richer, and more expressive temporal models, fostering their evolution from the basic "historical" models of the previous section. This evolution has inevitably modified the boundaries between the traditional ways of modeling time, often making them fuzzy. In particular, this happened with heterogeneous systems, which require the combination of different abstractions within the same model.

This section shows how the aforementioned models have been refined and adapted in order to meet more specific and advanced specification needs. These are particularly prominent in some classes of systems, such as hybrid, critical, and real-time systems [HM96]. As we discussed above in Section 1, these categories are independent but with large areas of overlap.

Keywords	Dimension	Section
discrete, dense, continuous, granularity	*Discrete vs. Dense*	3.1
qualitative, quantitative, metric(s)	*Ordering vs. Metric*	3.2
linear, branching, (non)deterministic	*Linear vs. Branching*	3.3
implicit(ly), explicit(ly)	*Implicit vs. Explicit*	3.4
(non)-Zeno, fairness, deadlock(ed)	*Time Advancement*	3.5
composing, composition, concurrency, synchrony, synchronous(ly), asynchronous(ly)	*Concurrency and Composition*	3.6
analysis, tool(set), verification, decision procedure	*Analysis and Verification*	3.7

Table 2: Keyword references to the "Dimensions" of Section 3.

As in the historical overview of Section 4, the main features of the models presented in this section are discussed along the dimensions introduced in Section 3. Such dimensions, however, have different relevance for different formalisms; in some cases a dimension may even be unrelated with some formalism. For this reason we avoid a presentation in the style of a systematic "tabular" cross-reference < Formalism/dimension >; rather, to help the reader match the features of a formalism with the coordinates of Section 3, we highlight the portions of the text where a certain dimension is specifically discussed by graphically emphasizing (in small caps) some related keywords. The correspondence between keywords and dimensions is shown in Table 2. Also, for the sake of conciseness, we do not repeat features of a *derived* formalism that are inherited unaffected from the "parent" notation.

The Computer- and System-Centric Views

As a preliminary remark we further generalize the need of adopting and combining different views of the same system and of its heterogeneous components. Going further — and, in some sense, back — in the path described in Section 4, which moved from the micro to the macro view of hardware, and then to the software view, we now distinguish between a *computer-centric* and *a system-centric*

view. As the terms themselves suggest, in a computer-centric view attention is focused on the computing device and its behavior, which may involve interaction with its environment through I/O operations; in a system-centric view, instead, attention is on a whole collection of heterogeneous components, and computing devices — hardware and software — are just a subset thereof. Most often, in such systems the dynamics of the components range over widely different time scales and time granularities (in particular continuous and discrete components are integrated).

In a *computer-centric* view we consider systems where time is inherently discrete, and which can be described with a (finite-)state model. Moreover, we usually adopt a strictly synchronous model of concurrency, where the global synchrony of the system is given by the global clock ticking. Nondeterminism is also often adopted to model concurrent computations at an abstract level.

Another typical feature of this view is the focus on the ease of — possibly automated — analysis to validate some properties; in general, it is possible and preferred to restrict and abstract away from many details of the time behavior in favor of a decidable formal description, amenable to automated verification.

An example of computer-centric view is the design and analysis of a field bus for process control: the attention is focused on discrete signals coming from several sensors and on their proper synchronization; the environment that generates the signals is "hidden" by the interface provided by the sensors.

Conversely, in the *system-centric* view, the aim is to model, design, and analyze the whole system; this includes the process to be controlled, the sensors and actuators, the network connecting the various elements, the computing devices, etc.

In the *system-centric* view, according to what kind of application domain we consider, time is sometimes continuous, and sometimes discrete. The concurrency model is often asynchronous, and the evolution of components is usually deterministic. For instance, a controlled chemical process would be described in terms of continuous time and asynchronous deterministic processes; on the other hand a logistic process — such as the description of a complex storage system — would be probably better described in terms of discrete time. Finally, the system-centric view puts particular emphasis on input/output variables, modular divisions among components, and the resulting "information flow", similarly to some aspects of dynamical systems. Thus, the traditional division between hardware and software is blurred, in favor of the more systemic aspects.

In practice, no model is usually taken to be totally computer-centric or system-centric; more often, some aspects of both views are united within the same model, tailored for some specific needs.

The remainder of this section presents some broad classes of formal languages, in order to discuss what kind of temporal models they introduce, and what kind of systems they are suitable to describe.

We first analyze a selected sample of operational formalisms. Then, we discuss descriptive formalisms based on logic, and devote particular attention to

some important ones. Finally, we present another kind of descriptive notations, the algebraic formalisms, that are mostly timed versions of successful untimed formal languages and methods.

To discuss some features of the formalisms surveyed we will adopt a simple running example based on a resource allocator. Let us warn the reader, however, that the various formalizations proposed for the running example do not aim at being different specifications of the same system; on the contrary, the semantics may change from case to case, according to which features of the formalism we aim at showing in that particular instance.

5.1   Operational Formalisms

We consider three broad classes of operational formalisms: synchronous state machines, Petri nets as the most significant exponent of asynchronous machines, and heterogeneous models.

5.1.1   Synchronous Abstract Machines

In Section 4 we presented some classes of (finite-)state machines that have a synchronous behavior. As we noticed there, those models are mainly derived from the synchronous "macro" view of hardware digital components, and they are suitable to describe "traditional" *sequential* computations.

The natural evolution of those models, in the direction of increasing complexity and sophistication, considers *concurrent* and *reactive* systems. These are, respectively, systems where different components operate in parallel, and open systems whose ongoing interaction with the environment is the main focus, rather than a static input/output relation. The models presented in this section especially tackle these new modeling needs.

Infinite-Word Finite-State Automata.   Perhaps the simplest extension of automata-based formalisms to deal with reactive computations consists in describing a semantics of these machines over infinite (in particular, denumerable) sequences of input/output symbols. This gives rise to finite-state models that are usually called "automata on infinite words" (or ω-words). The various flavors of these automata differ in how they define acceptance conditions (that is, how they distinguish between the "good" and "bad" interactions with the environment) and what kind of semantic models they adopt.

Normally these models are defined in a nondeterministic version, whose transition relation $\delta \subseteq \Sigma \times S \times S$ (where Σ is the input alphabet, and S is the state space) associates input symbol, current state and next state. Thus, for the same pair $\langle \sigma, s \rangle$ of input symbol and current state, more than one next state n may be in relation with it, that is, the automaton can "choose" any of the next states in the set $\{n \mid \langle \sigma, s, n \rangle \in \delta\}$. Nondeterminism and infinite words require the definition of different, more complex acceptance conditions than in the deterministic, finite word case. For instance, the Büchi acceptance condition is defined through a set of final states, some of which must be visited infinitely often in at least

one of the nondeterministically-chosen runs [Var96]. Other acceptance conditions are defined, for instance, in Rabin automata, Streett automata, parity automata, Muller automata, tree automata, etc. [Tho90].

As an example of use of infinite-word automata, let us model a simple resource manager. Before presenting the example, however, we warn the reader that we are not interested in making the resource manager as realistic as possible; rather, as our aim is to show through small-sized models the most relevant features of the formalisms presented, for the sake of brevity we introduce simplifications that a real-world manager would most probably avoid.

The behavior of the resource manager is the following: Users can issue a request for a resource either with high priority (hpr) or with low priority (lpr). Whenever the resource is free and a high-priority request is raised, the resource is immediately granted and it becomes occupied. If it is free and a low-priority request is received, the resource is granted after two time units. Finally, if a high-priority request is received while the resource is granted, it will be served as soon as the resource is released, while a low-priority request will be served two instants after the resource is released. Further requests received while the resource is occupied are ignored.

The above behavior can be modeled by the automaton of Figure 9, where the various requests and grant actions define the input alphabet (and noop defines a "waiting" transition); note that the automaton is actually deterministic. We assume that all states are accepting states.

Figure 9: A resource manager modeled by an infinite-word finite-state automaton.

Let us analyze the infinite-word finite-state automaton models with respect to our coordinates. First of all, these models can be considered as mainly "computer-centric", focusing on simplicity and abstractness. In particular, from the point of view of the computer scientist, they are particularly appealing, as they allow one to reason about time in a highly simplified way.

There is no explicit notion of quantitative time. As usual, however, a simple METRIC is implicitly defined by associating a time unit with the execution of a single transition; thus time is inherently DISCRETE. For example, in Figure 9, we measure IMPLICITLY the two time units after which a low priority request

31

is granted, by forcing the path from the request lpr to occ to pass through two intermediate states via two "wait" transitions noop.

The simplicity of the time model makes it amenable to automated VERIFICATION. Various techniques have been developed to analyze and verify automata, the most successful of whom is probably *model checking* [CGP00]. (See also Section 5.3).

The NONDETERMINISTIC versions of these automata are particularly effective for characterizing multiple computation paths. In defining its formal semantics one may exploit a BRANCHING time model. There are, however, relevant examples of nondeterministic automata that adopt a LINEAR time model, Büchi automata being the most noticeable instance thereof. In fact, modeling using linear time is usually considered more intuitive for the user; for instance, considering the resource manager described above, the linear runs of the automaton naturally represent the possible sequences of events that take place in the manager. This intuitiveness was often considered to be traded off with amenability to automatic verification, since the first model checking procedures were more efficient with branching logic [CGP00]. Later progresses have shown, however, that this trade off is often fictitious, and linear time models may be endowed with efficient verification procedures [Var01].

When COMPOSING multiple automata in a global system we must face the problem of CONCURRENCY. The two most common concurrency models used with finite automata are *synchronous* concurrency and *interleaving* concurrency.

- In *synchronous concurrency*, concurrent transitions of different composed automata occur simultaneously, that is the automata evolve with the same "global" time. This approach is probably the simpler one, since it presents a global, unique vision of time, and is more akin to the "synchronous nature" of finite-state automata. Synchronous concurrency is pursued in several languages that constitute extensions and specializations of the basic infinite-word finite-state automaton, such as Esterel [BG92] and Statecharts (see below).

- In *interleaving concurrency*, concurrent transitions are ordered arbitrarily. Then any two global orderings of the transitions that differ only for the ordering of concurrent transitions are considered equivalent. Interleaving semantics may be regarded as a way to introduce a weak notion of concurrency in a strictly synchronous system. The fact that interleaving introduces partially ordered transitions weakens however the intuitive notion of time as a total order. Also, the natural correspondence between the execution of a single transition and the elapsing of a time unit is lost and *ad hoc* rules are required to restate a time metric based on the transition execution sequence.

  Another problem introduced by adopting an interleaving semantic model lies in the FAIRNESS requirement, which prescribes that every concurrent request eventually gets satisfied. Usually, fairness is enforced explicitly *a priori* in the composition semantic.

The main strength of the infinite-word finite-state automata models, i.e., their simplicity, constitutes also their main limitation. When describing physical systems, adopting a strictly synchronous and discrete view of time might be an obstacle to a "natural" modeling of continuous processes, since discretization may be too strong of an abstraction. In particular, some properties may not hold after discretization, such as periodicity if the duration of the period is some irrational constant, incommensurable with the duration of the step assumed in the discretization.

Moreover, it is very inconvenient to represent heterogeneous systems with this formalism when different components run at highly different speeds and the time GRANULARITY problem arises. In more technical terms, for this type of models it is rather difficult to achieve *compositionality* [AFH96, AH92b].

Statecharts. Statecharts are an automata-based formalism, invented by David Harel [Har87]. They are a quite popular tool in the software engineering community, and a version thereof is part of the UML standard [UML05, UML04].

In a nutshell, Statecharts are an enrichment of classical finite-state automata that introduces some mechanisms for hierarchical abstraction and parallel composition (including synchronization and communication mechanisms). They may be regarded as an attempt to overcome some of the limitations of the bare finite-state automaton model, while retaining its advantages in terms of simplicity and ease of graphical representation. They assume a synchronous view of communication between parallel processes.

Let us use the resource manager running example to illustrate some of Statecharts' features; to this purpose we introduce some modifications to the initial definition. First, after any request has been granted, the resource must be released within 100 time units. To model such METRIC temporal constraints we associate a *timeout* to some states, namely those represented with a short squiggle on the boundary (such as hhr or wg in Figure 10).

Thus, for instance, the transition that exits state hhr must be taken within 100 time units after hhr has been entered: if no rel event has been generated within 100 time units, the timeout event to is "spontaneously-generated" exactly after 100 time units.[11] Conversely the lower bound of 0 in the same state indicates that the same transition cannot be taken immediately. We use the same mechanism to model the maximum amount of time a low-priority request may have to wait for the resource to become available; in this case, with respect to the previous example, we allow the low-priority request to be granted immediately, nondeterministically. Notice that modeling time constraints using timeouts (and exit events) implies an IMPLICIT modeling of a global system time, with respect to which timeouts are computed, just like in finite-state automata.

---

[11]Note that there are in fact two transitions from state hhr to state no-hr, one that is labeled to/rel, and one that is labeled rel; they are represented in Figure 10 with a single arc instead of two separate ones for the sake of readability. The transition labeled to/rel indicates that when the timeout expires (the to event), a rel event is triggered, which is then sensed by the other parts of the Statechart, hence producing other state changes (for example from glr to free).
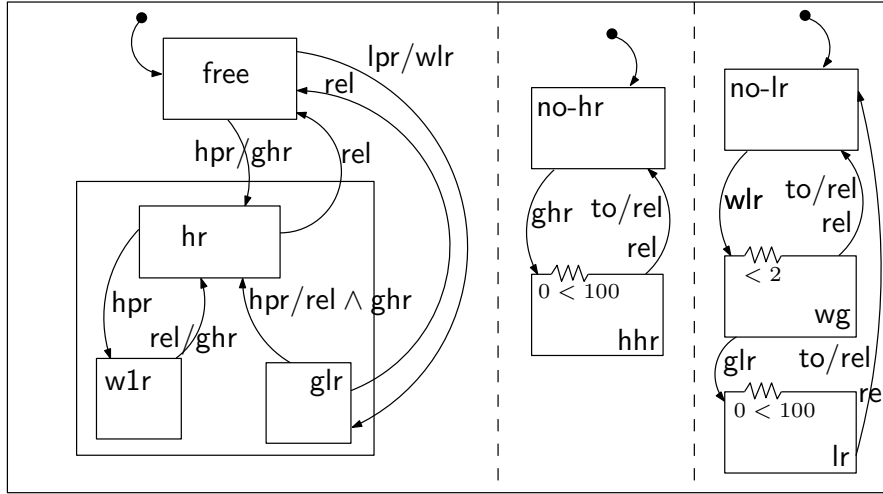
Figure 10: A resource manager modeled through a Statechart.

In fact, timeouts can be regarded as an enrichment of the discrete finite state automaton model with a CONTINUOUS feature.

The example of Figure 10 exploits Statecharts' so-called "AND (parallel) composition" to represent three logically separable components of the system, divided by dashed lines. The semantics of AND composition is obtained as the Cartesian product construction,[12] and it is usually called SYNCHRONOUS COMPOSITION;[13] however, Statecharts' graphical representation avoids the need to display all the states of the product construction, ameliorating the readability of a complex specification. In particular in our example, we choose to allow one pending high-priority request to be "enqueued" while the resource is occupied; thus the leftmost component is a finite-state automaton modeling whether the resource is free, serving a high-priority request with no other pending requests (state hr), or with one pending request (state w1r), or serving a low-priority request (state glr).

Since in Statecharts all transition events — both input and output — are "broadcast" over the whole system, labeling different transitions with the same name enforces synchronization between them. For instance, whenever the automaton is in the global state ⟨w1r, hhr, no-lr⟩, a release event rel triggers the global state to become ⟨hr, no-hr, no-lr⟩, and then cascading immediately to ⟨hr, hhr, no-lr⟩, because of the output event ghr triggered by the transition from w1r to hr. Note that we are implicitly assuming, in the example above, that

---

[12]In fact, the semantics of the AND composition of submachines in Statecharts differs slightly from the classic notion of Cartesian product of finite-state machines; however, in this article we will not delve any further in such details, and instead refer the interested reader to [Har87] for a deeper discussion of this issue.

[13]We warn the reader that the terminology often varies greatly among different areas; for instance [CL99] names the Cartesian product composition "completely asynchronous".

ghr and wlr are "internal events", i.e., they do not occur spontaneously in the environment but can only be generated internally for synchronization.

NONDETERMINISM can arise in three basic features of Statecharts models. First, we have the "usual" nondeterminism of two mutually exclusive transitions with the same input label (such as in Figure 11(a)). Second, states with timeout are exited nondeterministically *within* the prescribed bounds (Figure 11(b)). Third, Statechart modules may be composed with "XOR composition", that represents a nondeterministic choice between different modules (Figure 11(c)).


(a)   (b)   (c)

Figure 11: Nondeterminism in Statecharts.

The popularity of Statecharts has produced an array of different ANALYSIS TOOLS, mostly automated. For instance [HLN$^+$90, BDW00, GTBF03].

While overcoming some of the limitations of the basic finite-state automata models, Statecharts' rich syntax often hides subtle semantic problems that instead should be better exposed to avoid inconsistencies and faults in specifications. In fact, over the years several researches have tried to define formally the most crucial aspects of the temporal semantics of Statecharts. The fact itself that different problems were unveiled only incrementally by different contributors is an indication of the difficulty of finding a comprehensive, intuitive, non-ambiguous semantics to an apparently simple and plain language. We discuss here just a few examples, referring the interested reader to [HPSS87, PS91, von94, HN96] for more details.

The apparently safe "perfect SYNCHRONY" assumption — the assumption that all transition events occur simultaneously — and the global "broadcast" availability of all events — which are therefore non local — generate some subtle difficulties in obtaining a consistent semantics. Consider for instance the example of Figure 10, and assume the system is in the global state ⟨glr, no-hr, lr⟩. If a high-priority request takes place, and thus a hpr event is generated, the system shifts to the state ⟨hr, no-hr, lr⟩ in zero time. Simultaneously, the taken transition triggers the events rel and ghr. If we allow a zero-time residence in states, the former event moves the system to ⟨hr, no-hr, no-lr⟩, representing the low-priority request being forced to release the resource. Still simultaneously, the latter ghr event triggers the transition from no-hr to hhr in the middle sub-automaton. This is in conformity with our intuitive requirements; however the same rel generated event also triggers the first sub-automaton to the state free, which is instead against the intuition that suggests that the event is only a message sent to the other parts of the automaton.

If we refine the analysis, we discover that the picture is even more compli-

cated. The middle automaton is in fact in the state hhr, while the time has not advanced; thus we still have the rel event available, which should immediately switch the middle automaton back to the state no-hr. Besides being intuitively not acceptable, this is also in conflict with the lower bound on the residence time in hhr. Moreover, in general we may end up having multiple XOR states occupied at the same time. Finally, it is not difficult to conceive scenarios in which the simultaneous occurrence of some transitions causes an infinite sequence of states to be traversed, thus causing a ZENO behavior.

How to properly disentangle such scenarios is not obvious. A partial solution would be, for instance, to avoid instantaneous transitions altogether, attaching a non-zero time to transitions and forcing an ordering between them or, symmetrically, to disallow a zero-time residence in states. This (partially) asynchronous approach is pursued for instance in Timed Statecharts [KP92], or in other works [Per93]. Alternatively, other solutions disallow loops of zero-time transitions, but accept a finite number of them (for instance, by "consuming" each event spent by a transition [HN96]); the Esterel language, which is a "relative" of Statecharts', follows this approach.

Timed and Hybrid Automata. As we discussed above, the strictly discrete and synchronous view of finite-state automata may be unsuitable to model adequately and compositionally processes that evolve over a dense domain. Statecharts try to overcome these problems by adding some continuous features, namely timeout states. Timed and hybrid automata push this idea further, constituting models, still based on finite-state automata, that can manage continuous variables. Let us first discuss timed automata.

Timed automata enrich the basic finite-state automata with real-valued *clock* variables. Although the name "timed automata" could be used generically to denote automata formalisms where a description of time has been added (e.g., [LV96, AH96, Arc00]), here we specifically refer to the model first proposed by Alur and Dill [AD94], and to its subsequent enrichments and variations. We refer the reader to Alur and Dill's original paper [AD94] and to [BY04] for a formal, detailed presentation.

In timed automata, the total state is composed of two parts: a finite component (corresponding to the state of a finite automaton, which is often called *location*), and a continuous one represented by a finite number of positive real values assigned to variables called *clocks*. The resulting system has therefore an *infinite* state space, since the clock components take value in the infinite set $\mathbb{R}_{\geq 0}$. The evolution of the system is made of alternating phases of instantaneous synchronous discrete "jumps" and continuous clock increases. More precisely, whenever a timed automaton sits in some discrete state, each clock variable x increases as time elapses, that is it evolves according to the dynamic equation $\dot{x} = 1$, thus effectively measuring time. External input events cause the discrete state to switch; during the transition some clock variables may be reset to zero instantaneously. Moreover, both discrete states and transitions may have attached some constraints on clocks; each constraint must be satisfied while

sitting in the discrete state, and when taking the transition, respectively.[14]

To illustrate this notation, let us model the resource manager example through a timed automaton. We modify the system behavior of the State-chart example, by disallowing high-priority requests to preempt low-priority ones; moreover, let us assume that one low-priority request can be "enqueued" waiting for the resource to become free. The resulting timed automaton — using a single clock w — is pictured in Figure 12.
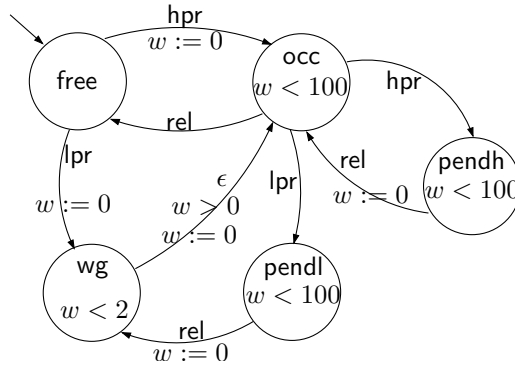

Figure 12: A resource manager modeled through a timed automaton.

The semantics of a timed automaton is usually formally defined by means of a timed transition system. The "natural" semantics is the *timed* semantics, which exactly defines the possible runs of one automaton over sequences of input symbols. More precisely, each symbol in the input sequence is paired with a *timestamp* that indicates the absolute time at which the symbol is received. Then, a run is defined by a sequence of total states (each one a pair ⟨location, clock value⟩ of the automaton, which evolve according to the times-tamped input symbols, in such a way that, for every pair of consecutive states $\langle l_i, c_i \rangle \xrightarrow{\text{in}, ts} \langle l_{i+1}, c_{i+1} \rangle$ in the run the constraints on the locations and the transition are met. For instance, the automaton of Figure 12 may go through the following run:

$$\langle \mathsf{free}, 0 \rangle \xrightarrow{\mathsf{hpr}, 4.7} \langle \mathsf{occ}, 0 \rangle \xrightarrow{\mathsf{lpr}, 53.9} \langle \mathsf{pendl}, 49.2 \rangle \xrightarrow{\mathsf{rel}, 64} \langle \mathsf{wg}, 0 \rangle \xrightarrow{\epsilon, 65.1} \langle \mathsf{occ}, 0 \rangle \cdots$$

In the run above state location occ is entered at time 4.7 and, since the corresponding transition resets clock w, the new state becomes $\langle \mathsf{occ}, 0 \rangle$; then, at time 53.9 (when clock w has reached value 49.2), location occ is exited and pendl is entered (this time, the clock w is not reset), which satisfies the constraint $w < 100$ of location occ, and so on.

Timed semantics introduces a METRIC treatment of time through times-tamps. Notice that, in some sense, the use of timestamps introduces "two dif-

---

[14]The original Alur and Dill's formalization [AD94] permitted constraints only on transitions; however, adding constraints to locations as well is a standard extension that does not impact on the salient features of the model (expressiveness, in particular) [BY04].

ferent notions of time": the inherently DISCRETE one, given by the position i in the run/input sequence, which defines a total ordering on events, and the CONTINUOUS and metric one, recorded by the timestamps and controlled through the clocks. This approach, though simple in principle, somewhat sacrifices naturalness, since a complete time modeling is no more represented as a unique flow but is two-fold.

Other, different semantics of timed automata have been introduced and analyzed in the literature. Subtle differences often arise depending on which semantics is adopted; for instance, interval-based semantics interprets timed automata over piecewise-constant functions of time, and the change of location is triggered by discontinuities in the input [AFH96, ACM02, Asa04].

Let us consider a few more features of time modeling for timed automata.

- While timed automata are in general NONDETERMINISTIC, their semantics is usually defined through LINEAR time models, such as the one outlined above based on run sequences. Moreover, deterministic timed automata are strictly less expressive than nondeterministic ones, but also more amenable to automated verification, so they may be preferred in some practical cases.

- *Absolute time is* IMPLICITLY assumed in the model and becomes apparent in the timestamps associated with the input symbols. The *relative time* measured by clocks, however, is EXPLICITLY measured and set.

- Timed automata may exhibit ZENO *behaviors*, when distances between times at which transitions in a sequence are taken become increasingly smaller, accumulating to zero. For instance, in the example of Figure 12, the two transitions hpr and rel may be taken at times $1, 1 + 2^{-1}, 1 + 2^{-1} + 2^{-2}, \ldots, \Sigma_{k=0}^{n} 2^{-k}, \ldots$, so that the absolute time would accumulate at $\Sigma_{k=0}^{\infty} 2^{-k} = 2$. Usually, these Zeno behaviors are ruled out *a priori* in defining the semantics of timed automata, by requiring that timestamped sequences are acceptable only when the timestamp values are unbounded.

  Moreover, in Alur and Dill's formulation [AD94] timed words have *strictly monotonic* timestamps, which implies that some time (however small) must elapse between two consecutive transitions; other semantics have relaxed this requirement by allowing weakly monotonic timestamps [BY04], thus permitting sequences of zero-time transitions.

Hybrid automata [ACHH93, NOSY93, Hen96] are a generalization of timed automata where the dense-valued variables — called "clocks" in timed automata — are permitted to evolve through more complicated timed behaviors. Namely, in hybrid automata one associates to each *discrete state* a set of possible *activities*, which are smooth functions (i.e., functions that are continuous together with all of their derivatives) from time to the dense domain of the *variables*, and a set of *invariants*, which are sets of allowed values for the variables. Activities specify possible variables' behaviors, thus generalizing the simple dynamics

of clock variables in timed automata. More explicitly, whenever a hybrid automaton sits in some discrete location, its variables evolve over time according to one activity, *nondeterministically* chosen among those associated with that state. However, the evolution can continue only as long as the variables keep their values within the invariant set of the state. Then, upon reading input symbols, the automaton instantaneously switches its discrete state, possibly resetting some variables according to the additional constraints attached to the taken transitions, similarly to timed automata.

Although in this general definition the evolution of the dense-valued variables can be represented by any function such that all its derivatives are continuous, in practice more constrained (and simply definable) subsets are usually considered. A common choice is to define the activities by giving a set of bounds on the first-order derivative, with respect to time, of the variables. For a variable y, the constraint $0.5 < \dot{y} < \pi$ is an example of a class of such activities (see Figure 13 for a visual representation).

Figure 13: Some behaviors compatible with the constraint $0.5 < \dot{y} < \pi$.

In both timed and hybrid automata, one typically defines a COMPOSITION semantics where concurrent automata evolve in parallel, but synchronize on transitions in response to input symbols, similarly to traditional automata and Statecharts.

The development of timed and hybrid automata was also motivated by the desire to extend and generalize the powerful and successful techniques of automatic VERIFICATION (and model checking in particular) based on the combination of infinite-word finite-state automata and temporal logic (see Section 5.3), to the metric treatment of time. However, the presence of real-valued variables renders the verification problem much more difficult and, often, undecidable. Thus, with respect to the general model, restrictions are introduced that make the models more tractable and amenable to verification — usually at the price of sacrificing some expressiveness.

In a nutshell, the verification problem is generally tackled by producing a *finite abstraction* of a timed/hybrid automaton, where all the relevant behaviors of the modeled system are captured by an equivalent, but finite, model, which is therefore exhaustively analyzable by model checking techniques. Such

procedures usually assume that all the numeric constraints on clocks and variables are expressed by *rational* numbers; this permits the partitioning of the space of all possible behaviors of the variables into a finite set of *regions* that describe equivalent behaviors, preserving verification properties such as reachability and emptiness. For a precise description of these techniques see e.g., [AM04, ACH$^+$95, HNSY94, HKPV98].

These analysis techniques have been implemented in some interesting tools, such as UPPAAL [LPY97], Kronos [Yov97], Cospan [AK95], IF [BGO$^+$04], and HyTech [HHWT97].

Timed Transition Models. Ostroff's *Timed Transition Models* (TTM) [Ost90] are another formalism that is based on enriching automata with time variables; they are a real-time METRIC extension of Manna and Pnueli's fair transition systems [MP92].

In TTMs, time is modeled EXPLICITLY by means of a clock variable t. t takes values in a DISCRETE time domain, and is updated explicitly and SYNCHRONOUSLY by the occurrence of a special *tick* transition. The clock variable, as any variable in TTMs, is global and thus shared by all transitions. All transitions other than *tick* do not change time but only update the other components of the state; therefore it is possible to have several different states associated with the same time instant. Transitions are usually annotated with lower and upper bounds l, u; this prescribes that the transition is taken at least l, and no more than u clock ticks (i.e., time units), after the transition has become enabled.

In practice, it is assumed that every TTM system includes a *global clock* subsystem, such as that pictured in Figure 14. Notice that this subsystem allows the special *tick* transition to occur at any time, making time advance one step. The *tick* transition is *a priori* assumed to be fairly scheduled, that is it must occur infinitely often to prevent ZENO behaviors where time stops.

tick $:\rightarrow t := t + 1$

clock

Figure 14: A Timed Transition Model for the clock.

We give a few more details of TTMs in Section 5.3 (where a TTM resource manager specification is also given) when discussing dual language approaches.

5.1.2 Asynchronous Abstract Machines: Petri nets

This section introduces Petri nets as one of the most popular examples of asynchronous abstract machines.

Petri nets owe their name to their inventor, Carl Adam Petri [Pet63]. Since their introduction they became rather popular both in the academic and, to

some extent, in the industrial world, as a fairly intuitive graphical tool to model concurrent systems. For instance, they inspired transition diagrams adopted in the UML standard [UML05, UML04, EPLF03]. There are a few slightly different definitions of such nets and of their semantics. Among them one of the most widely adopted is the following, which we present informally; the reader is referred to the literature [Pet81, Rei85] for a comprehensive treatment.

A *Petri net* consists of a set of places, and a set of transitions. Places store tokens and pass them to transitions. A transition is *enabled* whenever all of the incoming places hold at least one token. Whenever a transition is enabled a *firing* can occur; this happens nondeterministically. As a consequence of a firing, the enabling tokens are removed from the incoming places and moved to the outgoing places the transition is connected to. Thus, for any possible combination of nondeterministic choices, we have a *firing sequence*.

Let us consider again the example of the resource manager, using a Petri net model. We introduce the following modifications with respect to the previous examples. First, since we are now considering untimed Petri nets, we do not introduce any metric time constraint. Second, we disallow low-priority requests while the resource is occupied, or high-priority requests while there is a pending low-priority request. Conversely, we introduce a mechanism to "count" the number of consecutive high-priority requests that occur while the resource is occupied. Then, we make sure that all of them are served (consecutively) before the resource becomes free again. This behavior is modeled by the Petri net in Figure 15, where the places are denoted by the circles free, occ, pendh, wr, and wg, and the thick lines denote transitions. Notice that we allow an unbounded number of tokens in each place (actually, the only place where the tokens can accumulate is pendh, where each token represents a pending high-priority request). Finally, we have also chosen to introduce an *inhibiting arc*, from place pendh to transition rel_2, denoted by a small circle in place of an arrowhead: this means that the corresponding transition is enabled if and only if place pendh stores no tokens. This is a non-standard feature of Petri nets which is often added in the literature to increase the model's expressive power.

According to our taxonomy, Petri nets, as defined above, can be classified as follows:

- There is no explicit notion of time. However a time model can be IMPLIC-ITLY associated with the semantics of the net.

- There are at least two major approaches to formalizing the semantics of Petri nets.

  - The simpler one is based on *interleaving semantics*. According to this semantics the behaviors of a net are just its firing sequences. Interleaving semantics, however, introduces a total ordering in the events modeled by the firing of net transitions which fails to capture the asynchronous nature of the model. For instance, in the net of Figure 15 the two sequences $\langle hpr_1, hpr_2, hpr_3, rel_3, hpr_3, rel_3, rel_3, rel_2 \rangle$ and $\langle hpr_1, hpr_2, hpr_3, hpr_3, rel_3, rel_3, rel_3, rel_2 \rangle$ both belong to the set of
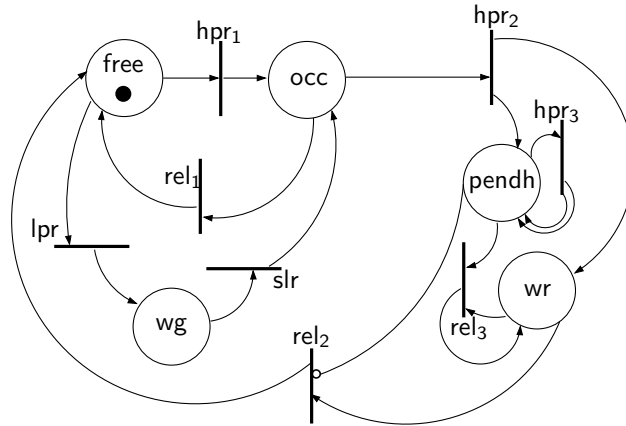
Figure 15: A resource manager modeled through a Petri net.

– possible net's behaviors; however, they both imply an order between the firing of transitions hpr_3 and rel_3, whereas the graphical structure of the net emphasizes that the two events can occur asynchronously (or simultaneously).

– For this reason, a *true concurrency* (i.e., fully ASYNCHRONOUS) approach is often preferred to describe the semantics of Petri nets. In a true concurrency approach it is natural to see the time model as a *partial order*, instead of a total order of the events modeled by transition firings. Intuitively, in a true concurrency modeling the two sequences above can be "collapsed" into $\langle \mathsf{hpr}_1, \mathsf{hpr}_2, \mathsf{hpr}_3, \{\mathsf{hpr}_3, \mathsf{rel}_3\}, \mathsf{rel}_3, \mathsf{rel}_3, \mathsf{rel}_2\rangle$, where the pair { } denotes the fact that the included items can be "shuffled" in any order.

• Petri nets are a NONDETERMINISTIC operational model. For instance, still in the net of Figure 15, whenever place occ holds some tokens, both transitions hpr_2 and rel_1 are enabled, but they are in *conflict*, so that only one of them can actually fire. Such a nondeterminism could be formalized by exploiting a BRANCHING-*time* model.

• In traditional Petri nets the time model has no METRICS, so that it should be seen only as a (possibly partial) order.[15]

• We also remark that Petri nets are usually "less compositional" than other operational formalisms, and synchronous automata in particular. While notions of COMPOSITION of Petri nets have been introduced in the literature, they are often less natural and more complicated than, for in-

---

[15]Unless one adopts the convention of associating one time unit to the firing of a single transition, as it is often assumed in other — synchronous — operational models such as finite state automata. Such an assumption, however, would contrast sharply with the asynchronous original nature of the model.

stance, Statecharts' synchronous composition; this is partly due to the asynchronous "nature" of the nets.

To model hard real-time systems a metric time model is necessary in most cases. To overcome this difficulty, many extensions have been proposed in the literature to introduce a metric time model. Here we report on Merlin and Farber's approach [MF76], which has been probably the first one of such extensions and is one of the most intuitive and popular ones. For a thorough and comprehensive survey of the many time extensions to Petri nets we refer to [CMS99, Cer93].

A Timed Petri net according to the Merlin and Farber's approach is simply a net where a minimum and a maximum firing time are attached to each transition (both firing times can be 0, and the maximum time can be ∞). Figure 16 shows how the net of Figure 15 can be augmented in such a way. The time bounds that have been introduced refine the specification of the resource manager by prescribing that each use of the resource must take no longer than 100 contiguous (i.e., since the last request occurred) time units, and that a low priority request is served within 2 time units.
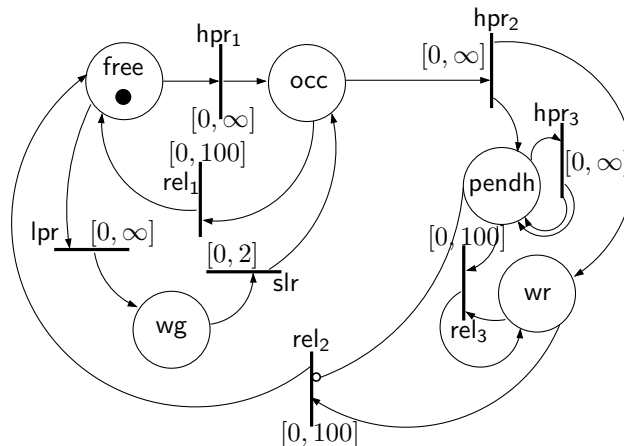
Figure 16: A resource manager modeled through a timed Petri net.

The fairly natural intuition behind this notation is that, since the time when a transition is enabled (i.e., all its input places have been filled with at least one token), the transition can fire — nondeterministically — at any time that is included in the specified interval, unless it is disabled by the firing of a conflicting transition. For instance, place wg becomes occupied after a low priority request is issued, thus enabling transition slr. The latter can fire at any time between 0 and 2 time instants after it has become enabled, thus expressing the fact that the request is served *within* 2 time units.

Despite its intuitive attractiveness, several intricacies are hidden in the previous informal definition, as has been pointed out in the literature when at-

tempting to formalize their semantics [FMM94, GMMP91]. Here we focus only on the main ones.

- Suppose that the whole time interval elapsed since the time when a transition became enabled: is at this point the transition *forced* to fire or not? In the negative case it will never fire in the future and the tokens in their input places will be wasted (at least for *that* firing). There are arguments in favor of both choices; normally — including the example of Figure 16 — the former one is assumed (it is often called *strong time semantics* (STS)) but there are also cases where the latter one (which is called *weak time semantics* (WTS) and is considered as more consistent with traditional Petri nets semantics, where a transition is never forced to fire) is preferred.[16]

- If the minimum time associated with a transition is 0, then the transition can fire immediately once enabled and we have a case of zero-time transition (more precisely we call this circumstance *zero-time firing* of the transition). As we pointed out in other cases, zero-time firing can be a useful abstraction whenever the duration of the event modeled by the firing of the transition can be neglected with respect to other activities of the whole process.[17] On the other hand zero-time firing can produce some intricate situations since two subsequent transitions (e.g., hpr_1 and rel_1 in Figure 16) could fire simultaneously. This can produce some ZENO behaviors if the net contains loops of transitions with 0 minimum time. For this reason "zero-time loops" are often forbidden in the construction of timed Petri nets.

Once the above semantic ambiguities have been clarified, the behavior of timed Petri nets can be formalized through two main approaches.

- A time stamp can be attached to each token when it is produced by the firing of some transition in an output place. For instance, with reference to Figure 16, we might have the sequence of transitions $\langle \mathsf{hpr}_1(2), \mathsf{hpr}_2(3), \mathsf{hpr}_3(4), \mathsf{rel}_3(5), \mathsf{hpr}_3(6), \cdots \rangle$ (that is, hpr_1 fires at time 2 producing a token with time stamp 2 in occ; this is consumed at time 3 by the firing of hpr_2 which also produces one token in pendh and one in wr, both timestamped 3, etc.). In this way time is EXPLICITLY modeled in a METRIC way — whether DISCRETE or CONTINUOUS — as a further variable describing

---

[16]In this regard, notice that the timed automata of Section 5.1.1 could be considered to have a *weak time semantics*. In fact, transitions in timed automata *are not forced* to be taken when the upper limit of some constraint is met; rather, all that is prescribed by their semantics is that *when* (if) a transition is taken by a timed automaton, its corresponding constraint (and those of the source and target locations) *must* be met.

[17]Normally the firing of a transition is considered as instantaneous. This assumption does not affect generality since an activity with a non-null duration can be easily modeled as a pair of transitions with a place between them: the first transition models the beginning of the activity and the second one models its end.

system's state and evolution (more precisely, as *many* further variables, one for each produced token).

As remarked in Section 5.1.1, this approach actually introduces two different time models in the formalism: the time implicitly subsumed by the firing sequence and the time modeled by the time stamps attached to tokens. Of course some restrictions should be applied to guarantee consistency between the two orderings: for instance, the same succession of firings described above could induce the timed sequence $\langle \mathsf{hpr}_1(2), \mathsf{hpr}_2(3), \mathsf{hpr}_3(4),$ $\mathsf{hpr}_3(6), \mathsf{rel}_3(5), \cdots \rangle$, that should however be excluded from the possible behaviors.

- The net could be described as a dynamical system as in the traditional approach described in Section 4.1. The system's state would be the net marking whose evolution should be formalized as a function of time. To pursue this approach, however, a few technical difficulties must be overcome:

  - First, tokens cannot be formalized as entities with no identity, as it happens with traditional untimed Petri nets. Here too, some kind of time stamp may be necessary. Consider, for instance, the net fragment of Figure 17, and suppose that one token is produced into place P at time 3 by transition t_1^i and another token is produced by t_2^i at time 4; then, according to the normal interpretation of such Petri nets (but different semantic formalizations could also be given, depending on the phenomenon that one wants to model) the output transition t° should fire once at time 6=3+3 and a second time at time 7=4+3. Thus, a state description that simply asserts that at time 4 there are two tokens in P would not be sufficient to fully describe the future evolution of the net.

Figure 17: An example Petri net fragment.

  - If zero-time firings are admitted, strictly speaking, system's state cannot be formalized as a function of the independent variable "time": consider, for example, the case in which, in the net of Figure 16, at time t both transitions lpr and slr fire (which can happen, since slr admits zero-time firing); in this case, it would happen that at time t

both a state (marking) where place wg is marked and a state where place occ is marked — and wg is not marked anymore — would hold.

In [FMM94] this problem has been solved by forbidding "zero-time loops" and by stating the convention that in case of a "race" of zero-time firings (which is always finite) only the places at the "end of the race" are considered as marked, whereas tokens flow instantaneously through other places without marking them.

In [GMM99] a more general approach is proposed, where zero-time firings are considered as an abstraction of a non-null but *infinitesimal* firing time. By this way it has been shown that mathematical formalization and analysis of the net behavior become simpler and — perhaps — more elegant.

Timed Petri nets have also been the object of a formalization through the dual language approach (see Section 5.3).

As for other formalisms of comparable expressive power, Petri nets suffer intrinsic limitations in the techniques for (semi-)automatic ANALYSIS and VERIFICATION. In fact, let us consider the reachability problem, i.e., the problem of stating whether a given marking can be reached by another given marking. This is the main analysis problem for Petri nets since most other properties can be reduced to some formulation of this basic problem [Pet81]. For normal, untimed Petri nets with no inhibitor arcs the reachability problem has been shown intractable though decidable; if Petri nets are augmented with some metric time model and/or inhibitor arcs, then they reach the expressive power of Turing machines and all problems of practical interest become undecidable.[18] Even building interpreters for Petri nets to analyze their properties through simulation faces problems of combinatorial explosion due to the intrinsic nondeterminism of the model.

Nevertheless interesting tools for the analysis of both untimed and timed Petri nets are available. Among them we mention [BD91], which provides an algorithm for the reachability problem of timed Petri nets assuming the set of rational numbers as the time domain. This work has pioneered further developments. For a comprehensive survey of tools based on Petri nets see [TGI].

Before closing this section let us also mention the Abstract State Machines (ASM) formalism [BS03], whose generality subsumes most types of operational formalisms, whether synchronous or asynchronous. However, ASM have not received, to the best of our knowledge, much attention in the realm of real-time computing until recently, when the Timed Abstract Machine notation [OL07b] and its tools [OL07a] have been developed.

---

[18]Of course, interesting particular cases are always possible, e.g., the case of bounded nets, where the net is such that during its behavior the number of tokens in every place never exceeds a given bound.

5.2 Descriptive Formalisms

Let us now consider *descriptive* (or *declarative*) formalisms. In descriptive formalisms a system is formalized by declaring the fundamental properties of its behavior. Most often, this is done by means of a language based on mathematical logic; more seldom algebraic formalisms (e.g., process algebras) are exploited. As we saw in Section 2, descriptive notations can be used alone or in combination with operational ones, in a dual language approach. In the former case, both the requirements and the system specification are expressed within the same formalism; therefore verification consists of proving that the axioms (often expressed in some logic language) that constitute the system specification imply the formulas that describe the requirements. In the latter case, the verification is usually based on some *ad hoc* techniques, whose features may vary significantly depending on the adopted combination of descriptive and operational notations. We treat dual-language approaches in Section 5.3.

When considering the description of the timed behavior of a system through a logic formalism, it is natural to refer to *temporal logics.* A distinction should be made here. Strictly speaking, temporal logics are a particular family of modal logics [Kri63, RU71] possessing specific operators — called *modalities* — apt to express temporal relationships about time-dependent propositions. The modalities usually make the treating of time-related information quite intuitive as they avoid the EXPLICIT reference to absolute time values and mirror the way the human mind intuitively reasons about time; indeed, temporal logics were initially introduced by philosophers [Kam68]. It was Pnueli who first observed [Pnu77] that they could be effectively used to reason about temporal properties of programs, as well. Some temporal logics are discussed in the following Section 5.2.1.

In the computer science communities, however, the term "temporal logic" has been used in a broader sense, encompassing all logic-based formalisms that possess some mechanism to express temporal properties and to reason about time, even when they introduce some EXPLICIT reference to a dedicated variable representing the current value of time or some sort of clock and hence adopt a style of description that is different in nature from the original temporal logic derived from modal logic. Many of these languages have been used quite successfully for modeling time-related features: some of them are described in Section 5.2.2 below.

We emphasize that there is a wide variety of different styles and flavors when it comes to temporal logics. As usual, we do not aim to be exhaustive in the presentation of temporal logics (we refer the reader to other papers specifically on temporal logics, e.g., [Eme90, AH93, AH92b, Ost92, Hen98, BMN00, FPR08b]), but to highlight some significant approaches to the problem of modeling time in logic.

Finally, a different approach to descriptive modeling of systems, based on the calculational aspects of specifications, is the algebraic one. We discuss algebraic formalisms in Section 5.2.3.

5.2.1 Temporal Logics

In this section we deal with temporal logics with essentially IMPLICIT time, and we focus our discussion on a few key issues, namely the distinction between linear-time and branching-time logics, the adoption of a discrete or non-discrete time model, the use of a metric on time to provide means to express temporal properties in a quantitatively precise way, the choice of using solely temporal operators that refer to the future versus introducing also past-tense operators, and the assumption of time points or time intervals as the fundamental time entities. In our discussion we will go from simple to richer notations and occasionally combine the treatment of some of the above mentioned issues. Finally, some VERIFICATION issues about temporal logics will be discussed while presenting dual language approaches in Section 5.3.

Linear-Time Temporal Logic. As a first, simplest example of temporal logic, let us consider propositional Linear-Time Temporal Logic (LTL) with discrete time. In LTL, formulas are composed from the atomic propositions with the usual Boolean connectives and the temporal connectives X (*next*, also denoted with the symbol $\bigcirc$), F (*eventually in the future*, also $\Diamond$), G (*globally — i.e., always — in the future*, also $\Box$), and U (*until*). These have a rather natural and intuitive interpretation, as the formulas of LTL are interpreted over LINEAR sequences of states: the formula Xp means that proposition p holds at the state that immediately follows the one where the formula is interpreted, Fp means that p will hold at some state following the current one, Gp that p will hold at all future states, pUq means that there is some successive state such that proposition q will hold then, and that p holds in all the states between the current and that one.

Notice that the presence of the "next" operator X implies that the logic refers to a DISCRETE temporal domain: by definition, there would be no "next state" if the interpretation structure domain were not discrete. On the other hand, depriving LTL of the next operator would "weaken" the logic to a pure ordering without any metrics (see below).

To illustrate LTL's main features, let us consider again the resource manager introduced in the previous sections: the following formula specifies that, if a low priority request is issued at a time when the resource is free, then it will be granted at the second successive state in the sequence.

$$G(\text{free} \wedge \text{lpr} \Rightarrow \text{XXocc})$$

LTL is well-suited to specify qualitative time relations, for instance ordering among events: the following formula describes a possible assumption about incoming resource requests, i.e., that no two consecutive high priority requests may occur without a release of the resource between them (literally, the formula reads as: if a high priority request is issued then the resource must be eventually released and no other similar request can take place until the release occurs).

$$G(\text{hpr} \Rightarrow \text{X}(\neg\text{hpr U rel}))$$

Though LTL is not expressly equipped with a METRIC on time, one might use the next operator X for this purpose: for instance, $X^3 p$ (i.e., XXXp) would mean that proposition p holds 3 time units in the future. The use of X^k to denote the time instant at k time units in the future is only possible, however, under the condition that there is a one-to-one correspondence between the states of the sequence over which the formulas are interpreted and the time points of the temporal domain. Designers of time-critical systems should be aware that this is not necessarily the case: there are linear discrete-time temporal logics where two consecutive states may well refer to the same time instant whereas the first following state associated with the successive time instant is far away in the state sequence [Lam94, MP92, Ost89]. We already encountered this critical issue in the context of finite state automata and the FAIRNESS problem (see Section 5.1.1) and timed Petri nets when zero-time transitions are allowed (see Section 5.1.2) and will encounter it again in the dual language approach (Section 5.3).

Metric Temporal Logics. Several variations or extensions of LINEAR time temporal logic have been defined to endow it with a metric on time, and hence make it suitable to describe strict real-time systems. Among them, we mention Metric Temporal Logic (MTL) [Koy90] and TRIO [GMM90, MMG92]. They are commonly interpreted both over DISCRETE and over DENSE (and CONTINUOUS) time domains.

MTL extends LTL by adding to its operators a QUANTITATIVE time parameter, possibly qualified with a relational symbol to imply an upper bound for a value that typically represents a distance between time instants or the length of some time interval. For instance the following simple MTL formula specifies bounded response time: there is a time distance d such that an event p is always followed by an event q with a delay of at most d time units (notice that MTL is a first-order logic).

$$\exists d : G(p \Rightarrow F_{<d} q) \tag{2}$$

The following formula asserts that p eventually takes place, and then periodically occurs with period d.

$$F(p \wedge G(\neg p\, U_d\, p))$$

TRIO introduces a quantitative notion of time by adopting a single basic modal operator, called *Dist*. The simplest formula $Dist(p, d)$ means that proposition p holds at a time instant exactly d time units from the current one; notice that this formula may refer to the future, if $d > 0$, or to the past, if $d < 0$, or even to the present time if $d = 0$. All the operators of LTL, their quantitative-time counterparts and also other operators not found in traditional temporal logic are defined in TRIO by means of first-order quantification over the time parameter of the basic operator *Dist*. We include in Table 3 a list of some of the most significant ones (and especially those used in the following).

Referring again to the example of the resource manager, the following TRIO formula asserts that any low priority resource request is satisfied within 100

OPERATOR	DEFINITION	DESCRIPTION
$\text{Futr}(F, t)$	$t \geq 0 \wedge \text{Dist}(F, t)$	F holds t time units in the future
$\text{Past}(F, t)$	$t \geq 0 \wedge \text{Dist}(F, -t)$	F held t time units in the past
$\text{Alw}(F)$	$\forall d : \text{Dist}(F, d)$	F holds always
$\text{Lasts}(F, t)$	$\forall d \in (0, t) : \text{Futr}(F, d)$	F holds for t time units in the future
$\text{Lasted}(F, t)$	$\forall d \in (0, t) : \text{Past}(F, d)$	F held for t time units in the past
$\text{WithinF}(F, t)$	$\exists d \in (0, t) : \text{Futr}(F, d)$	F holds within t time units in the future
$\text{Until}(F, G)$	$\exists d > 0 : \text{Lasts}(F, d) \wedge \text{Futr}(G, d)$	F holds until G holds
$\text{NowOn}(F)$	$\exists d > 0 : \text{Lasts}(F, d)$	F holds for some non-empty interval in the future
$\text{UpToNow}(F)$	$\exists d > 0 : \text{Lasted}(F, d)$	F held for some non-empty interval in the past

Table 3: TRIO derived temporal operators.

time units
$$\text{Alw}(\mathsf{lpr} \Rightarrow \text{WithinF}(\mathsf{occ}, 100))$$
while the next one states that any two high priority requests must be at least 50 time units apart.
$$\text{Alw}(\mathsf{hpr} \Rightarrow \text{Lasts}(\neg\mathsf{hpr}, 50))$$

We note incidentally that both in MTL and in TRIO the interpretation structure associates one single state with every time instant and no EXPLICIT state component needs to be devoted to the representation of the current value of "time": quantitative timing properties can be specified using the modal operators embedded in the language. Other approaches to the quantitative specification of timing properties in real-time systems are based on the use of the operators of (plain) LTL in combination with assertions that refer to the value of some *ad hoc* introduced clock predicates or explicit time variable [Ost89]. For instance the following formula of Real Time Temporal Logic (RTTL, a logic that will be discussed in Section 5.3) states the same property expressed by MTL Formula (2) above specifying bounded response time (in the formula the variable t represents the current value of the time state component).

$$\forall \text{T} ((\mathsf{p} \wedge t = \text{T}) \Rightarrow \text{F}(\mathsf{q} \wedge t \leq \text{T} + d))$$

Dealing with different time granularities Once suitable constructs are available to denote in a quantitatively precise way the time distance among events and the length of time intervals, then the problem may arise of describing systems that include several components that evolve, possibly in a partially independent fashion, on different time scales. This is dealt in the temporal logic TRIO described above by adopting syntactic and semantic mechanisms that enable dealing with different levels of TIME GRANULARITY [CCM$^+$91]. Syntactically, temporal expressions can be labeled in such a way that they may be interpreted in different time domains: for instance, 30_D denotes 30 days whereas 3_H denotes 3 hours. They key issue is the possibility given to the user to specify a semantic mapping between time domains of different granularity; hence, the truth of a predicate at a given time value at higher (coarser) level of granularity is defined in terms of the interpretation in an interval at the lower (finer) level

associated with the value at the higher level. For instance, Figure 18 specifies that, say, working during the month of November means working from the 2$^{\text{nd}}$ through the 6$^{\text{th}}$, from the 9$^{\text{th}}$ through the 13$^{\text{th}}$, etc.


Figure 18: Interpretation of an upper-level predicate in the lower-level domain. Solid lines denote the intervals in the lower domain where the predicate holds.

As with derived TRIO temporal operators, suitable predefined mappings help the user specify a few standard situations. For instance given two temporal domains $\mathbb{T}_1$ and $\mathbb{T}_2$, such that $\mathbb{T}_1$ is coarser than $\mathbb{T}_2$, p *event in* $\mathbb{T}_1 \to \mathbb{T}_2$ means that predicate p is true in any $t \in \mathbb{T}_1$ if and only if it is true in just one instant of the interval of $\mathbb{T}_2$ corresponding to t. Similarly, p *complete in* $\mathbb{T}_1 \to \mathbb{T}_2$ means that p is true in any $t \in \mathbb{T}_1$ if and only if it is true in the whole corresponding interval of $\mathbb{T}_2$.

By this way the following TRIO formula

$$\text{Alw}_{\text{M}}(\forall emp(\mathsf{work}(emp) \Rightarrow \mathsf{get\_salary}(emp)))$$

which formalizes the sentence "every month, if an employee works, then she gets her salary" introduced in Section 3.1 is given a precise semantics by introducing the mapping of Figure 18 for predicate work, and by stating that $\mathsf{get\_salary}$ *event in* $M \to D$.

In some applicative domains having administrative, business, or financial implications, the change of time granularity is often paired with a reference to a global time calendar that evolves in a *synchronous way*. For instance, time units such as days, weeks, months and years change in a synchronized way at certain predefined time instants (e.g., midnight or new year) that are conventionally established in a global fashion.

On the contrary, when a process evolves in a way such that its composing events are related directly with one another but are unrelated with any global time scale, time distances can be expressed in a time scale with no intended reference to a global time scale: in such cases we say that time granularity is managed in an *asynchronous way*. Quite often the distinction of the two intended meanings is implicit in natural language sentences and depends on some conventional knowledge that is shared among the parties involved in the described process; thus, in the formalization stage, it needs to be made explicit.

Consider for instance the following description of a procedure for carrying out written exams: "Once the teacher has completed the explanation of the exercise,

the students must solve it within exactly three hours. Then, the teacher will collect their solutions and will publish and register the grades after three days". Clearly, the former part of the sentence must be interpreted in the asynchronous way (students have to complete their job within 180 minutes starting from the minute when the explanation ended). The latter part, however, is normally intended according to the synchronous interpretation: results will be published before midnight of the third "calendar day" following the one when the exam was held.

This notion of synchronous vs. asynchronous refinement of predicates can be made explicit by adding an indication (S for synchronous, A for asynchronous) denoting the intended mode of granularity refinement for the predicates included in the subformula. Hence the above description of the written examination procedure could be formalized by the following formula, where H stands for "hours", and D for "days":

$$\mathrm{Alw}_{\mathrm{H,A}}(\mathsf{exerciseDelivery} \Rightarrow \mathrm{Futr}(\mathsf{solutionCollect}, 3))$$
$$\wedge$$
$$\mathrm{Alw}_{\mathrm{D,S}}(\mathsf{exerciseDelivery} \Rightarrow \mathrm{Futr}(\mathsf{gradesPublication}, 3))$$

To the best of our knowledge only few other languages in the literature approach the granularity problem in a formal way [BB06, Rom90]. Among these [Rom90] addresses the problem both for space and time in formal models of geographic data processing requirements.

Dense Time Domains and the Non-Zenoness Property. The adoption of a dense, possibly continuous time domain allows one to model asynchronous systems where the occurrence of distinct, independent events may be at time instants that are arbitrarily close. As a consequence, ZENO behaviors, where for instance an unbounded number of events takes place in a bounded time interval, become possible and must be ruled out by means of suitable axioms or through the adoption of *ad hoc* underlying semantic assumptions. The axiomatic description of non-Zenoness is immediate for a first order, metric temporal logic like MTL or TRIO, when it is applied to simple entities like predicates or variables ranging over finite domains. It can be more complicated when non-Zenoness must be specified in the most general case of variables that are real-valued functions of time [GM01].

Informally, a predicate is non-Zeno if it has finite variability, i.e., its truth value changes a finite number of times over any finite interval. Correspondingly, a general predicate P can be constrained to be non-Zeno by requiring that there always exists a time interval before or after every time instant, where P is constantly true or it is constantly false. This constraint can be expressed by the following TRIO formula (see [HR04, LWW07] for formulations in other similar logics):

$$\mathrm{Alw}((\mathrm{UpToNow}(\mathsf{P}) \vee \mathrm{UpToNow}(\neg\mathsf{P})) \wedge (\mathrm{NowOn}(\mathsf{P}) \vee \mathrm{NowOn}(\neg\mathsf{P}))) \quad (3)$$

The additional notion of non-Zeno *interval-based* predicate is introduced to model a property or state that holds continuously over time intervals of length strictly greater than zero. Suppose, for instance, that the "occupied state" for the resource in the resource manager example is modeled in the specification through a predicate occ; to impose that occ be an interval-based (non-Zeno) predicate, one can introduce, in addition to Formula (3), the following TRIO axiom (which eliminates the possibility of occ being true in isolated time instants).

$$\text{Alw}((\mathsf{occ} \Rightarrow \text{UpToNow}(\mathsf{occ}) \vee \text{NowOn}(\mathsf{occ})) \wedge$$
$$(\neg\mathsf{occ} \Rightarrow \text{UpToNow}(\neg\mathsf{occ}) \vee \text{NowOn}(\neg\mathsf{occ})))$$

A complementary category of non-Zeno predicates corresponds to properties that hold at *isolated time points*, and therefore can naturally model instantaneous events. If, in the resource manager specification, predicate hpr represents the issue of a high priority request, it can be constrained to be a point-based predicate by introducing the following formula in addition to Axiom (3).

$$\text{Alw}(\text{UpToNow}(\neg\mathsf{hpr}) \wedge \text{NowOn}(\neg\mathsf{hpr}))$$

Finally, non-Zenoness for a time dependent variable T (representing for instance the current temperature in a thermostat application) ranging over an uncountable domain D essentially coincides with T being piecewise analytic,[19] as a function of time. Analyticity is a quite strong "smoothness" requirement on functions which guarantees that the function intersects any constant line only finitely many times over any finite interval. Hence, any formula of the kind $T = \mathsf{v}$, where v is a constant value in D, is guaranteed to be non-Zeno according to the previous definitions for predicates. Formally, non-Zenoness for T can be constrained by the following TRIO formula (where $\mathsf{r}, \mathsf{l} : \mathbb{R} \to D$ are functions that are analytic at 0).

$$\text{Alw}(\exists d > 0 : \forall t : 0 < t < d \Rightarrow (\text{Dist}(T = \mathsf{r}(t), t) \wedge \text{Dist}(T = \mathsf{l}(t), -t)))$$

In [GM06a] it is shown that the adoption of a small set of predefined categories of specification items like the point- and interval-based predicates outlined above can make the modeling of real-time hybrid systems quite systematic and amenable to automated verification.

Future and Past Operators. While the Linear Temporal Logic LTL, as originally proposed by Pnueli [Pnu77] to study the correctness of programs, has only future operators, one may consider additional modalities for the past tense, e.g., P (for *previous*) as the operator corresponding in the past to the next operator X, or O (for *once*) as opposed to F, S (for *since*) as the past

---

[19]A function is analytic at a given point if it possesses derivatives of all orders and agrees with its Taylor series about that point [Wei, Kno96]. It is piecewise analytic if it is analytic over finitely many contiguous (open) intervals.

version of the *until* operator U, etc. The question then arises, whether the past operators are at all necessary (i.e., if they actually increase the expressiveness of the logic) or useful in practice (i.e., if there are significant classes of properties that can be described in a more concise and transparent way by using also past operators than by using future operators only).

Concerning the question of expressiveness, it is well known from [GPSS80] that LTL with past operators does not add expressive power to future-only LTL. Moreover, the separation theorem by Gabbay [Gab87] allows for the elimination of past operators, producing an LTL formula to be evaluated in the initial instant only: therefore, LTL with past operators is said to be *initially equivalent* to future-only LTL [Eme90].[20]

On the other hand, it is widely recognized that the extension of LTL with past operators [Kam68] allows one to write specifications that are easier, shorter, and more intuitive [LPZ85]. A customary example, taken from [Sch02], is the specification: *Every alarm is due to a fault*, which, using the *globally* operator G and the *previously* operator O (*once*), may be very simply written as:

$$G(\mathsf{alarm} \Rightarrow O\,\mathsf{fault})$$

whereas the following is one of the simplest LTL versions of the same specification, using the *until* operator.

$$\neg(\neg\mathsf{fault}\,\mathsf{U}\,(\mathsf{alarm} \wedge \neg\mathsf{fault}))$$

In [LMS02], it has been shown that the elimination of past operators may yield an exponential growth of the length of the derived formula.

These expressiveness results change significantly when we consider logics interpreted over dense time domains. In general, past operators add expressive power when the time domain is dense, even if we consider mono-infinite time lines such as $\mathbb{R}_{\geq 0}$. For instance, [BCM05] shows that, over the reals, propositional MTL with past operators is strictly more expressive than its future-only version. The question of the expressiveness of past operators over dense time domains was first addressed, and shown to differ from the discrete case, in [AH92a, AH93].

Branching-Time Temporal Logic. As discussed in Section 3.3, in BRANCHING-*time temporal logic* every time instant may split into several future ones and therefore formulas are interpreted over *trees* of states; such trees represent all possible computations of the modeled system. The branching in the interpretation structure naturally represents the NONDETERMINISTIC nature of the model, which may derive from some intrinsic feature of the device under construction or from some feature of the stimuli coming from the environment with which the device interacts. When interpreting a branching temporal logic formula at some current time, the properties asserted for the future may be evaluated with

---

[20] As it is customary in the literature, we consider one-sided infinite time discrete domains (i.e., $\mathbb{N}$). The bi-infinite case (i.e., $\mathbb{Z}$) is much less studied [PP04].

reference to *all* future computations (i.e., branches of the state tree) starting from the current time or only to *some* of them. Therefore, branching time temporal logic possesses modal operators that allow one to quantify universally or existentially over computations starting from the current time.

The Computation Tree Logic (CTL) [EH86] has operators that are similar to LTL, except that every temporal connective must be preceded by a *path quantifier*: either E (which stands for *there exists a computation*, sometimes also denoted with the quantification symbol $\exists$) or A (*for all computations*, also $\forall$). With reference to the usual resource manager example, the formula below asserts that in every execution a low priority request (predicate lpr) will be eventually followed by the resource being occupied (predicate occ) in some of the evolutions following the request:

$$AG\,(\mathsf{lpr} \Rightarrow EF\,\mathsf{occ})$$

while the following formula asserts that there exists a computation of the resource manager where all low priority requests are certainly (i.e., in every possible successive evolution) eventually followed by the resource being occupied:

$$EG\,(\mathsf{lpr} \Rightarrow AF\,\mathsf{occ})$$

These examples, though very simple, show that in branching time temporal logics temporal and path quantifiers may interact in quite a subtle way.

Not surprisingly, branching temporal logic has been extended in a METRIC version (TCTL, timed CTL) by adding to its operators quantitative time parameters, much in the same way MTL extends Linear Temporal Logic [ACD93, HNSY94].

We refer the reader to [Var01] for a deep analysis of the mutual pros and cons of linear time versus branching time logics.

Interval-Based Temporal Logics. All temporal logics we have considered so far adopt time *points* as the fundamental entities: every state is associated with a time instant and formulas are interpreted with reference to some time instant. By contrast, the so-called *interval temporal logics* assume time *intervals*, rather than time instants, as the original temporal entity, while time points, if not completely ignored, are considered as derived entities.

In principle, from a purely conceptual viewpoint, choosing intervals rather than points as the elementary time notion may be considered as a matter of subjective preference, once it is acknowledged that an interval may be considered as a set of points, while, on the other hand, a point could be viewed as a special case of interval having null length [Koy92]. In formal logic, however, apparently limited variations in the set of operators may make a surprisingly significant difference in terms of expressiveness and complexity or decidability of the problems related with analysis and verification. Over the years, interval temporal logics have been a quite rich research field, producing a mass of formal notations with related analysis and verification procedures and tools.

A few relevant ones are: the Interval-based Temporal Logic of Schwartz et al. [SMV83], the Interval Temporal Logic of Moszkowski [Mos83, Mos86], the Duration Calculus of Chachoen et al. [CHR91], the Metric Interval Temporal Logic (MITL) of Alur et al. [AFH96].

Among them, Duration Calculus (DC) refers to a CONTINUOUS LINEAR sequence of time instants as the basic interpretation structure. The significant portions of the system state are modeled by means of suitable functions from time (i.e., from the nonnegative reals) to Boolean values, and operators measuring accumulated durations of states are used to provide a METRIC over time. For instance, in our resource manager example, the property that the resource is never occupied for more than 100 time units without interruption (except possibly for isolated instants) would be expressed with the DC formula:

$$\Box(\lceil \mathsf{occ} \rceil \Rightarrow \ell \leq 100)$$

where $\lceil \mathsf{occ} \rceil$ is a shorthand for $\int \mathsf{occ} = \ell \wedge \ell > 0$, which formalizes the fact that the predicate occ stays true continually (except for isolated points) over an interval of length ℓ.

Another basic operator of Duration Calculus (and of several other interval logics as well) is the *chop* operator ; (sometimes denoted as $^\frown$). Its purpose it to join two formulas predicating about two different intervals into one predicating about two adjacent intervals. For example, if we wanted to formalize the property that any client occupies the resource for at least 5 time units, we could use the chop operator as follows:

$$\Box(\lceil \neg \mathsf{occ} \rceil ; \lceil \mathsf{occ} \rceil ; \lceil \neg \mathsf{occ} \rceil \Rightarrow \ell > 5)$$

where the symbol ℓ in the right-hand side of the implication now refers to the length of the overall interval, obtained by composition through the *chop* operator.

Duration Calculus also embeds an underlying semantic assumption of finite variability for state functions that essentially corresponds to the previously discussed non-ZENO requirement: each (Boolean-valued) interpretation must have only finitely many discontinuity points in any finite interval.

5.2.2 Explicit-Time Logics

Another category of descriptive formalisms adopts a "timestamp" EXPLICIT view of time. This is typically done by introducing an *ad hoc* feature (e.g., a variable that represents the current time, or a time-valued function providing a timestamp associated with every event occurrence). In this section we focus on the distinguishing features of Lamport's Temporal Logic of Actions (TLA) [Lam94], and Alur and Henzinger's Timed Propositional Temporal Logic (TPTL) [AH94]. Other relevant examples of explicit-time logics are the Real Time Logic (RTL) of Mok et al. [JM86] and Ostroff's Real-Time Temporal Logic (ESM/RTTL) [Ost89] (which will be presented in the context of the dual language approach in Section 5.3).

Temporal Logic of Actions. TLA formulas are interpreted over LINEAR, DISCRETE state sequences, and include variables, first order quantification, predicates and the usual modal operators F and G to refer to some or all future states. While basic TLA does not have a QUANTITATIVE treating of time, in [AL94] Abadi and Lamport show how to introduce a distinguished state variable *now* with a CONTINUOUS domain, representing the current time, so that the specification of temporal properties consists of formulas predicating explicitly on the values of *now* in different states, thus describing its expected behavior with respect to the events taking place.

With reference to the resource manager example, to formally describe the behavior in case of a low-priority request an action lpr would be introduced, describing the untimed behavior of this request. An *action* is a predicate about two states, whose values are denoted by unprimed and primed variables, for the current and next state, respectively. Therefore, the untimed behavior of an accepted low-priority request would simply be to change the value of the state of the resource (indicated by a variable *res*) from free to occupied, as in the following definition.

$$\mathsf{lpr} \triangleq res = \text{free} \ \wedge \ res' = \text{occ}$$

Then, the timed behavior associated with this action would be specified by setting an upper bound on the time taken by the action, specifying that the action must happen within 2 time units whenever it is continuously enabled. Following the scheme in [AL94], a *timer* would be defined by means of two formulas (which we do not report here for the sake of brevity: the interested reader can find them in [AL94]). The first one defines predicate $\text{MaxTime}(t)$, which holds in all states whose timestamp (represented by the state variable *now*) is less than or equal the absolute time t. The second formula defines predicate $\text{VTimer}(t, A, \delta, v)$, where A is an action, δ is a delay, v is the set of all variables, and t is a state variable representing a timer. Then, $\text{VTimer}(t, A, \delta, v)$ holds if and only if either action A is not currently enabled and t is ∞, or A is enabled and t is $now + \delta$ (and it will stay so until either A occurs, or A is disabled, see [AL94, Sec. 3] for further details).

Finally, the timed behavior of low-priority requests would be defined by the following action lpr^t, where T_{gr} is a state variable representing the maximum time within which action lpr must occur.

$$\mathsf{lpr}^t \triangleq \mathsf{lpr} \ \wedge \ \text{VTimer}(T_{\text{gr}}, \mathsf{lpr}, 2, v) \ \wedge \ \text{MaxTime}(T_{\text{gr}})$$

More precisely, the formula above states that after action lpr is enabled, it must occur before time surpasses value $now + 2$.

It is interesting to discuss how TLA solves the problem of ZENO behaviors. Zeno behaviors are possible because TLA formulas involving time are simply satisfied by behaviors where the variable *now*, being a regular state variable, does not change value. There are at least two mechanisms to ensure non-Zenoness. The first, simpler one introduces explicitly in the specification the requirement

that time always advances, by the following formula NZ.

$$\text{NZ} \triangleq \forall t \in \mathbb{R} : \text{F}(now > t)$$

An alternative *a posteriori* approach, which we do not discuss in detail, is based on a set of theorems provided in [AL94] to infer the non-Zenoness of specifications written in a certain canonical form, after verifying some semantic constraints regarding the actions included in the specification.

It is worth noticing that also in TLA, like in other temporal logics discussed above, two consecutive states may refer to the same time instant, so that the logic departs from the notion of time inherited from classical physics and from traditional dynamical system theory. In every timed TLA specification, it is thus customary to explicitly introduce a formula that states the separation of time-advancing steps from ordinary program steps (see [AL94] for further details). This approach is somewhat similar in spirit to that adopted in TTM/RTTL. which is presented in Section 5.3.

Timed Propositional Temporal Logic. The TPTL logic by Alur and Henzinger represents a quite interesting example of how a careful choice of the operators provided by a temporal logic can make a great difference in terms of expressiveness, decidability, and complexity of the verification procedures. TPTL may be roughly described as a "half-order" logic, in that it is obtained from propositional LINEAR time logic by adding variables that refer to time, and allowing for a *freeze quantification* operator: for a variable x, the freeze quantifier (denoted as x.) bounds the variable x to the time when the sub-formula in the scope of the quantification is evaluated. One can think of it as the analogue, for logic languages, of clock resets in timed automata (see Section 5.1.1). The freeze quantifier is combined with the usual modal operators F and G: if $\phi(x)$ is a formula in which variable x occurs free, then formula $\text{F}x.\phi(x)$ asserts that there is some future instant, with some absolute time k, such that $\phi(k)$ will hold in that instant; similarly, $\text{G}x.\phi(x)$ asserts that $\phi(h)$ will hold in any future instant, h being the absolute time of that instant.

The familiar property of the resource manager, that any low priority resource request is satisfied within 100 time units would be expressed in TPTL as follows.

$$\text{G}x.\,(\textsf{lpr} \Rightarrow \text{F}y.\,(\textsf{occ} \wedge y < x + 100))$$

In [AH94] the authors show that the logic is decidable over DISCRETE time, and define a doubly exponential DECISION PROCEDURE for it; in [AH92b] they prove that adding ordinary first order quantification on variables representing the current time, or adding past operators to TPTL, would make the decision procedure of the resulting logic non-elementary. Therefore they argue that TPTL constitutes the "best" combination of expressiveness and complexity for a temporal logic with METRIC on time.

5.2.3 Algebraic Formalisms

Algebraic formalisms are descriptive formal languages that focus on the *axiomatic* and *calculational* aspects of a specification. In other words, they are based on axioms that define how one can symbolically derive consequences of basic definitions [Bae04, Bae03]. From a software engineering viewpoint, this means that the emphasis is on *refinement* of specifications (which is formalized through some kind of algebraic *morphism*).

In algebraic formalisms devoted to the description of concurrent activities, the basic behavior of a system is usually called *process*. Hence, algebraic formalisms are often named with the term *process algebras*. A process is completely described by a set of (abstract) events occurring in a certain order. Therefore, a process is also called a *discrete event system*.

In order to describe concurrent and reactive systems, algebraic formalisms usually provide a notion of *parallel composition* among different, concurrently executing, processes. Then, the semantics of the global system is fully defined by applications of the transformation axioms of the algebra on the various processes. Such a semantics — given axiomatically as a set of transformations — is usually called *operational semantics*, not to be confused with operational formalisms (see Section 5.1).

Untimed Process Algebras. Historically, the first process algebraic approaches date back to the early work by Bekič [Bek71] and to Milner's comprehensive work on the Calculus of Communicating Systems (CCS) formalism [Mil80, Mil89]. Basically, they aimed at extending the axiomatic semantics for sequential programs to concurrent processes. In this section, we focus on Communicating Sequential Processes (CSP), another popular process algebra, introduced by Hoare [Hoa78, Hoa85] and subsequently developed into several formalisms. As usual, we refer the reader to [BPS01] for a more detailed and comprehensive presentation of process algebras, and to the historical surveys [Bae04, Bae03].

Communicating Sequential Processes are a process algebra based on the notion of *communication* between processes. The basic process is defined by the sequences of events it can generate or accept; to this end the $\rightarrow$ operator is used, which denotes a sequence of two events that occur in order. Definitions are typically recursive, and infinite behaviors can consequently arise However, a pre-defined process *SKIP* always terminates as soon as it is executed. In the following examples we denote primitive events by lowercase letters, and processes by uppercase letters.

Processes can be unbounded in number, and parametric with respect to numeric parameters, which renders the formalism very expressive. We exploit this fact in formalizing the usual resource manager example (whose complete CSP specification is shown in Table 4) by allowing an unbounded number of pending high-priority requests, similarly to what we did with Petri nets in Section 5.1.2.

In CSP two *choice* operators are available. One is *external* choice, denoted by the box operator $\Box$; this is basically a choice where the process that is

actually executed is determined by the first (prefix) event that is available in the environment. In the resource manager example, external choice is used to model the fact that a *FREE* process can stay idle for one transition (behaving as process P_N), or accept a high-priority request or a low-priority one (behaving as processes P_H and P_L, respectively). On the other hand, *internal* choice, denoted by the $\sqcap$ operator, models a nondeterministic choice where the process chooses between one of two or more possible behaviors, independently of externally generated events. In the resource manager example, the system's process WG internally chooses whether to skip once or twice before granting the resource to a low-priority request. A special event, denoted by τ, is used to give a semantics to internal choices: the τ event is considered invisible outside the process in which it occurs, and it leads to one of the possible internal choices.

Concurrently executing processes are modeled through the parallel composition operator $\|$.[21] In our example, we represent the occupied resource by a parallel composition of an OCC process and a counter $CNT(k)$ counting the number of pending high-priority requests. The former process turns back to behaving as a *FREE* process as soon as there are no more pending requests. The latter, instead, reacts to release and high-priority request events. In particular, it signals the number of remaining enqueued processes by issuing the parametric event enqueued!k (which is received by an OCC process, as defined by the incoming event enqueued?k of OCC).

$$
\begin{aligned}
\mathsf{FREE} \quad &= \quad \square_{k \in \{\mathsf{H},\mathsf{L},\mathsf{N}\}} \mathsf{P}_k \\
\mathsf{P_N} \quad &= \quad \mathit{SKIP} \longrightarrow \mathsf{FREE} \\
\mathsf{P_H} \quad &= \quad \mathsf{hpr} \longrightarrow \mathsf{P_O} \\
\mathsf{P_O} \quad &= \quad \mathsf{OCC} \ _{\{\mathsf{enqueued}\}} \| _{\{\mathsf{enqueued},\mathsf{rel},\mathsf{hpr}\}} \ \mathsf{CNT}(0) \\
\mathsf{OCC} \quad &= \quad \mathsf{enqueued}?0 \longrightarrow \mathsf{FREE} \quad \square \quad \mathsf{enqueued}?k : \mathbb{N}_{>0} \longrightarrow \mathsf{OCC} \\
\mathsf{CNT}(-1) \quad &= \quad \mathit{SKIP} \\
\mathsf{CNT}(k) \quad &= \quad \mathsf{rel} \longrightarrow \mathsf{DEQ}(k) \quad \square \quad \mathsf{hpr} \longrightarrow \mathsf{CNT}(k+1) \\
\mathsf{DEQ}(k) \quad &= \quad \mathsf{enqueued}!k \longrightarrow \mathsf{CNT}(k-1) \\
\mathsf{P_L} \quad &= \quad \mathsf{lpr} \longrightarrow \mathsf{WG} \\
\mathsf{WG} \quad &= \quad \mathsf{WG}_1 \sqcap \mathsf{WG}_2 \\
\mathsf{WG}_1 \quad &= \quad \mathit{SKIP} \ ; \ \mathsf{P_O} \\
\mathsf{WG}_2 \quad &= \quad \mathit{SKIP} \ ; \ \mathit{SKIP} \ ; \ \mathsf{P_O}
\end{aligned}
$$

Table 4: The resource manager modeled through CSP.

Let us now discuss the characteristics of the process algebraic models in

---

[21] $\mathsf{P}_{1\,A}\|_B\mathsf{P}_2$ denotes the parallel composition of processes P_1 and P_2 such that P_1 only engages in events in A, P_2 only engages in events in B, and they both synchronize on events in $A \cap B$.

general — and CSP in particular — with respect to the time modeling issues presented in Section 3.

- Basic process algebras usually have no QUANTITATIVE notion of time, defining simply an *ordering* among different events. In particular, time is typically DISCRETE [Bae04]. Variants of this basic model have been proposed to introduce metric and/or dense time; we discuss them in the remainder of this section.

- The presence of the silent transition τ is a way of modeling NONDETER-MINISTIC behaviors; in particular, the nondeterministic internal choice operator $\sqcap$ is based on the τ event.

- Even if process algebras include nondeterministic behaviors, their semantics is usually defined on LINEAR time models. There are two basic approaches to formalize the semantics of a process algebra: the *operational* one has been briefly discussed above; for the *denotational* one we refer the interested reader to [Sch00].

- The parallel COMPOSITION operation is a fundamental primitive of process algebras. The semantics which is consequently adopted for concurrency is either based on *interleaving* or it is *truly asynchronous*. Whenever interleaving concurrency is chosen, it is possible to represent a process by a set of classes of equivalent linear *traces* (see the timed automata subsection of Section 5.1.1). Therefore, the semantics of the parallel composition operator can be expressed solely in terms of the other operators of the algebra; the rule that details how to do this is called *expansion theorem* [Bae04]. On the contrary, whenever a truly asynchronous concurrency model is chosen no expansion theorem holds, and the semantics of the parallel composition operator is not reducible to that of the other operators.

- Processes described by algebraic formalisms may include DEADLOCKED *behaviors* where the state does not advance as some process is blocked. Let us consider, for instance, the following process P_i, which internally chooses whether to execute $hpr \rightarrow P_i$ or $lpr \rightarrow P_i$:

$$P_i \quad = \quad hpr \longrightarrow P_i \ \sqcap \ lpr \longrightarrow P_i$$

Process P_i may *refuse* an lpr event offered by the environment, if it internally (i.e., independently of the environment) chooses to execute $hpr \rightarrow P_i$. In such a case, P_i would deadlock. It is therefore the designer's task to prove *a posteriori* that a given CSP specification is deadlock-free.

Among other popular process algebras, let us just mention the Algebra of Communicating Processes (ACP) [BW90] and other approaches based on the integration of data description into process formalization, the most widespread approach being probably that of LOTOS [vEVD89, Bri89].

Timed Process Algebras. Quantitative time modeling is typically introduced in process algebras according to the following general schema, presented and discussed by Nicollin and Sifakis in [NS91]. First of all, each process is augmented with an *ad hoc* variable that EXPLICITLY represents time and can be CONTINUOUS. Time is global and all cooperating processes are synchronized on it.

Then, each process's evolution consists of a sequence of two-phase steps. During the first phase, an arbitrarily long — but *finite* — sequence of events occurs, while time does not change; basically, this evolution phase can be fully described by ordinary process algebraic means. During the second phase, instead, the time variable is incremented while all the other state variables stay unchanged, thus representing time progressing; all processes also SYNCHRONOUSLY update their time variables by the same amount, which can possibly be infinite (divergent behavior).

Time in such a timestamp model is usually called *abstract* to denote the fact that it does not correspond to concrete or physical time. Notice that several of the synchronous operational formalisms, e.g., those presented in Section 5.1.1, can also be described on the basis of such a time model. For instance, in synchronous abstract machines *à la* Esterel [BG92] the time-elapsing phase corresponds implicitly to one (discrete) time unit.

Assuming the general time model above, the syntax of process algebras is augmented with constructs allowing one to EXPLICITLY refer to QUANTITATIVE time in the description of a system. This has been first pursued for CSP in [RR88], and has been subsequently extended to most other process algebras. We refer the reader to [BM02, NS91, Bae03] — among others — for more references, while briefly focusing on Timed CSP (TCSP) in the following example.

Example 5 (Timed CSP). The CSP language has been modified [DS95, Sch00] by extending a minimal set of operators to allow the user to refer to metric time. In our resource manager example (whose complete Timed CSP specification is shown in Table 5), we only consider two metric constructs: the special process *WAIT* and the so-called timed timeout $\rhd^t$.

The former is a quantitative version of the untimed *SKIP*: *WAIT* t is a process which just delays for t time units. We use this to model explicitly the acceptance of a low-priority request, which waits for two time units before occupying the resource (note that we modified the behavior with respect to the untimed case, by removing the nondeterminism in the waiting time).

The timed timeout $\rhd^t$ a modification of the untimed timeout $\rhd$ (not presented in the previous CSP example). The semantics of a formula $\mathsf{P} \rhd^t \mathsf{Q}$ is that of a process that behaves as P if any of P's initial events occurs within t time units; otherwise, it behaves as Q after t time units. In the resource manager example, we exploit this semantics to prescribe that the resource cannot be occupied continuously for longer than 100 time units: if no release (rel) or high-priority request (hpr) events occur within 100 time units, the process $\mathsf{CNT}(k)$ is timed out and the process DEQ is forcefully executed.

Finally, it is worth discussing how TCSP deals with the problem of ZENO

$$\begin{aligned}
\mathsf{FREE} &= \square_{k \in \{\mathsf{H,L,N}\}} \mathsf{P}_k \\
\mathsf{P_N} &= \mathit{SKIP} \longrightarrow \mathsf{FREE} \\
\mathsf{P_H} &= \mathsf{hpr} \longrightarrow \mathsf{P_O} \\
\mathsf{P_O} &= \mathsf{OCC}\ _{\{\mathsf{enqueued}\}} \|_{\{\mathsf{enqueued,rel,hpr}\}}\ \mathsf{CNT}(0) \\
\mathsf{OCC} &= \mathsf{enqueued?0} \longrightarrow \mathsf{FREE}\ \square\ \mathsf{enqueued?}k : \mathbb{N}_{>0} \longrightarrow \mathsf{OCC} \\
\mathsf{CNT}(-1) &= \mathit{SKIP} \\
\mathsf{CNT}(k) &= (\mathsf{rel} \longrightarrow \mathsf{DEQ}(k) \square \mathsf{hpr} \longrightarrow \mathsf{CNT}(k+1)) \rhd^{100} \mathsf{DEQ}(k) \\
\mathsf{DEQ}(k) &= \mathsf{enqueued!}k \longrightarrow \mathsf{CNT}(k-1) \\
\mathsf{P_L} &= \mathsf{lpr} \longrightarrow \mathit{WAIT}\ 2 \longrightarrow \mathsf{P_O}
\end{aligned}$$

Table 5: The resource manager modeled through Timed CSP.

behaviors. The original solution of TCSP (see [DS95]) was to rule out Zeno processes *a priori* by requiring that any two consecutive actions be separated by a fixed delay of δ time units, thus prohibiting simultaneity altogether. This solution has the advantage of being simple and of totally ruling out problems of Zenoness; on the other hand, it forcefully introduces a discretization in behavior description, and it yields complications and lack of uniformity in the algebra axioms. Therefore, subsequent TCSP models have abandoned this strong assumption by allowing for simultaneous events and arbitrarily short delays. Consequently, the non-Zenoness of any given TCSP specification must be checked explicitly *a posteriori*.

Several ANALYSIS and VERIFICATION techniques have been developed for, and adapted to, process algebraic formalisms. For instance, let us just mention the FDR2 refinement checker [Ros97], designed for CSP, and the LTSA toolset [MK99] for the analysis of dual-language models combining process-algebraic descriptions with labeled transition systems.

5.3 Dual Language Approaches

The dual language approach, as stated in the introduction of Section 5.2, combines an operational formalism, useful for describing the system behavior in terms of states and transitions, with a descriptive notation suitable for specifying its properties. It provides a methodological support to the designer, in that it constitutes a unified framework for requirement specification, design, and verification. Although a dual language approach often provides methods and tools for VERIFICATION (e.g., for model checking), we point out that effectiveness or efficiency of verification procedures are not necessarily a direct consequence of the presence of two, heterogeneous notations (an operational and a descriptive one), but can derive from other factors, as the case of SPIN, discussed below, shows. In recent years a great number of frameworks to specify, design and

verify critical, embedded, real-time systems have been proposed, which may be considered as applications of the dual language approach. As usual we limit ourselves to mention the most significant features of a few representative cases.

The TTM/RTTL Framework

The work of Ostroff [Ost89] is among the first ones addressing the problem of formal specification, design, and verification of real-time systems by pursuing a dual language approach. It proposes a framework based on Extended State Machines and Real-Time Temporal Logic (ESM/RTTL). In later works, ESM have been extended to Timed Transition Models (TTM) [Ost90, Ost99].

The operational part of the framework (TTM) associates transitions with lower and upper bounds, referred to the value of a global, DISCRETE time clock variable. We briefly discussed the time model introduced by this formalism in Section 5.1.1.

Here, let us illustrate TTM through the usual resource manager example. Figure 19 represents a system similar to the Timed Petri net example of Section 5.1.2: the number of low-priority requests is not counted, while that of high-priority ones is. Each transition is annotated with lower and upper bounds, a *guard*, and a variable update rule. For instance, the transition rel_2 can be taken whenever the guard $occ > 1$ evaluates to true; notice that we exploit an integer-valued state variable to count the number of pending high-priority requests. The effect of rel_2 is to update the occ variable by incrementing it. Finally, when rel_2 becomes enabled, it *must* be taken within a maximum of 100 clock ticks, *unless* the state is left (and possibly re-entered) by taking another (non tick) enabled transition (such as hpr_2, which is always enabled, since it has no guard).
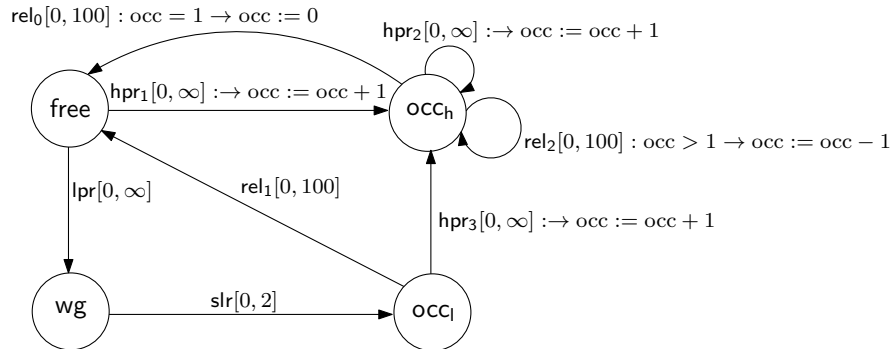

Figure 19: A resource manager modeled through a Timed Transition Model.

The descriptive part of the TTM/RTTL framework (RTTL) is based on Manna and Pnueli's temporal logic: it assumes LINEAR time and it adopts the usual operators of future-only propositional LTL. Real-time (i.e., QUANTITA-TIVE) temporal properties are expressed by means of (in)equalities on simple

arithmetic expressions involving the clock variable, as discussed in Section 5.2.1. For instance, the familiar requirement that a low priority request is followed, within 100 time units, by the resource being occupied would be expressed as follows.

$$\forall T\,((\mathsf{lpr} \wedge t = T) \Rightarrow F(\mathsf{occ} \wedge t \leq T + 100))$$

RTTL formulas are interpreted over TTM *trajectories*, i.e., sequences of states corresponding to TTM computations: [Ost89] provides both a proof system and VERIFICATION procedures based on reachability analysis techniques.

The TTM/RTTL framework is also supported by the StateTime TOOLSET [Ost97], which in turn relies on the STeP tool [BBC+00].

Model Checking Environments

The SPIN model checking environment [Hol03] is based, for the operational part, on Büchi automata, which are edited by the designer using a high-level notation called ProMeLa. The syntax of ProMeLa closely resembles that of the C programming language (and therefore is — perhaps deceptively — amenable to C programmers) and, in addition to the traditional constructs for sequential programming, provides features like parallel processes, communication channels, nondeterministic conditional instructions. The descriptive notation is plain future-only LTL, with the known limitations concerning the possibility to express complex properties and quantitative time constraints already pointed out in Section 5.2.1. Model checking in SPIN is performed by translating the LTL formula expressing the required property into a Büchi automaton and then checking that the languages of the two automata (that obtained from the ProMeLa program and the one coming from the LTL formula) are disjoint. It is therefore apparent that the distinction between the operational and the descriptive parts is maintained only in the user interface for methodological purposes, and it blurs during verification.

UPPAAL [LPY97] is another prominent framework supporting model-checking in a dual language approach. The operational part consists of a network of timed automata combined by the CCS parallel composition operator, and it provides both synchronous communication and asynchronous communication. The descriptive notation uses CTL in a restricted form, allowing only formulas of the kind $AG\phi$, $AF\phi$, $EG\phi$, $EF\phi$, and $AG(\phi \Rightarrow AF\psi)$, where ϕ and ψ are "local" formulas, i.e., Boolean expressions over state predicates and integer variables, and clock constraints.

Other Dual Language Approaches

Among the numerous other dual language frameworks [JM94] we mention [FMM94], which combines timed Petri nets and the TRIO temporal logic: it provides a systematic procedure for translating any timed Petri net into a set of TRIO axioms that characterize its behavior, thus making it possible to derive required properties of the Petri net within the TRIO proof system.

[FM02] introduces a real-time extension of the Object Constraint Language (OCL, [WK99]), which is a logic language that allows users to state (and verify through model checking) properties of transitions of UML state diagrams (which, as mentioned in Section 5.1.1, are a variation of Harel's Statecharts), especially temporal ones.

6   Conclusions

In computer science, unlike other fields of science and engineering, the modeling of time is often restricted to the formalization and analysis of specific problems within particular application fields, if not entirely abstracted away. In this paper we have analyzed the historical and practical reasons of this fact; we have examined various categories under which formalisms to analyze timing aspects in computing can be classified; then we surveyed — with no attempt at exhaustiveness, but with the goal of conceptual completeness — many of such formalisms; in doing so we analyzed and compared them with respect to the above categories.

The result is a quite rich and somewhat intricate picture of different but often tightly connected models, certainly much more variegate than the way time modeling is usually faced in other fields of science and engineering. As in other cases, in this respect, too, computing science has much to learn from other, more established, cultural fields of engineering, but also the converse is true [GM06b].

Perhaps, the main lesson we can extract from our study is that *despite the common understanding that time is a basic, unique conceptual entity, there are "many notions of time" in our reasoning; this is reflected in the adoption of different formal models when specifying and analyzing any type of system where timing behavior is of any concern.*

In some sense the above claim could be seen as an application of a principle of relativity to the abstractions required by modern — heterogeneous — system design. Whereas traditional engineering could comfortably deal with a unique abstract model of time as an independent "variable" flowing in an autonomous and immutable way to which all other system's variables had to be related, the advent of computing and communication technologies, with elaboration speeds that are comparable with the light's speed produced, and perhaps imposed, a fairly sharp departure from such a view:

- Often a different notion of time must be associated with different system's components. This may happen not only because the various components (possibly social organizations) are located in different places and their evolution may take place at a speed such that it is impossible to talk about "system state at time t", but also because the various components may have quite different nature — typically, a controlled environment and a controller subsystem based on some computing device — with quite different dynamics.

66

- In particular, even inside the same computing device, it may be necessary to distinguish between an "internal time", defined and measured by device's clock, and an "external time", which is the time of the environment with which the computing apparatus must interact and synchronize. The consequence of this fact is that often, perhaps in a hidden way, two different notions of time coexist in the same model (for instance, the time defined by the sequence of events and the time defined by a more or less explicit variable — a clock — whose value may be recorded and assigned just like other program variables).

- A different abstraction on time modeling may be useful depending on the type of properties one may wish to analyze: for instance, in some cases just the ordering of events matters, whereas in other cases a precise quantitative measure of the distance among them is needed. As a consequence many different mathematical approaches have been pursued to comply with the various modeling needs, the distinction between discrete and continuous time domains being only "the tip of the iceberg" of this issue.

Whether future evolutions will produce a better unification of the present state of the art or even more diversification and specialization in time modeling is an open and challenging question.

References

[ACD93]   Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking in dense real-time. *Information and Computation*, 104(1):2–34, 1993.

[ACH$^+$95]   Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995.

[ACHH93]   Rajeev Alur, Costas Courcoubetis, Thomas A. Henzinger, and Pei-Hsin Ho. Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In Robert L. Grossman, Anil Nerode, Anders P. Ravn, and Hans Rischel, editors, *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*, pages 209–229. Springer-Verlag, 1993.

[ACM02]   Eugene Asarin, Paul Caspi, and Oded Maler. Timed regular expressions. *Journal of the ACM*, 49(2):172–206, 2002.

[AD94]   Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

[AFH96]   Rajeev Alur, Tomás Feder, and Thomas A. Henzinger. The benefits
            of relaxing punctuality. *Journal of the ACM*, 43(1):116–146, 1996.

[AH92a]   Rajeev Alur and Thomas A. Henzinger. Back to the future: To-
            wards a theory of timed regular languages. In *Proceedings of
            the 33rd Annual Symposium on Foundations of Computer Science
            (FOCS'92)*, pages 177–186. IEEE Computer Society Press, 1992.

[AH92b]   Rajeev Alur and Thomas A. Henzinger. Logics and models of real
            time: A survey. In *Real Time: Theory in Practice*, volume 600 of
            Lecture Notes in Computer Science, pages 74–106. Springer-Verlag,
            1992.

[AH93]   Rajeev Alur and Thomas A. Henzinger. Real-time logics: Complex-
            ity and expressiveness. *Information and Computation*, 104(1):35–
            77, 1993.

[AH94]   Rajeev Alur and Thomas A. Henzinger. A really temporal logic.
            Journal of the ACM, 41(1):181–204, 1994.

[AH96]   Myla Archer and Constance L. Heitmeyer. Mechanical verifica-
            tion of timed automata: a case study. In *Proceedings of the IEEE
            Real Time Technology and Applications Symposium*, pages 192–203,
            1996.

[AK95]   Rajeev Alur and Robert P. Kurshan. Timing analysis in COSPAN.
            In Rajeev Alur, Thomas A. Henzinger, and Eduardo D. Sontag, ed-
            itors, *Proceedings of the 3rd DIMACS/SYCON Workshop on Ver-
            ification and Control of Hybrid Systems*, volume 1066 of *Lecture
            Notes in Computer Science*, pages 220–231. Springer-Verlag, 1995.

[AL94]   Martín Abadi and Leslie Lamport. An old-fashioned recipe for real
            time. *ACM Transactions on Programming Languages and Systems*,
            16(5):1543–1571, 1994.

[AM04]   Rajeev Alur and P. Madhusudan. Decision problems for timed au-
            tomata: A survey. In Marco Bernardo and Flavio Corradini, editors,
            *Revised Lectures from the International School on Formal Meth-
            ods for the Design of Computer, Communication and Software Sys-
            tems: Formal Methods for the Design of Real-Time Systems (SFM-
            RT'04)*, volume 3185 of *Lecture Notes in Computer Science*, pages
            1–24. Springer-Verlag, 2004.

[Ant00]   Panos J. Antsaklis, editor. *Special Issue on Hybrid Systems: Theory
            and Applications*, volume 88 of *Proceedings of the IEEE*. IEEE
            Press, 2000.

[Arc00]   Myla Archer. TAME: Using PVSstrategies for special-purpose the-
            orem proving. *Annals of Mathematics and Artificial Intelligence*,
            29(1–4):139–181, 2000.

[Asa04]   Eugene Asarin. Challenges in timed languages: from applied theory to basic theory. *Bulletin of the EATCS*, 83:106–120, 2004. (Column: concurrency).

[Bae03]   J. C. M. Baeten. Over thirty years of process algebra: Past, present and future. In L. Aceto, Z. Ésik, W. J. Fokkink, and A. Ingólfsdóttir, editors, *Process Algebra: Open Problems and Future Directions*, volume NS–03–3 of *BRICS Notes Series*, pages 7–12. 2003.

[Bae04]   J. C. M. Baeten. A brief history of process algebra. Technical Report CSR 04–02, Department of Mathematics and Computer Science, Technische Universiteit Eindhoven, 2004.

[BB94]   Grady Booch and Doug Bryan. *Software Engineering with ADA*. Addison-Wesley, 1994.

[BB06]   Alan Burns and Gordon Baxter. Time bands in systems structure. In D. Besnard, C. Gacek, and C. B. Jones, editors, *Structure for dependability*. Springer, 2006.

[BBC$^+$00] Nikolaj S. Bjørner, Anca Browne, Michael Colón, Bernd Finkbeiner, Zohar Manna, Henny B. Sipma, and Tomás E. Uribe. Verifying temporal properties of reactive systems: A STeP tutorial. *Formal Methods in System Design*, 16(3):227–270, 2000.

[BBM98]   Michael S. Branicky, Vivek S. Borkar, and Sanjoy K. Mitter. A unified framework for hybrid control: Model and optimal control theory. *IEEE Transactions on Automatic Control*, 43(1):31–45, 1998.

[BCM05]   Patricia Bouyer, Fabrice Chevalier, and Nicolas Markey. On the expressiveness of TPTL and MTL. In R. Ramanujam and Sandeep Sen, editors, *Proceedings of the 25th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'05)*, volume 3821 of *Lecture Notes in Computer Science*, pages 432–443. Springer-Verlag, 2005.

[BD91]   Bernard Berthomieu and Michel Diaz. Modeling and verification of time dependent systems using time Petri nets. *IEEE Transactions on Software Engineering*, 17(3):259–273, 1991.

[BDW00]   Tom Bienmüller, Werner Damm, and Hartmut Wittke. The STATEMATE verification environment – making it real. In *Proceedings of the 12th International Conference on Computer Aided Verification (CAV'00)*, volume 1855 of *Lecture Notes in Computer Science*, pages 561–567. Springer-Verlag, 2000.

[Bek71]   Hans Bekič. Towards a mathematical theory of processes. Technical Report Technical Report TR 25.125, IBM Laboratory, Wien, 1971. Published in [Bek84].

[Bek84]   Hans Bekič. Programming languages and their definition. In C. B. Jones, editor, *Selected Papers by Hans Bekič*, volume 177 of *Lecture Notes in Computer Science*. Springer-Verlag, 1984.

[BG92]   Gérard Berry and Georges Gonthier. The Esterel synchronous programming language: Design, semantics, implementation. *Science of Computer Programming*, 19(2):87–152, 1992.

[BGO+04] Marius Bozga, Susanne Graf, Ileana Ober, Iulian Ober, and Joseph Sifakis. The IF toolset. In Marco Bernardo and Flavio Corradini, editors, *Revised Lectures from the International School on Formal Methods for the Design of Computer, Communication and Software Systems: Formal Methods for the Design of Real-Time Systems (SFM-RT'04)*, volume 3185 of *Lecture Notes in Computer Science*, pages 237–267. Springer-Verlag, 2004.

[BL74]   Walter S. Brainerd and Lawrence H. Landweber. *Theory of Computation*. John Wiley & Sons, 1974.

[BM02]   J. C. M. Baeten and C. A. Middelburg. *Process Algebra with Timing*. Monographs in Theoretical Computer Science. Springer-Verlag, 2002.

[BMN00]   Pierfrancesco Bellini, Riccardo Mattolini, and Paolo. Nesi. Temporal logics for real-time system specification. *ACM Computing Surveys*, 32(1):12–42, March 2000.

[BPS01]   Jan A. Bergstra, Alban Ponse, and Scott A. Smolka, editors. *Handbook of Process Algebra*. Elsevier, 2001.

[Bri89]   Ed Brinksma, editor. *Information Processing Systems — Open Systems Interconnection — LOTOS — A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour*. ISO, 1989. ISO 8807:1989.

[BS03]   Egon Börger and Robert F. Stärk. *Abstract State Machines. A Method for High-Level System Design and Analysis*. Springer-Verlag, 2003.

[BW90]   Jos C. M. Baeten and W. P. Weijland. *Process Algebra*, volume 18 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1990.

[BY04]   Johan Bengtsson and Wang Yi. Timed automata: Semantics, algorithms and tools. In Jörg Desel, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Lectures on Concurrency and Petri Nets, Advances in Petri Nets (from ACPN'03)*, volume 3098 of *Lecture Notes in Computer Science*, pages 87–124. Springer-Verlag, 2004.

[CCM⁺91]   Edoardo Corsetti, Ernani Crivelli, Dino Mandrioli, Angelo Morzenti, Angelo Montanari, Pierluigi San Pietro, and Elena Ratto. Dealing with different time scales in formal specifications. In *In Proceedings of the 6th International Workshop on Software Specification and Design*, pages 92–101, 1991.

[Cer93]   Antonio Cerone. *A Net-Based Approach for Specifying Real-Time Systems*. PhD thesis, Università degli Studi di Pisa, Dipartimento di Informatica, 1993. TD-16/93.

[CGP00]   Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. *Model Checking*. MIT Press, 2000.

[CHR91]   Zhou Chaochen, C. A. R. Hoare, and Anders P. Ravn. A calculus of duration. *Information Processing Letters*, 40(5):269–276, 1991.

[CHR02]   Franck Cassez, Thomas A. Henzinger, and Jean-François Raskin. A comparison of control problems for timed and hybrid systems. In Claire Tomlin and Mark R. Greenstreet, editors, *Proceedings of the 5th International Workshop on Hybrid Systems: Computation and Control (HSCC'02)*, volume 2289 of *Lecture Notes in Computer Science*, pages 134–148. Springer-Verlag, 2002.

[CL99]   Christos G. Cassandras and Stephane Lafortune. *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, 1999.

[CMS99]   Antonio Cerone and Andrea Maggiolo-Schettini. Time-based expressivity of time Petri nets for system specification. *Theoretical Computer Science*, 216(1–2):1–53, 1999.

[Coo04]   Matthew Cook. Universality in elementary cellular automata. *Complex Systems*, 15:1–40, 2004.

[DK05]   Pedro R. D'Argenio and Joost-Pieter Katoen. A theory of stochastic systems (parts i and ii). *Information and Computation*, 203(1):1–74, 2005.

[DS95]   Jim Davies and Steve Schneider. A brief history of timed CSP. *Theoretical Computer Science*, 138:243–271, 1995.

[EH86]   E. Allen Emerson and Joseph Y. Halpern. "Sometimes" and "not never" revisited: On branching versus linear time temporal logic. *Journal of the ACM*, 33(1):151–178, 1986.

[Eme90]   E. Allen Emerson. Temporal and modal logic. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 996–1072. Elsevier Science, 1990.

[EPLF03]   Hans-Erik Eriksson, Magnus Penker, Brian Lyons, and David Fado. *UML 2 Toolkit*. John Wiley & Sons, 2003.

[FM02]   Stephan Flake and Wolfgang Mueller. An OCL extension for real-time constraints. In T. Clark and J. Warmer, editors, *Object Modeling with the OCL*, volume 2263 of *Lecture Notes in Computer Science*, pages 150–171. Springer-Verlag, 2002.

[FMM94]   Miguel Felder, Dino Mandrioli, and Angelo Morzenti. Proving properties of real-time systems through logical specifications and Petri net models. *IEEE Transactions on Software Engineering*, 20(2):127–141, 1994.

[FPR08a]   Carlo A. Furia, Matteo Pradella, and Matteo Rossi. Automated verification of dense-time MTL specifications via discrete-time approximation. In Jorge Cuéllar, Tom Maibaum, and Kaisa Sere, editors, *Proceedings of the 15th International Symposium on Formal Methods (FM'08)*, volume 5014 of *Lecture Notes in Computer Science*, pages 132–147. Springer-Verlag, May 2008.

[FPR08b]   Carlo A. Furia, Matteo Pradella, and Matteo Rossi. Comments on "temporal logics for real-time system specification". *ACM Computing Surveys*, 2008. Accepted for publication (April 2008).

[FR06]   Carlo A. Furia and Matteo Rossi. Integrating discrete- and continuous-time metric temporal logics through sampling. In Eugene Asarin and Patricia Bouyer, editors, *Proceedings of the 4th International Conference on the Formal Modeling and Analysis of Timed Systems (FORMATS'06)*, Lecture Notes in Computer Science. Springer-Verlag, 2006.

[Fra86]   Nissim Francez. *Fairness*. Monographs in Computer Science. Springer-Verlag, 1986.

[Gab87]   Dov M. Gabbay. The declarative past and imperative future. In Behnam Banieqbal, Howard Barringer, and Amir Pnueli, editors, *Proceeding of Temporal Logic in Specification (TLS'87)*, volume 398 of *Lecture Notes in Computer Science*, pages 409–448, Altrinchamm, UK, April 1987. Springer-Verlag.

[GJM02]   Carlo Ghezzi, Mehdi Jazayeri, and Dino Mandrioli. *Fundamentals of Software Engineering*. Prentice Hall, 2nd edition, 2002.

[GM01]   Angelo Gargantini and Angelo Morzenti. Automated deductive requirements analysis of critical systems. *ACM Transactions on Software Engineering and Methodology*, 10(3):255–307, 2001.

[GM06a]   Angelo Gargantini and Angelo Morzenti. Automated verification of continuous time systems by discrete temporal induction. In *Proceedings of the 13th International Symposium on Temporal Representation and Reasoning (TIME'06)*. IEEE Computer Society Press, 2006.

[GM06b]   Carlo Ghezzi and Dino Mandrioli. The challenges of software engineering education. In P. Inverardi and M. Jazayeri, editors, *Proceedings of the ICSE 2005 Education Track*, volume 4309 of *Lecture Notes in Computer Science*, pages 115–127. Springer-Verlag, 2006.

[GMM90]   Carlo Ghezzi, Dino Mandrioli, and Angelo Morzenti. TRIO: A logic language for executable specifications of real-time systems. *The Journal of Systems and Software*, 12(2):107–123, 1990.

[GMM99]   Angelo Gargantini, Dino Mandrioli, and Angelo Morzenti. Dealing with zero-time transitions in axiom systems. *Information and Computation*, 150(2):119–131, 1999.

[GMMP91]   Carlo Ghezzi, Dino Mandrioli, Sandro Morasca, and Mauro Pezzè. A unified high-level Petri net formalism for time-critical systems. *IEEE Transactions on Software Engineering*, 17(2):160–172, 1991.

[GPSS80]   Dov M. Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. On the temporal basis of fairness. In *Proceedings of the 7th Annual ACM Symposium on Principles of Programming Languages (POPL'80)*, pages 163–173. ACM Press, 1980.

[GTBF03]   Holger Giese, Matthias Tichy, Sven Burmester, and Stephan Flake. Towards the compositional verification of real-time UML designs. In *Proceedings of ESEC/SIGSOFT FSE 2003*, pages 38–47, 2003.

[Har87]   David Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8(3):231–274, 1987.

[Hen96]   Thomas A. Henzinger. The theory of hybrid automata. In *Proceedings of the 11th Annual Symposium on Logic in Computer Science (LICS)*, pages 278–292. IEEE Computer Society Press, 1996.

[Hen98]   Thomas A. Henzinger. It's about time: Real-time logics reviewed. In Davide Sangiorgi and Robert de Simone, editors, *Proceedings of the 9th International Conference on Concurrency Theory (CONCUR'98)*, volume 1466 of *Lecture Notes in Computer Science*, pages 439–454. Springer-Verlag, 1998.

[HHWT97]   Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. HYTECH: A model checker for hybrid systems. *International Journal on Software Tools for Technology Transfer*, 1(1–2), 1997.

[HKPV98]   Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What's decidable about hybrid automata? *Journal of Computer and System Sciences*, 57(1):94–124, 1998.

[HLN+90]   David Harel, Hagi Lachover, Amnon Naamad, Amir Pnueli, Michal Politi, Rivi Sherman, Aharon Shtull-Trauring, and Mark B. Trakhtenbrot. STATEMATE: A working environment for the development

of complex reactive systems. *IEEE Transactions on Software Engineering*, 16(4):403–414, 1990.

[HM96]   Constance L. Heitmeyer and Dino Mandrioli, editors. *Formal Methods for Real-Time Computing*. John Wiley & Sons, 1996.

[HMU00]   John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison Wesley, 2nd edition, 2000.

[HN96]   David Harel and Amnon Naamad. The STATEMATE semantics of statecharts. *ACM Transactions on Software Engineering and Methodology*, 5(4):293–333, 1996.

[HNSY94] Thomas A. Henzinger, Xavier Nicollin, Joseph Sifakis, and Sergio Yovine. Symbolic model checking for real-time systems. *Information and Computation*, 111(2):193–244, 1994.

[Hoa78]   C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666—677, 1978.

[Hoa85]   C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.

[Hol03]   Gerard J. Holzmann. *The SPIN Model Checker: Primer and Reference Manual*. Addison-Wesley, 2003.

[HPSS87] David Harel, Amir Pnueli, Jeanette P. Schmidt, and Rivi Sherman. On the formal semantics of statecharts. In *Proceedings of the 2nd IEEE Symposium on Logic in Computer Science (LICS'87)*, pages 54–64, 1987.

[HR04]   Yoram Hirshfeld and Alexander Rabinovich. Logics for real time: Decidability and complexity. *Fundamenta Informormaticae*, 62(1):1–28, 2004.

[JM86]   Farnam Jahanian and Aloysius K. Mok. Safety analysis of timing properties in real-time systems. *IEEE Transactions on Software Engineering*, 12(9):890–904, 1986.

[JM94]   Farnam Jahanian and Aloysius K. Mok. Modechart: A specification language for real-time systems. *IEEE Transactions on Software Engineering*, 20(12):933–947, 1994.

[Kam68]   Johan Anthony Willem Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, University of California at Los Angeles, 1968.

[KB04]   Randy H. Katz and Gaetano Borriello. *Contemporary Logic Design*. Prentice Hall, 2nd edition, 2004.

[Kha95]   Hassan Khalil. *Nonlinear Systems*. Prentice-Hall, 2nd edition, 1995.

[Kno96]   Konrad Knopp. *Theory of Functions, Parts 1 and 2, Two Volumes Bound as One*, chapter 8, pages 83–111. Dover, 1996.

[Koy90]   Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.

[Koy92]   Ron Koymans. (real) time: A philosophical perspective. In J. W. de Bakker, Cornelis Huizing, Willem P. de Roever, and Grzegorz Rozenberg, editors, *Proceedings of the REX Workshop: "Real-Time: Theory in Practice"*, volume 600 of *Lecture Notes in Computer Science*, pages 353–370. Springer-Verlag, 1992.

[KP92]   Yonit Kesten and Amir Pnueli. Timed and hybrid statecharts and their textual representation. In Jan Vytopil, editor, *Proceedings of the 2nd International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'92)*, volume 571 of *Lecture Notes in Computer Science*, pages 591–620. Springer-Verlag, 1992.

[Kri63]   Saul Aaron Kripke. Semantical analysis of modal logic I. *Zeitschrift fur Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.

[Lam80]   Leslie Lamport. "Sometime" is sometimes "not never": On the temporal logic of programs. In *Proceedings of the 7th ACM Symposium on Principles of Programming Languages (SIGPLAN-SIGACT)*, pages 174–185. ACM Press, 1980.

[Lam83]   Leslie Lamport. What good is temporal logic? In R. E. A. Mason, editor, *Proceedings of the 9th IFIP World Congress*, volume 83 of *Information Processing*, pages 657–668. North-Holland, 1983.

[Lam94]   Leslie Lamport. The temporal logic of actions. *ACM Transations on Programming Languages and Systems*, 16(3):872–923, 1994.

[LMS02]   François Laroussinie, Nicolas Markey, and Philippe Schnoebelen. Temporal logic with forgettable past. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science (LICS'02)*, pages 383–392. IEEE Computer Society Press, 2002.

[LPY97]   Kim G. Larsen, Paul Pettersson, and Wang Yi. UPPAAL in a nutshell. *International Journal on Software Tools for Technology Transfer*, 1(1–2), 1997.

[LPZ85]   Orna Lichtenstein, Amir Pnueli, and Lenore D. Zuck. The glory of the past. In *Proceedings of 3rd Workshop on Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 196–218. Springer-Verlag, 1985.

[LV96]   Nancy Lynch and Frits W. Vaandrager. Forward and backward simulations – part II: Timing-based systems. *Information and Computation*, 128(1):1–25, 1996.

[LWW07]   Carsten Lutz, Dirk Walther, and Frank Wolter. Quantitative temporal logics over the reals: PSPACE and below. *Information and Computation*, 205(1):99–123, 2007.

[Mea55]   George H. Mealy. A method for synthesizing sequantial circuits. *Bell System Technical Journal*, 34:1045–1079, 1955.

[Men97]   Elliott Mendelson. *Introduction to Mathematical Logic*. Chapman and Hall, fourth edition, 1997.

[MF76]   P. M. Merlin and D. J. Farber. Recoverability and communication protocols: Implications of a theoretical study. *IEEE Transactions on Communications*, 24(9):1036–1043, 1976.

[MG87]   Dino Mandrioli and Carlo Ghezzi. *Theoretical Foundations of Computer Sciences*. John Wiley & Sons, 1987.

[Mil80]   Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer-Verlag, 1980.

[Mil89]   Robin Milner. *Communication and Concurrency*. Prentice Hall, 1989.

[MK99]   Jeff Magee and Jeff Kramer. *Concurrency: State Models & Java Programs*. John Wiley & Sons, 1999.

[MMG92]   Angelo Morzenti, Dino Mandrioli, and Carlo Ghezzi. A model parametric real-time logic. *ACM Transactions on Programming Languages and Systems*, 14(4):521–573, 1992.

[Moo56]   Edward F. Moore. Gedanken-experiments on sequential machines. In *Automata Studies*, volume 34 of *Annals of Mathematical Studies*, pages 129–153. Princeton University Press, 1956.

[Mos83]   Ben Moszkowski. *Reasoning about Digital Circuits*. PhD thesis, Department of Computer Science, Stanford University., 1983. Techincal Report STAN–CS–83–970.

[Mos86]   Ben Moszkowski. *Executing temporal logic programs*. Cambridge University Press, 1986.

[MP92]   Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, 1992.

[NOSY93]   Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In Robert L. Grossman, Anil Nerode, Anders P. Ravn, and Hans Rischel, editors, *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*, pages 149–178. Springer-Verlag, 1993.

[NS91]   Xavier Nicollin and Joseph Sifakis. An overview and synthesis of timed process algebras. In Kim G. Larsen and Arne Skou, editors, *Proceedings of the 3rd International Workshop on Computer Aided Verification (CAV'91)*, volume 575 of *Lecture Notes in Computer Science*, pages 376–398. Springer-Verlag, 1991.

[Odi99]   Piergiorgio Odifreddi. *Classical Recursion Theory*. North Holland, 1999.

[OL07a]   Martin Ouimet and Kristina Lundqvist. The TASM toolset: Specification, simulation, and formal verification of real-time systems. In W. Damm and H. Hermanns, editors, *Proceedings of the 19th International Conference on Computer-Aided Verification (CAV'07)*, volume 4590 of *Lecture Notes in Computer Science*, pages 126–130. Springer-Verlag, 2007.

[OL07b]   Martin Ouimet and Kristina Lundqvist. The timed abstract state machine language: Abstract state machines for real-time system engineering. In *Proceedings of the 14th International Workshop on Abstract State Machines (ASM'07)*, 2007.

[Ost89]   Jonathan S. Ostroff. *Temporal Logic for Real Time Sytems*. Advanced Software Development Series. John Wiley & Sons, 1989.

[Ost90]   Jonathan S. Ostroff. Deciding properties of timed transition models. *IEEE Transactions on Parallel and Distributed Systems*, 1(2):170–183, 1990.

[Ost92]   Jonathan S. Ostroff. Formal methods for the specification and design of real-time safety critical systems. *Journal of Systems and Software*, 18(1):33–60, 1992.

[Ost97]   Jonathan S. Ostroff. A visual toolset for the design of real-time discrete-event systems. *IEEE Transactions on Control Systems Technology*, 5(3):320–337, 1997.

[Ost99]   Jonathan S. Ostroff. Composition and refinement of discrete real-time systems. *ACM Transactions on Software Engineering and Methodology*, 8(1):1–48, 1999.

[Pap94]   Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[Per93]   Adriano Peron. Synchronous and asynchronous models for state-charts. Technical Report TD-21/93, Dipartimento di Informatica, Università di Pisa, 1993.

[Pet63]   Carl A. Petri. Fundamentals of a theory of asynchronous informa-tion flow. In *Proceedings of IFIP Congress*, pages 386–390. North Holland Publishing Company, 1963.

[Pet81]   James L. Peterson. *Petri Net theory and the Modelling of Systems.* Prentice-Hall, 1981.

[Pnu77]   Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th IEEE Symposium on Fundations of Computer Science (FOCS'77)*, pages 46–67, 1977.

[PP04]   Dominique Perrin and Jean-Éric Pin. *Infinite Words*, volume 141 of *Pure and Applied Mathematics*. Elsevier, 2004.

[PS91]   Amir Pnueli and Michal Shalev. What is in a step: On the seman-tics of statecharts. In Takayasu Ito and Albert R. Meyer, editors, *Proceedings of the International Conference on Theoretical Aspects of Computer Software (TACS'91)*, volume 526 of *Lecture Notes in Computer Science*, pages 244–264. Springer-Verlag, 1991.

[Rei85]   Wolfgang Reisig. *Petri Nets: An Introduction.* EATCS Monographs on Theoretical Computer Science. Springer-Verlag, 1985.

[RKNP04]   J. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathemat-ical Techniques for Analyzing Concurrent and Probabilistic Systems*, volume 23 of *CRM Monograph Series*. American Mathematical So-ciety, 2004.

[Rog87]   Hartley Rogers, Jr. *Theory of Recursive Functions and Effective Computability.* MIT Press, 1987.

[Rom90]   Gruia-Catalin Roman. Formal specification of geographic data pro-cessing requirements. *IEEE Transaction on Knowledge and Data Engineering*, 2(4):370–380, 1990.

[Ros97]   A. William Roscoe. *The Theory and Practice of Concurrency.* Prentice-Hall International, 1997.

[RR88]   George M. Reed and A. William Roscoe. A timed model for Com-municating Sequential Processes. *Theoretical Computer Science*, 58(1–3):249–261, 1988.

[RU71]   Nicholas Rescher and Alasdair Urquhart. *Temporal Logic.* Springer-Verlag, 1971.

[Sch00]   Steven Schneider. *Concurrent and Real-Time Systems: The CSP Approach.* John Wiley & Sons, 2000.

[Sch02]   Philippe Schnoebelen. The complexity of temporal logic model checking. In Philippe Balbiani, Nobu-Yuki Suzuki, Frank Wolter, and Michael Zakharyaschev, editors, *Proceedings of the 4th Conference on Advances in Modal Logic*, pages 393–436. King's College Publications, 2002.

[Sip05]   Michael Sipser. *Introduction to the Theory of Computation*. Course Technology, 2nd edition, 2005.

[SMV83]   Richard L. Schwartz, P. M. Melliar-Smith, and Friedrich H. Vogt. An interval logic for higher-level temporal reasoning. In *Proceedings of the 2nd Annual ACM Symposium on Principles of Distributed Computing (PODC'83)*, pages 198–212. ACM Press, 1983.

[Som04]   Ian Sommerville. *Software Engineering*. Addison Wesley, 7th edition, 2004.

[SP05]   Sigurd Skogestad and Ian Postlethwaite. *Multivariable Feedback Control: Analysis and Design*. Wiley, 2nd edition, 2005.

[TGI]   Petri nets tools database. `http://www.informatik.uni-hamburg.de/TGI/PetriNets/tools/quick.html`.

[Tho90]   Wolfgang Thomas. Automata on infinite objects. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 133–164. Elsevier Science, 1990.

[UML04]   Unified modeling language specification. Technical Report formal/04-07-02, Object Management Group, 2004.

[UML05]   UML 2.0 superstructure specification. Technical Report formal/05-07-04, Object Management Group, 2005.

[Var96]   Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Proceedings of the 8th BANFF Higher Order Workshop Conference on Logics for Concurrency: Structure versus Automata*, volume 1043 of *Lecture Notes in Computer Science*, pages 238–266. Springer-Verlag, 1996.

[Var01]   Moshe Y. Vardi. Branching vs. linear time: Final showdown. In Tiziana Margaria and Wang Yi, editors, *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'01)*, volume 2031 of *Lecture Notes in Computer Science*, pages 1–22. Springer-Verlag, 2001.

[vEVD89]   Peter H. J. van Eijk, C. A. Vissers, and Michel Diaz, editors. *The formal description technique LOTOS*. Elsevier Science, 1989.

[von94]   Michael von der Beeck. A comparison of statecharts variants. In *Proceedings of the 3rd International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*, volume 863 of *Lecture Notes in Computer Science*, pages 128–148. Springer-Verlag, 1994.

[vS00]   Arjan van der Schaft and Hans Schumacher. *An Introduction to Hybrid Dynamical Systems*, volume 251 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag, 2000.

[Wei]   Eric W. Weisstein. Real analytic function. From MathWorld – A Wolfram Web Resource. `http://mathworld.wolfram.com/RealAnalyticFunction.html`.

[Wir77]   Niklaus Wirth. Toward a discipline of real-time programming. *Communications of the ACM*, 20(8):577–583, 1977.

[WK99]   Jos B. Warmer and Anneke G. Kleppe, editors. *The Object Constraint Language*. Addison-Wesley, 1999.

[Wol94]   Stephen Wolfram. *Cellular automata and complexity*. Perseus Books Group, 1994.

[Yov97]   Sergio Yovine. Kronos: A verification tool for real-time systems. *International Journal on Software Tools for Technology Transfer*, 1(1–2), 1997.