

A Methodology for the Construction of Scheduled Systems

K. Altisen, G. Gößler, and J. Sifakis

VERIMAG, 2 av. Vignates, 38610 Gières, France
{altisen, goessler, sifakis}@imag.fr

Abstract. We study a methodology for constructing scheduled systems by restricting successively the behavior of the processes to be scheduled. Restriction is used to guarantee the satisfaction of two types of constraints: schedulability constraints characterizing timing properties of the processes, and constraints characterizing particular scheduling algorithms including process priorities, non-idling, and preemption. The methodology is based on a controller synthesis paradigm. The main results deal with the characterization of scheduling policies as safety constraints and the simplification of the synthesis process by applying a composability principle.

1 Introduction

Scheduling coordinates the execution of application and system activities, so as requirements about their temporal behavior are met. Guaranteeing correctness of schedulers is essential for the development of dependable real-time systems. In many application areas, well established theory and scheduling algorithms have been successfully applied to real-time systems development.

Existing scheduling theory is limited because it requires the system to fit into the mathematical framework of the schedulability criterion (e.g. all tasks are supposed periodic, worst case execution times are known). Studies to relax such hypotheses have been carried out but they generalize one hypothesis at a time, and no unified approach has been proposed.

To overcome limitations of scheduling theory, it is important to study its connections to specification theory and take advantage of their complementarity [8, 13, 4]. The specification based approach consists in building a timed model of the scheduled system or of an abstraction of it. Then, timed analysis tools are used either to check that the exact model meets scheduling requirements or to extract from the abstraction a scheduler [6, 10].

A major difficulty in applying this approach is the generation of the timed model from some description of the scheduling method. In fact, scheduling deals with the very dynamic nature of real-time systems, and behavior modeling requires a deep understanding of mechanisms such as priorities and preemption, as well of concepts such as urgency, idling, timeliness.

In this paper we propose a methodology for modeling scheduling algorithms that constructs compositionally the scheduled system from a global timed model based on

1. A functional description of the processes to be scheduled, their resources, and the associated synchronization and management constraints;
2. Timing requirements added to the functional description and relating in particular execution speed with the dynamics of the external environment;
3. A description of a scheduling algorithm consisting of three types of requirements about
 - (a) Fixed or dynamic priorities, for choosing between pending requests of the processes,
 - (b) Possibility of idling, meaning that the scheduler may not satisfy a pending request anticipating the satisfaction of a forthcoming higher priority request,
 - (c) Preemption, that is, for a given preemption order between processes, a process of lower priority is preempted when a process of higher priority raises a request.

In previous papers [3, 2] we have shown how a functional description can be extended into a timed one by preserving progress properties. In this paper we study a methodology for constructing a scheduled system from scheduling requirements and a timed specification of the processes to be scheduled. The methodology is based on the controller synthesis paradigm [11, 9, 1]. A scheduler is considered as a controller of the processes to be scheduled which restricts their behavior by triggering their controllable actions. The restricted behavior must respect the timing constraints of the processes as well as constraints characterizing the scheduling requirements.

We have shown in [1] how schedulers can be computed by applying a synthesis algorithm to timed automata. The synthesis algorithm computes iteratively from a constraint K characterizing scheduling requirements, the maximal *control invariant* K' , $K' \Rightarrow K$. The latter denotes the set of states from which K is guaranteed. The behavior of the scheduled system is obtained by restricting the controllable actions of the processes so as to respect the control invariant K' .

The application of synthesis techniques is limited for two reasons. First, the practical complexity of the synthesis algorithm is high even in the case of timed automata without scheduling policy constraints. Second, scheduling with preemption requires the use of automata with integrators [5] which implies that iterative computation of control invariants may not terminate.

The proposed methodology allows to decompose the global controller synthesis procedure into the application of simpler steps. At each step a control invariant corresponding to a particular class of constraints is applied to further restrict the behavior of the system of processes to be scheduled. The presented results can be summarized as follows:

1. Global scheduling requirements can be characterized by a constraint K of the form $K = K_{\text{algo}} \wedge K_{\text{sched}}$ where K_{algo} specifies a particular scheduling algorithm, and K_{sched} characterizes schedulability requirements of the processes. Furthermore, K_{algo} is a conjunction of constraints about the scheduling policy, the possibility of non-idling, and preemption;

2. A step of the method corresponds to the computation of a controller for some constraint. The control invariant corresponding to a constraint can be computed in a straightforward manner (without iterative fixpoint computation);
3. The scheduled system can be obtained by successive applications of steps restricting the process behavior by control invariants implying all the scheduling constraints, provided that some composability conditions are satisfied. In fact, the restriction by a control invariant does not necessarily preserve previously imposed control invariants.

The methodology allows an incremental construction of a scheduled system, or of an abstraction of it, if some steps fail.

The paper is composed of two sections. The first section presents basic results about control invariants and their composability. The second section shows how scheduling requirements can be expressed as constraints which are control invariants in some cases. The application of the methodology is illustrated by examples.

2 Control Invariants and Composability

2.1 Timed System

To model scheduling algorithms, we use reactive timed systems with two kinds of actions as in [1]: controllable actions that can be triggered by the scheduler, and uncontrollable actions that can be considered as internal actions of the processes to be scheduled. Controllable actions are typically resource allocations and process preemption, while uncontrollable actions are process arrival and termination.

Both controllable and uncontrollable actions are submitted to timing constraints expressed in terms of real-valued variables called *timers*. The derivatives of timers may take the values 0 or 1, as specified by a boolean vector.

Definition 2.1. (*X-constraint*). Let X be a finite set of timers, $\{x_1, \dots, x_m\}$, real-valued variables defined on the set of non-negative reals \mathbb{R}_+ . A predicate C generated by the grammar $C ::= x \# d \mid x - y \# d \mid C \wedge C \mid \neg C$, where $x, y \in X$, d is an integer, and $\# \in \{\leq, <\}$, is called a *X-constraint*.

Definition 2.2. (*Falling Edge*). Let C be a X-constraint, and b be a boolean derivative vector of $\{0, 1\}^m$. The *closed* (resp. *open*) *falling edge* of C w.r.t. b , written $\Downarrow_b C$ (resp. $\Downarrow_b^o C$) is defined as $\forall x \in \mathbb{R}_+^m$,

$$\begin{aligned} \Downarrow_b C(x) &\Leftrightarrow C(x) \wedge \exists t > 0 . \forall t' \in (0, t] . \neg C(x + t'b) \\ \Downarrow_b^o C(x) &\Leftrightarrow \neg C(x) \wedge \exists t > 0 . \forall t' \in (0, t] . C(x - t'b) . \end{aligned}$$

Example 2.1. Let $X = \{x_1, x_2\}$ be the set of real valued variables. $C = x_1 \leq 3$ and $C' = 2 \leq x_1 < 6 \wedge x_1 - x_2 \leq 4$ are X-constraints. For $b = (1, 1)$ and $b' = (1, 0)$, we have:

$$\begin{array}{l} \Downarrow_b C = (x_1 = 3) \quad \Downarrow_b^o C = false \\ \Downarrow_{b'} C = (x_1 = 3) \quad \Downarrow_{b'}^o C = false \end{array} \left| \begin{array}{l} \Downarrow_b C' = false \\ \Downarrow_{b'} C' = (x_1 - x_2 = 4) \\ \quad \wedge (x_2 < 2) \end{array} \right. \begin{array}{l} \Downarrow_b C' = (x_1 = 6) \wedge (x_2 \geq 2) \\ \Downarrow_{b'} C' = (x_1 = 6) \wedge (x_2 \geq 2) \end{array}$$

Definition 2.3. (*Timed System*). A *timed system* is

1. An untimed labeled transition system (S, A, T) where S is a finite set of control states; A is a finite vocabulary of actions partitioned into two sets of controllable and uncontrollable actions noted A^c and A^u ; $T \subseteq S \times A \times S$ is an untimed transition relation;
2. A finite set of timers $X = \{x_1, \dots, x_m\}$, as in definition 2.1;
3. A function b mapping S into $\{0, 1\}^m$. The image of $s \in S$ by b denoted b_s is a boolean derivative vector;
4. A labeling function h mapping untimed transitions of T into timed transitions: $h(s, a, s') = (s, a, g, \tau, r, s')$, where the guard g is a X -constraint; the reset $r \subseteq X$ is a set of timers to be reset; $\tau \in \{\lambda, \delta, \epsilon\}$ is an urgency type, respectively lazy, delayable, eager.

Semantics. A timed system defines a *transition graph* $(\mathcal{V}, \mathcal{E})$ constructed as follows. $\mathcal{V} = S \times \mathbb{R}_+^m$, that is, vertices (s, x) are states of the timed system.

The set $\mathcal{E} \subseteq \mathcal{V} \times (A \cup \mathbb{R}_+^*) \times \mathcal{V}$ of the edges of the graph is partitioned into three classes of edges: \mathcal{E}^c controllable, \mathcal{E}^u uncontrollable, and \mathcal{E}^t timed, corresponding respectively to the case where the label is a controllable action, an uncontrollable action, and a (strictly) positive real.

Given $s \in S$, let J be the set of indices such that $\{(s, a_j, s_j)\}_{j \in J}$ is the set of all the untimed transitions departing from s . Also let $h(s, a_j, s_j) = (s, a_j, g_j, \tau_j, r_j, s_j)$. For all $j \in J$, $((s, x), a_j, (s_j, x[r_j])) \in \mathcal{E}^c \cup \mathcal{E}^u$ iff $g_j(x)$ and $x[r_j]$ is the timer valuation obtained from x when all the timers in r_j are set to zero and the others are left unchanged.

To define \mathcal{E}^t , we use the predicate φ , called *time progress function*. The notation $\varphi((s, x), t)$ means that time can progress from state (s, x) by t .

$$\varphi((s, x), t) \Leftrightarrow \bigwedge_{j \in J} \begin{cases} \tau_j = \delta \Rightarrow \forall t' \in [0, t] . \neg \downarrow_{b_s} g_j(x + t'b_s) \wedge \\ \tau_j = \epsilon \Rightarrow \forall t' \in (0, t] . \neg \downarrow_{b_s} g_j(x + t'b_s) \end{cases}$$

If $\varphi((s, x), t)$, then $((s, x), t, (s, x + tb_s)) \in \mathcal{E}^t$ where $x + tb_s$ is the valuation obtained from x by increasing by t the timer values for which b_s elements are equal to one.

The above definition means that at control state s , time cannot progress whenever an eager transition is enabled, or beyond the falling edge of a delayable guard.

We will usually denote by TS a timed system. TS^c (resp. TS^u) represents the timed system composed of the controllable (resp. uncontrollable) transitions of TS only.

Proposition 2.1. If φ , φ^c , and φ^u are respectively the time progress functions of TS, TS^c , and TS^u then $\varphi = \varphi^c \wedge \varphi^u$.

Example 2.2. (*A Periodic Process*). Let us model a periodic non-preemptible process P as a timed system. P is of period $T > 0$ and uses the CPU for an execution time E . It also has a relative deadline of D ($D \leq T$).

As shown in Fig. 2, the timed system has three control states, s , w , and e where P is respectively sleeping, waiting for the CPU, and executing on the CPU. The actions a , b , and f stand for arrive, begin, and finish. The timer x is used to measure execution time while the timer t measures the time elapsed since the process has arrived. In all states, both timers progress. The only controllable action is b .




Fig. 2: A periodic process.

By convention, transition labels are of the form a^x, g^τ, r , where x can be u (uncontrollable) or c (controllable), and τ is an urgency type. The set r is omitted if it is empty.

Notice that since the transition b is delayable, the process might wait for a non-zero time although the CPU is free: idling is permitted. A non-idling process is modeled by changing the urgency type of the transition b to eager (see example 2.5 for further details). A preemptive periodic process is modeled in section 3.3.

2.2 Restriction and Control Invariants

Definition 2.4. (*Constraint*). Given a timed system with a set of timers X and a set of control states $\{s_1, \dots, s_n\}$, a *constraint* is a state predicate represented as an expression of the form $\bigvee_{i=1}^n s_i \wedge C_i$ where C_i is a X -constraint and s_i is (also) the boolean denoting presence at state s_i .

Definition 2.5. (*Restriction*). Let TS be a timed system and K be a constraint. The *restriction* of TS by K denoted TS/K , is the timed system TS where each guard g of a controllable transition (s, a, g, τ, r, s') , is replaced by

$$g'(x) = g(x) \wedge K(s', x[r]) .$$

Notice that in the restriction TS/K , the states reached right after execution of a controllable transition satisfy K . Moreover, it follows from the definition that $(TS/K_1)/K_2 = TS/(K_1 \wedge K_2)$.

Definition 2.6. (*Proper Invariant*). Let TS be a timed system and K be a constraint. We say that K is a *proper invariant* of TS , denoted by $TS \models \text{inv}(K)$, if K is preserved by the edges of \mathcal{E} , i.e., $\forall (s, x) . K(s, x) \Rightarrow \forall ((s, x), \gamma, (s', x')) \in \mathcal{E} . K(s', x')$.

Proper invariants, called simply invariants for closed systems, are constraints preserved by all the transitions of the system. We use the term “proper” to distinguish them from control invariants introduced in the following definition. Control invariants are constraints that are satisfied by the restricted system.

Definition 2.7. (*Control Invariant*). Let TS be a timed system and K be a constraint. K is a *control invariant* of TS if $TS/K \models \text{inv}(K)$.

Proposition 2.2. If K is a proper invariant of a timed system TS , then K is a control invariant of TS .

This property follows from the trivial observation that if TS and TS/K are initialized in K , then they have the same behavior. However, notice that control invariants are not proper invariants, in general.

Proposition 2.3. For any timed system TS and constraint K such that $TS^u \models \text{inv}(K)$, K is a control invariant of TS (i.e. $TS/K \models \text{inv}(K)$).

Proof (sketch). Assume $K(s, x)$ for some state (s, x) . To prove $TS/K \models \text{inv}(K)$ it must be shown that K is preserved in TS/K by (1) controllable, (2) uncontrollable, and (3) timed edges of TS/K . By construction of TS/K , (1) is true. From $TS^u \models \text{inv}(K)$, (2) and (3) follow.

Definition 2.8. (*Timed System of Processes*). A timed system of processes is a timed system $TS = (S, A, T, X, b, h)$ obtained by composition of processes where a process P_i is a timed system $(S_i, A_i, T_i, X_i, b_i, h_i)$. TS is the timed system of n processes $\{P_1, \dots, P_n\}$ if

$$\begin{aligned} S &= S_1 \times \dots \times S_n ; A = A_1 \cup \dots \cup A_n ; X = X_1 \cup \dots \cup X_n ; \\ \text{For } s &= (s_1 \dots s_n) \in S \text{ and } x \in X_i, b_s[x] = b_{i, s_i}[x]; \\ \text{For } s &= (s_1 \dots s_i \dots s_n), \text{ and } s' = (s_1 \dots s'_i \dots s_n) \in S, \\ \left| \begin{array}{l} t = (s, a_i, s') \in T \quad \Leftrightarrow \quad t_i = (s_i, a_i, s'_i) \in T_i \\ h(t) = (s, a_i, g_i, \tau_i, r_i, s') \Leftrightarrow h_i(t_i) = (s_i, a_i, g_i, \tau_i, r_i, s'_i) . \end{array} \right. \end{aligned}$$

We assume that processes have disjoint sets of control states, and timers. Moreover, we accept that guards are general constraints on timers and control states as in the definition 2.4.

Example 2.3. (*Mutual Exclusion*). Consider a timed system of n periodic non-preemptible processes $\{P_1, \dots, P_n\}$, instances of the generic process of Fig. 2, and the constraint

$$K_{\text{mutex}} = \bigwedge_{i \neq j} \neg e_i \vee \neg e_j$$

expressing mutual exclusion. It is trivial to check that K_{mutex} is a control invariant, as $TS^u \models \text{inv}(K_{\text{mutex}})$. In fact, K_{mutex} is time invariant and is preserved by uncontrollable transitions.

If TS is the timed system of two processes — as in Fig. 2 for which the parameters (E, T, D) are equal to $(5, 15, 15)$ and $(2, 5, 5)$, resp. — and if $K_{\text{mutex}} = \neg e_1 \vee \neg e_2$, then $TS_1 = TS/K_{\text{mutex}}$ is obtained by restricting the controllable guards g_{b_1} and g_{b_2} to

$$\begin{aligned} g'_{b_1} &= (t_1 \leq D_1 - E_1) \wedge \neg e_2 = (t_1 \leq 10) \wedge \neg e_2 \\ g'_{b_2} &= (t_2 \leq D_2 - E_2) \wedge \neg e_1 = (t_2 \leq 3) \wedge \neg e_1 . \end{aligned}$$

2.3 Control Invariants and Synthesis

Following ideas in [11], synthesis is used to partially restrict the non-determinism of a system so as it satisfies a given invariant.

Problem 2.1. (SYNTH). Solving the *synthesis problem* for a timed system TS and a constraint K amounts to giving a non-empty control invariant K' of TS which implies K, i.e. $K' \Rightarrow K$, $TS/K' \models \text{inv}(K')$.

We assume that the processes to be scheduled and their timing constraints are represented by a timed system of processes TS. Furthermore, we consider that scheduling requirements can be expressed as a constraint (safety property) K. A scheduled system can be obtained by solving the synthesis problem for TS and K, as explained in [1]. If K' is a control invariant implying K, then TS/K' describes a scheduled system.

We assume that the constraint K is in general the conjunction of two constraints $K = K_{\text{algo}} \wedge K_{\text{sched}}$. K_{algo} is an optional constraint characterizing a particular scheduling algorithm. We provide in section 3, a general framework for the decomposition of K_{algo} and the modeling of different scheduling policies.

K_{sched} expresses the fact that the timing requirements of the processes are satisfied. We consider that the processes to be scheduled are structurally timelock-free [2]. This property means that time always eventually progresses. It is implied by the fact that at any control state, if no action is enabled then time can progress, and the requirement that in any circuit of the control graph a timer is reset and tested against some positive lower bound. For example, the periodic process of example 2.2 is structurally timelock-free.

Notice that structural timelock-freedom is preserved by restriction. For timelock-free timed systems, K_{sched} can be formulated as a constraint expressing the property that each process always eventually executes some action. This property implies fairness of the scheduling algorithm.

Definition 2.9. (\diamond). Let C be a X-constraint, $s \in S$ a control state, and $k \in \mathbb{N} \cup \{\infty\}$. We will use the notation

$$(\diamond_k^s C)(x) = \exists t \in [0, k] . C(x + tb_s)$$

to express the property “eventually C within k in s ”. If the state s is clear from the context, we write \diamond_k instead of \diamond_k^s . We use $(\diamond C)(x)$ for $\exists t \geq 0 . C(x + t)$.

For a timed system of processes as in definition 2.8,

$$K_{\text{sched}} = \bigwedge_{P_i} K_{\text{sched}_i} \quad \text{where} \quad K_{\text{sched}_i} = \bigvee_{s \in S_i} s \wedge \left(\bigvee_{(s, a, s') \in T_i} \diamond g_a \right) .$$

It can be shown that in general, K_{sched} is not a control invariant. We have shown in [1] how maximal schedulers for timed automata and their schedulability constraints can be computed. The synthesis algorithm has been implemented in the KRONOS tool.

Example 2.4. (*Schedulability*). The schedulability constraint for the timed system of n periodic processes TS as in example 2.3 is

$$K_{\text{sched}} = \bigwedge_{P_i} (s_i \wedge \diamond g_{a_i} \vee e_i \wedge \diamond g_{f_i} \vee w_i \wedge \diamond g_{b_i}) .$$

We consider the timed system of two processes described in example 2.3 where the mutual exclusion constraint has been applied. We have

$$K_{\text{sched}} = \left[\begin{array}{l} s_1 \wedge t_1 \leq 15 \\ \vee e_1 \wedge x_1 \leq 5 \\ \vee w_1 \wedge t_1 \leq 10 \end{array} \right] \bigwedge \left[\begin{array}{l} s_2 \wedge t_2 \leq 5 \\ \vee e_2 \wedge x_2 \leq 2 \\ \vee w_2 \wedge t_2 \leq 3 \end{array} \right] .$$

The maximal control invariant implying K_{sched} computed by KRONOS is

$$K'_{\text{sched}} = \left[\begin{array}{l} (s_1 \wedge s_2 \wedge t_1 \leq 15 \wedge t_2 \leq 5) \\ \vee (w_1 \wedge s_2 \wedge (t_2 \leq 3 \wedge t_1 \leq 10 \vee t_2 \leq 5 \wedge t_1 \leq t_2 + 3)) \\ \vee (s_1 \wedge w_2 \wedge t_1 \leq 15 \wedge t_2 \leq 3) \\ \vee (e_1 \wedge s_2 \wedge t_2 \leq 5 \wedge x_1 \leq 5 \wedge t_1 \leq x_1 + 10 \wedge t_2 \leq x_1 + 3) \\ \vee (w_1 \wedge w_2 \wedge (t_1 \leq 8 \wedge t_2 \leq 1 \vee t_2 \leq 3 \wedge t_1 \leq t_2 + 3)) \\ \vee (s_1 \wedge e_2 \wedge t_1 \leq 15 \wedge x_2 \leq 2 \wedge t_2 \leq x_2 + 3) \\ \vee (e_1 \wedge w_2 \wedge x_1 \leq 5 \wedge t_1 \leq x_1 + 10 \wedge t_2 + 2 \leq x_1) \\ \vee (w_1 \wedge e_2 \wedge (x_2 \leq 2 \wedge t_1 \leq x_2 + 8 \wedge t_2 \leq x_2 + 1 \vee \\ x_2 \leq 2 \wedge t_1 \leq t_2 + 3 \wedge t_2 \leq x_2 + 3)) \end{array} \right] .$$

In the rest of the paper, we show how to construct control invariants for some frequently used scheduling algorithms *without* fixpoint computation.

2.4 Control Invariant Composability

Contrary to proper invariants, control invariants are not composable by conjunction. In general, it can not be inferred from $\text{TS}/K_i \models \text{inv}(K_i), i = 1, 2$ that $\text{TS}/(K_1 \wedge K_2) \models \text{inv}(K_1 \wedge K_2)$. We study a notion of control invariant composability.

Definition 2.10. (*Composable Invariant*). Let TS be a timed system and K_1 be a constraint. K_1 is a *composable invariant* of TS if for all constraints K_2 , K_1 is a control invariant of TS/K_2 (i.e. if $\text{TS}/(K_1 \wedge K_2) \models \text{inv}(K_1)$).

Proposition 2.4. Let TS be a timed system and K_1 be a constraint on TS. K_1 is a composable invariant of TS iff $\text{TS}^u \models \text{inv}(K_1)$.

Proof. Let K_1 be a composable invariant of TS. By applying definition 2.10 with $K_2 = \text{false}$, we obtain: $\text{TS}/\text{false} = \text{TS}^u \models \text{inv}(K_1)$.

Conversely, assume that $\text{TS}^u \models \text{inv}(K_1)$ and let K_2 be some constraint. We show that $\text{TS}/(K_1 \wedge K_2) \models \text{inv}(K_1)$. Let (s, x) be a state of TS such that $K_1(s, x)$. (1) If there exists a controllable edge $((s, x), a_c, (s', x'))$ in the transition graph of $\text{TS}/(K_1 \wedge K_2)$, then by definition 2.5 of restriction, $(K_1 \wedge K_2)(s', x')$, thus $K_1(s', x')$. (2) An uncontrollable edge $((s, x), a_u, (s', x'))$ of $\text{TS}/(K_1 \wedge K_2)$ is also

an uncontrollable edge of TS^u , thus $K_1(s', x')$. (3) Let $\varphi_{(K_1 \wedge K_2)}$ be the time progress function of $\text{TS}/(K_1 \wedge K_2)$. According to the property 2.1, we have

$$\varphi_{(K_1 \wedge K_2)} = \varphi_{(K_1 \wedge K_2)}^c \wedge \varphi_{(K_1 \wedge K_2)}^u = \varphi_{(K_1 \wedge K_2)}^c \wedge \varphi^u .$$

If $((s, x), t, (s, x + tb_s))$ is a timed edge of $\text{TS}/(K_1 \wedge K_2)$, then it is also a timed edge of TS^u because $\varphi_{(K_1 \wedge K_2)} = \varphi_{(K_1 \wedge K_2)}^c \wedge \varphi^u$. Thus, $K_1(s, x + tb_s)$ from $\text{TS}^u \models \text{inv}(K_1)$.

Corollary 2.1. For a timed system TS and constraints K_1 and K_2 , $\text{TS}^u \models \text{inv}(K_1)$ and $(\text{TS}/K_1)/K_2 \models \text{inv}(K_2)$ implies that $\text{TS}/(K_1 \wedge K_2) \models \text{inv}(K_1 \wedge K_2)$.

That is, if K_1 is composable and if K_2 is a control invariant of TS/K_1 then $(K_1 \wedge K_2)$ is control invariant of TS .

This corollary justifies the incremental methodology for restricting a timed system. To impose a control invariant $K_1 \wedge K_2$ on TS , if K_1 is a composable invariant of TS , the restriction by a control invariant K_2 does not destroy the invariance of K_1 .

Example 2.5. (*Non-idling Constraint*). A scheduling algorithm is said to be non-idle if the CPU cannot remain free when there is a pending request. Let us consider the timed system of n processes as in example 2.3. As $\text{TS}^u \models \text{inv}(K_{\text{mutex}})$, K_{mutex} is composable which means that K_{mutex} is a proper invariant of any system obtained by restriction of $\text{TS}_1 = \text{TS}/K_{\text{mutex}}$.

In order to model non-idling, as remarked in example 2.2, all transitions b_i must have the urgency type *eager*. The non-idling constraint $K_{\text{non-idle}}$ specifies that an enabled b_i action is fired as soon as the CPU is free.

$$K_{\text{non-idle}} = \bigvee_{P_i} (e_i \vee x_i = E_i) \vee \bigwedge_{P_j} (s_j \vee w_j \wedge t_j = 0)$$

means that in a non-idling system, if no process P_i is executing or has just finished its execution, then any process P_j is either sleeping or waiting for zero time.

It can be shown that $K_{\text{non-idle}}$ is a proper invariant of TS_1 . However, it fails to be composable, in general. For the timed system of two processes described in example 2.3, the constraint $K_{\text{non-idle}}$ becomes

$$K_{\text{non-idle}} = \left[\begin{array}{l} (e_1 \vee e_2) \vee (x_1 = 5 \vee x_2 = 2) \\ \vee (s_1 \vee w_1 \wedge t_1 = 0) \wedge (s_2 \vee w_2 \wedge t_2 = 0) \end{array} \right] .$$

Notice that $\text{TS}_1/K_{\text{non-idle}} = \text{TS}_1$, that is, restricting by $K_{\text{non-idle}}$ does not change controllable transitions of TS_1 . It is easy to check that $\text{TS}_1/(K_{\text{non-idle}} \wedge K_{\text{sched}}) \not\models \text{inv}(K_{\text{non-idle}})$: consider for instance the eager transition b_1 from the control state $(w_1 s_2)$ to $(e_1 s_2)$ with guard $g'_{b_1} = \underbrace{t_1 \leq 10}_{g_{b_1}} \wedge \underbrace{t_2 \leq 3}_{K'_{\text{sched}}}$. When the system

reaches the state $(w_1 s_2)$ with timer values $(t_1 = 0, t_2 = 4)$, the action b_1 is not enabled although the CPU is free due to the restriction $t_2 \leq 3$ imposed by K_{sched} . Thus, $K_{\text{non-idle}}$ is violated.

Imposing K_{sched} has destroyed the property of the system to be non-idle. Thus the non-idling constraint is not composable. This is a consequence of the observation that a given scheduling problem with an idling solution may have no non-idling schedule.

The notion of composability described in this section allows to apply restrictions sequentially to build a system more and more close to the correct scheduler at each step.

3 Modeling Scheduling Algorithms

Timed systems with priorities are timed systems of processes with an associated set of priority orders on actions. They have been defined and studied in [3, 2]. We show how to model scheduling algorithms by specifying a timed system with priorities and that applying priorities is equivalent to restricting by a composable invariant.

3.1 Timed Systems with Priorities

Definition 3.1. (*Priority Order*). Let $\prec \subseteq A \times (\mathbb{N} \cup \{\infty\}) \times A$ be a relation. $a_1 \prec_k a_2$ is written for $(a_1, k, a_2) \in \prec$. The relation \prec is a *priority order* if $\forall k \in \mathbb{R}_+ \cup \{\infty\}$,

1. \prec_k is a partial order;
2. $a_1 \prec_k a_2 \Rightarrow \forall k' < k . a_1 \prec_{k'} a_2$;
3. $a_1 \prec_k a_2 \wedge a_2 \prec_l a_3 \Rightarrow a_1 \prec_{k+l} a_3$.

Definition 3.2. (*Timed System with Priorities*). A timed system with priorities (TS, pr) is the timed system of processes TS equipped with a *priority rule*, i.e., a finite set of pairs $pr = \{(C^i, \prec^i)\}_i$, where \prec^i is a priority order, and C^i is a X-constraint that specifies when the priority order applies, such that

1. $C^i \wedge C^j \neq \text{false} \Rightarrow \prec^i \cup \prec^j$ is a priority order;
2. No uncontrollable action is dominated in \prec^i ;
3. $(C^i, \prec^i) \in pr$ and $(a, k, b) \in \prec^i$ imply that transitions labeled by a do not reset any timer occurring in C^i .

For each state $s \in S$, let $\{(s, a_i, s_i)\}_{i \in I}$ be the set of the transitions departing from s , and $h(s, a_i, s_i) = (s, a_i, g_i, \tau_i, r_i, s_i)$. The timed system with priorities (TS, pr) represents a timed system TS' obtained from TS by replacing the guards g_j of TS by g'_j defined as follows:

$$g'_j = g_j \wedge \bigwedge_{(C, \prec) \in pr} \left(\neg C \vee \bigwedge_{\substack{\exists i \in I. \\ a_j \prec_k a_i}} \neg \diamond_k^s g_i \right) .$$

This formula says that an action a_j is allowed if there is no transition a_i leaving s that has priority over a_j , and that will become enabled within a delay of k .

Example 3.1. (*The edf Policy*). Consider the timed system TS_1 of n non-preemptible periodic processes, on which K_{mutex} has already been applied, as in example 2.3.

We show how the basic earliest deadline first (edf, [7]) mechanism can be specified by using a priority rule. A scheduler follows an edf policy if the CPU is granted to the waiting process that is closest to its relative deadline.

The edf policy is partially specified by

$$pr_{\text{edf}}^1 = \{(D_i - t_i < D_j - t_j, \{b_j \prec_0 b_i\})\}_{i \neq j} ,$$

i.e., whenever there are two processes P_i and P_j waiting for the CPU, the action b_i has immediate priority over the action b_j if P_i is closer to its relative deadline than P_j (namely, $D_i - t_i < D_j - t_j$).

It is easy to check that pr satisfies the requirements of definition 3.2. In particular, note that the constraints $D_i - t_i < D_j - t_j$ define a partial order on the set of b_i actions. The complete specification of the edf policy is given in example 3.3.

3.2 Priorities as Restriction

We show that applying a priority rule amounts to restricting by a particular constraint. To obtain this result, we construct from (TS, pr) a timed system TS' that is *strongly equivalent* to TS , and a constraint K_{pr} such that (TS, pr) is strongly equivalent to TS'/K_{pr} . Strong equivalence means that for any state of TS there exists a state of TS' such that the transition graphs are strongly bisimilar from these states, and conversely. The construction has only a theoretical interest and is used to show that K_{pr} is a composable invariant.

Let (TS, pr) be a timed system with priorities. In order to interpret priorities on TS as a constraint, we have to identify the states reached right after firing a restricted transition.

For this we transform $TS = (S, A, T, X, b, h)$ into a strongly equivalent timed system $TS' = (S', A, T', X, b', h')$ with $S' \subseteq S \cup (S \times A)$, by iterative application of a state splitting procedure which creates for each transition a unique target control state.

For each state $s \in S$ with an incident transition of the form $t = (ss, a_j, s)$ where $ss \in S'$ and $a_j \in A$, the splitting procedure removes t and creates a new transition $t' = (ss, a_j, (s, a_j))$. t' is labeled as t with in addition a reset of a new timer z_j . Notice that in TS' the set of states reached right after the execution of a_j is characterized by $((s, a_j) \wedge z_j = 0)$. For all states $s \in S'$, we take $b'_s[z_j] = 1$.




Fig. 4: The splitting procedure.

Proposition 3.1. Let (TS, pr) be a timed system with priorities, and TS' be the result of the splitting procedure on TS . The constraint

$$K_{pr} = \bigwedge_{s \in S'} \bigwedge_{(C, \prec) \in pr} \bigwedge_{\substack{i, j \in I \\ a_j \prec_k a_i}} \left(s_j \wedge z_j = 0 \Rightarrow (\neg \diamond_k^s g_i \vee \neg C) \right),$$

is a composable invariant of TS' , and $(TS', pr) = TS'/K_{pr}$, where for a given s , $\{(s, a_i, g_i, \tau_i, r_i, s_i)\}_{i \in I}$ is the set of transitions departing from s .

Proof. Notice that K_{pr} contains all the states but the ones that would be reached by firing a transition violating the priority rule.

$(TS', pr) = TS'/K_{pr}$ is obtained immediately by comparing syntactically the result of restriction by K_{pr} with the application of the priority rule pr .

To prove composability, we show that $TS'^u \models \text{inv}(K_{pr})$. Let (s, x) be a state of TS'^u such that $K_{pr}(s, x)$. (2) If there exists an uncontrollable edge $((s, x), a_u, (s', x'))$ in TS'^u , then K_{pr} cannot contain a constraint of the form $s' \wedge z = 0 \Rightarrow \neg C \vee \neg \diamond_k^{s''} g$, since a_u is the only transition leading to s' in TS'^u . Thus, $K_{pr}(s', x')$. (3) If time can progress by $t > 0$ from (s, x) in TS'^u , then $K_{pr}(s, x + tb_s)$ obviously holds.

Corollary 3.1. Let (TS, pr) be a timed system with priorities, K be a control invariant of (TS, pr) , $((TS, pr)/K)'$ be the result of the splitting procedure on $(TS, pr)/K$, and K_{pr} the constraint associated to pr . Then $((TS, pr)/K)' \models \text{inv}(K_{pr})$.

These results say that applying a priority rule can be seen as a restriction of a strongly equivalent timed system by a control invariant. Furthermore, whenever some other control invariant K is applied to (TS, pr) , then $(TS, pr)/K$ still satisfies the priority rule pr . In some cases, the property 3.1 holds without applying the splitting procedure, as shown in the following examples.

3.3 Basic Scheduling Algorithms

Example 3.2. (*The fifo Policy*). A scheduler follows a first in first out policy (fifo) if the CPU is granted to the process that has been waiting for the longest time. For non-preemptible processes, fifo is specified by using priorities as follows

$$pr_{\text{fifo}} = \{(t_j < t_i, \{b_j \prec_0 b_i\})\}_{i \neq j}.$$

This means that whenever two processes P_i and P_j are both waiting for the CPU, b_i has priority over b_j if process P_i has been waiting for longer time than process P_j , i.e. $t_j < t_i$.

Proposition 3.2. $(TS, pr_{\text{fifo}}) = TS/K_{\text{fifo}}$, where

$$K_{\text{fifo}} = \bigwedge_{i \neq j} (w_i \wedge e_j \wedge x_j = 0 \Rightarrow t_i \leq t_j)$$

is the constraint associated with pr_{fifo} . Moreover, K_{fifo} is a composable control invariant for TS . (Proof omitted.)

Example 3.3. (*The edf Policy*). We showed in example 3.1 how to model partially the edf policy on TS as a priority rule, pr_{edf}^1 . But this specification has to be completed since in case a process P_i arrives (transition a_i) exactly when the decision to allot the CPU to another process is made, this might be wrong depending on whether P_i was taken into account or not. This confusion situation can be prevented by a priority rule ensuring that the set of waiting processes is up to date before any decision is made. Therefore, processes arrival actions a_i are given priority over b_j actions:

$$pr_{\text{edf}} = pr_{\text{edf}}^1 \cup \{(t_i = T_i, \{b_j \prec_0 a_i\})\}_{i \neq j} .$$

Let K_{edf} be the constraint associated with pr_{edf} . Thus,

$$\begin{aligned} K_{\text{edf}} = \bigwedge_{i \neq j} w_i \wedge e_j \wedge x_j = 0 &\Rightarrow D_j - t_j \leq D_i - t_i \\ \wedge s_i \wedge e_j \wedge x_j = 0 &\Rightarrow t_i \neq T_i . \end{aligned}$$

Proposition 3.3. $(\text{TS}, pr_{\text{edf}}) = \text{TS}/K_{\text{edf}}$, and K_{edf} is a composable control invariant. (Proof omitted.)

Preemptive Fixed-priority Scheduling Preemptive fixed-priority scheduling assigns the CPU according to some fixed priority order between the processes to be scheduled. If the CPU is free, the highest priority process among the waiting processes is scheduled. An arriving process can preempt a running process of lower priority.

Fig. 6 shows the model of a preemptible process. It has an additional control state p (preempted), and two more transitions: pr (preempt) and rs (resume). The timer x is stopped in control state p , i.e. $b_p[x] = 0$. Everywhere else, timers progress. The timer x_{pr} measures the time elapsed since the process has been preempted.




Fig. 6: A preemptible process.

Consider the timed system of n processes P_1, \dots, P_n as shown in fig. 6 with the given fixed priorities π_1, \dots, π_n , where $\pi_i < \pi_j$ means that P_j has priority over P_i . As before, mutual exclusion is achieved by application of K_{mutex} . We construct the scheduled system of these processes according to the preemptive policy with the priorities π_1, \dots, π_n as follows.

Process Priorities. Priorities between the processes are specified by the priority rule pr_π on the CPU allocating actions b and rs :

$$pr_\pi = \{(true, \{b_i \prec_0 b_j, b_i \prec_0 rs_j, rs_i \prec_0 b_j, rs_i \prec_0 rs_j\}), (t_j = T_j, \{b_i \prec_0 a_j, rs_i \prec_0 a_j\})\}_{\pi_i < \pi_j} .$$

The first line says that the CPU is granted — by an action b_j or rs_j — to a process P_j that has highest priority among the waiting processes. Here, the constraint that specifies when the priority order applies is *true*, since the priorities are fixed and do not depend on timer valuations. The second line guarantees that the set of waiting processes is up to date before a new process is scheduled.

It is easy to show that pr_π satisfies the definition of a priority rule.

Proposition 3.4. Let K_{pr_π} be the constraint associated to pr_π . Then, $(TS, pr_\pi) = TS/K_{pr_\pi}$, and K_{pr_π} is a composable control invariant of TS. (Proof omitted.)

Preemption. pr_π only specifies the CPU allocation policy, but not the mechanism preempting a running process, which will be enforced by a further constraint

$$K_{\text{pmtn}} = \bigwedge_i (p_i \wedge x_{pr_i} = 0 \Rightarrow \exists j . \pi_j > \pi_i \wedge t_j = 0) .$$

Notice that for given process priorities π_1, \dots, π_n , the term $\exists j . \pi_j > \pi_i \wedge t_j = 0$ is a X-constraint. The constraint means that a process P_i must not take the pr_i action unless there is a higher priority process P_j that has just arrived. It implies that a running process is preempted as soon as a process of higher priority arrives. Immediately after that, since the a_i actions are eager, the CPU is assigned to a waiting process according to pr .

K_{pmtn} is a control invariant for TS, thus from corollary 2.1, $K_{\text{pmtn}} \wedge K_\pi$ is also a control invariant of TS. But K_{pmtn} is not composable, and neither is $K_{\text{pmtn}} \wedge K_\pi$.

Example 3.4. (*The rms Policy with Preemption*). The algorithm of preemptive rate-monotonic scheduling (rms, [7]) assigns to each process a fixed priority such that processes with shorter period have higher priority, i.e., $T_i > T_j \Rightarrow \pi_i < \pi_j$.

The invariant K_π can be obtained from pr_π as before. As remarked above, $K_{\text{pmtn}} \wedge K_\pi$ is not composable. However, the rms policy makes the scheduled system $(TS, pr_\pi)/K_{\text{pmtn}}$ nearly deterministic since π defines a total order. Therefore, there is no need to further restrict the system — it is either schedulable or not.

4 Conclusion

This work aims at bridging the gap between scheduling theory and timed systems specification and analysis. From the general idea that a scheduler is a controller of the scheduled processes, we elaborate a methodology for the construction of a scheduled system. The methodology is illustrated on periodic processes but it can be applied to arbitrary systems of structurally timelock-free processes.

A contribution of this work is the decomposition of scheduling requirements into classes of requirements that can be expressed as safety constraints. We believe that the decomposition allows better understanding of scheduling problems and clarification of the differences between the two approaches. Scheduling theory studies sufficient conditions guaranteeing K_{sched} for particular scheduling algorithms characterized by some K_{algo} . On the contrary, timed systems specification

and analysis have focused so far on the extraction of behaviors satisfying K_{sched} from a global model.

This work relates controller synthesis by means of the notion of control invariant, to a methodology for constructing a scheduled system satisfying given requirements. The existence of composable control invariants allows the automatic application of the corresponding synthesis steps. Not surprisingly, finding control invariants for schedulability is the hard problem that deserves further investigation. Possible directions are the development of specific synthesis algorithms or the use of constructive correctness techniques as in [2].

This work is developed in the framework of a project on real-time systems modeling and validation. We have applied the methodology to the description of the ceiling protocol [12] and are currently developing tools supporting the methodology.

References

1. K. Altisen, G. Göbller, A. Pnueli, J. Sifakis, S. Tripakis, and S. Yovine. A framework for scheduler synthesis. In *IEEE RTSS 1999 proceedings*, 1999.
2. S. Bornot, G. Göbller, and J. Sifakis. On the construction of live timed systems. In *TACAS 2000*, volume 1785 of *LNCS*. Springer-Verlag, 2000.
3. S. Bornot and J. Sifakis. On the composition of hybrid systems. In *International NATO School on "Verification of Digital and Hybrid Systems"*, LNCS. Springer Verlag, 1997.
4. P.A. Hsiung, F. Wang, and Y.S. Kuo. Scheduling system verification. In *TACAS'99*, volume 1597 of *LNCS*. Springer-Verlag, 1999.
5. Y. Kesten, A. Pnueli, J. Sifakis, and S. Yovine. Integration graphs: A class of decidable hybrid systems. *Information and Computation*, 736, 1992.
6. H.-H. Kwak, I. Lee, A. Philippou, J.-Y. Choi, and O. Sokolsky. Symbolic schedulability analysis of real-time systems. In *IEEE RTSS 1998 proceedings*, 1998.
7. C.L. Liu and J.W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, 20(1), 1973.
8. Z. Liu and M. Joseph. Specification and verification of fault-tolerance, timing, and scheduling. *ACM Transactions on Programming Languages and Systems*, 21(1):46–89, 1999.
9. O. Maler, A. Pnueli, and J. Sifakis. On the synthesis of discrete controllers for timed systems. In *STACS'95*, volume 900 of *LNCS*. Springer Verlag, 1995.
10. P. Niebert and S. Yovine. Computing optimal operation schemes for chemical plants in multi-batch mode. In *Hybrid Systems, Computation and Control*, volume 1790 of *LNCS*. Springer Verlag, 2000.
11. P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete event systems. *Journal of Control and Optimization*, 25(1), 1987.
12. L. Sha, R. Rajkumar, and J. P. Lehoczky. Priority inheritance protocols: An approach to real-time synchronization. *IEEE Transactions on Computers*, 39(9), 1990.
13. S. Vestal. Modeling and verification of real-time software using extended linear hybrid automata. In *Fifth NASA Langley Formal Methods Workshop*, 2000.