HAL
open science

Automatic generation of safe handlers for multi-task systems

Eric Rutten, Hervé Marchand

To cite this version:

Eric Rutten, Hervé Marchand. Automatic generation of safe handlers for multi-task systems. Journal of Embedded Computing, 2009, 3 (4), pp.255-276. 10.3233/JEC-2009-0097 . inria-00510023

HAL Id: inria-00510023
https://inria.hal.science/inria-00510023v1

Submitted on 25 Aug 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Automatic Generation of Safe Handlers for Multi-Task Systems

Eric Rutten *

INRIA Grenoble - Rhône-Alpes

Inovallée, 655 av. de l'Europe, MONTBONNOT, 38334 ST ISMIER Cedex, FRANCE

tel:+33 4 76 61 55 50, fax:+33 4 76 61 52 52, Eric.Rutten@inria.fr

Hervé Marchand

INRIA Rennes-Bretagne Atlantique

Campus de beaulieu, 35042 RENNES Cedex, FRANCE

tel: +33 2 99847509, fax:+33 2 99847171, Herve.Marchand@inria.fr

September 17, 2009

Abstract

We are interested in the programming of real-time embedded control systems, such as in robotic, automotive or avionic systems. They are designed with multiple tasks, each with multiple modes. It is complex to design task handlers that control the switching of activities in order to insure safety properties of the global system. We propose a model of tasks in terms of transition systems, designed especially with the purpose of applying existing discrete controller synthesis techniques. This provides us with a systematic methodology, for the automatic generation of safe task handlers, with the support of synchronous languages and associated tools.

Keywords: Real-time systems, safe design, discrete control synthesis, synchronous programming.

*corresponding author

1 Multi-task systems

1.1 Real-time control systems

The design of real-time control systems involves different aspects like e.g., device (sensor and actuator) management, numerical computation, and the discrete, event-based switching between modes of activity. Such systems are designed with a number of tasks, executed in a cyclic way, each computing a closed loop control. An example of such a control system architecture is ORCCAD [5]. Each task can have different modes, which can be different phases (initialization, nominal, exception handling), or versions implementing the same functionality with different levels of quality (e.g., computation approximation), and cost (e.g., computation time, energy, memory, bandwidth of communication, side-effects on the environment) [21]. Elaborate control systems offer multiple functionalities, managed by a number of such tasks, which can be invoked in sequence or parallel. They may then have to share processors or devices, and to be managed w.r.t. their interactions. Instances of such systems in applications are:

- control systems (e.g. robotics) ; a notion of quality of service in these systems is that the numerical computations involved, e.g., matrix computation for kinematic model update, can be performed with different levels of accuracy, by making approximations e.g., by limiting the development of series, or avoiding the computation of terms that have a negligible value under some conditions. The control laws can also differ by the way energy is consumed by actuators or side-effects on the environment.

- telecommunication and portable devices like cellular telephones present many functionalities; e.g., signal processing as in voice recognition or encodings for transmission, can have different implementations according to energy spending, volume of data, use of bandwidth.

- automobile embedded equipments, in every of their activation configurations, have a given energy consumption: this latter has to be bounded, because of the battery, so that the autonomy of the car is not jeopardized. At the same time, priority tasks (e.g., related to safety) have to be respected.

It is complex to handle tasks and to control the switching of modes in order to insure properties characterizing the safety of the control systems. This safety is often critical in the framework of transportation systems, energy production and, in general, in all systems with a strong and possibly hazardous interaction with people or the environment. The properties to be ensured can concern the safe access to resources, according to particular constraints of sequencing or exclusions, or managing more quantitative aspects, like bounding cost while maximizing quality (i.e., limiting degradation). Fault tolerance in these systems can be seen as switching between configurations, upon the occurrence of failure of e.g., resources, while maintaining a certain level of functionality, possibly in degraded modes.

hence, in a multi-task system, the control of activations and mode switching must restrict the system within allowed combinations, with respect to statical or dynamical conditions.

1.2 Task handlers in a layered architecture

Such systems present various aspects and layers: low-level interfaces with the controlled physical system, resources (computation, communication, or controlled physical system), control tasks (with possibly different versions w.r.t. performance), and, on top of all this, applications sequencing such tasks in order to fulfill some functionality. Figure 1 depicts a general architecture of the considered systems [2], which this paper considers in the more particular case of control systems and tasks. In this framework, the library of control tasks is managed by a task handler, that receives requests from the application, and monitors execution in relation with the physical system (in particular, the termination of tasks).

Resources managed by the task handler can be:

- physical resources in the controlled system: actuators (which should typically not be controlled by two different tasks, or should always be under control of some task, or should be used only under conditions concerning their physical interaction with other parts of the system), sensors (which might require to be used with particular parameter settings making them unsharable by several tasks, or which might involve some control themselves w.r.t., e.g., orientation), the way they consume energy, or the performance or quality with which they

can achieve their function, ...

- computing and communication resources for the tasks: processor(s) (which have to perform cyclic computation within a bounded period, which can have different characteristics), memory (which can be bounded too), communications (following some topology), bandwidth, ...

All these aspects involve interactions and constraints to be managed by the task handling controller. The complexity of the resulting system calls for some structuring, and assistance in the design of controllers. We are looking for ways of separating concerns between local, device-dependent aspects of control, and more global, system-level interactions, while keeping the application-independent aspects as re-usable as possible.

1.3 Summary of our approach

We want to assist in the design of safe task handlers, by proposing a method automating their generation. Our purpose is practical, and the formalization we give of the problem is to be hidden from users, who manipulate a few particular but meaningful patterns of model and properties [24].

At the level of abstraction of the task handler, we need to model a discrete behavior, characterized by the discrete states of the resources (e.g., occupied or free) and of the task (active or not, in which mode, ...), and also possibly of the environment or of some external user interaction. A notion of global state can be derived from local states. The dynamical behavior of the system, observed from that point of view, proceeds from state to state, following transitions that occur typically upon the occurrence of significant events, like requests from a user or application, sensor events, ... Therefore, we will use a formalization in terms of finite state automata, with labeled transitions, to model the systems under study.

The tasks handler has to handle requests from the application, while caring for properties of the computation and physical resources. Properties generic to control systems are typically:

- for the physical resources (e.g., actuators): mutual exclusion of control w.r.t. actuators, permanent control of an actuator, logical constraints between actions corresponding to physical aspects (e.g., cooling action between two heat-intensive ones).

- for the computing resources (processor, memory): boundedness of the global cycle time (for a time-sharing execution), switching between versions or modes of the control tasks (with different characteristics of cost in time, energy, memory, quality).

In order to be able to ensure these properties, we have to design the local controllers of the tasks in such a way that they offer the appropriate control states and events to represent the possible behaviors, and to enable the building of a correct global behavior by appropriately synchronizing them.

The control tasks must be structured and equipped so that properties can be enforced. They have a local behavior described as a finite state machine or automaton, representing typically:

- activity state: in reaction to invocation, requests from the application level, and reception of termination signals, a task can be active or not;
- handling of requests, with different responses, whether they can be delayed or not, and if they are, whether several can be memorized and queued, whether they can be forced or not;
- activity modes and commutations: when the same task can be performed according to different versions of an algorithm, or configurations of an execution platform, the automaton describes them, as well as switches between them.

Different patterns for tasks can be imagined, corresponding to different intended behaviors. The global behavior of the system is described by the composition of all such tasks. Controlling these behaviors so that the properties mentioned before are satisfied is our motivation.

Given the context explained above, the problem we address is the safe handling of actions of control systems, for a general architecture characterized by a set of tasks, and a set of properties to ensure, concerning their logical constraints and some more quantitative aspects.

Our approach consists of:

- building an automaton-based model of the possible behaviors of the set of tasks, and distinguishing, in the conditions labeling the transitions, between events that can be controlled by the handler, and the others, uncontrollable ones;

- formulating the properties on the basis of this model, in terms of state to be reached or avoided, or of sequences to be imposed or forbidden;
- automating the building of the handler by using discrete controller synthesis; it consists of generating the functions which give, for a current state and uncontrollable events, the set of controllable events such that the properties are satisfied.

1.4 Related work

This work unifies first results obtained in more particular contexts. Particularly, the class of tasks constituting the systems takes its inspiration in robot control systems, where a control law, designed following techniques of continuous control theory, is executed cyclically. Such computations can be controlled according to given events and states, control patterns, as proposed in e.g. the ORCCAD robot programming environment [5]. A complete system hosts a number of such tasks, which can be activated in complex sequences. Most formal methods approaches applied on such models concern verification. In our approach we explore the use of discrete controller synthesis, as a more constructive technique.

A first approach inspired by robotic systems was proposed, using a teleoperation application as illustration [23]. Another approach considered tasks with multiple modes or versions, and their characterization by weights associated to states, and used in optimal synthesis [18]. This paper proposes a unified formalization, and complements in the model. Such an automated construction of the task handler can be generalized as a way of compiling a controller from a mixed language [2]. Such a compilation is not classical, in that it makes use not only of a formal model of behaviors of the compiled program, which is already the case in the compilation of synchronous languages concerning invariants [9], but of a model of its dynamics.

Concerning applications of discrete controller synthesis, related approaches working with models of task systems can be found, often using timed models and synthesis techniques [13]. The orientation then is more towards the synthesis of application specific schedulings for real-time tasks [11]. We concentrate on rather logical and discrete aspects (even when considering static weights), where we feel the algorithmic complexity costs are much less, and allow for a greater scalability potential.

In a component-based design setting [3], synthesis can be used to compute the allowed interactions, given constraints and properties to be enforced, of components seen through an interface characterizing their input and output conditions; it produces a form of partial protocol, leaving some non-determinism to be resolved by an on-line scheduler in a fair manner. Also related are works where interfaces of components are determined as an abstraction of internal behaviors, and used in order to compose them while preserving properties of interest [6].

2 Automata and Controller Synthesis

The basic models are transition systems, in which events can occur simultaneously (A transition between two consecutive states can be labeled by a vector or conjunction of events). Given a property involving states or events, discrete control synthesis consists of determining the constraints that make the resulting automaton satisfy the property, by inhibiting the transitions leading to its violation. In the following, we give definitions inspired by [12, 19, 2], and introduce graphical notations which will be used for convenience further.

2.1 Synchronous Automata

Given a set of boolean variables $\mathcal{V} = \{\mathcal{V}_1, \cdots, \mathcal{V}_n\}$, we first define a *valuation* of the set of variables $\mathcal{V}$ as a function: $val : \mathcal{V} \longrightarrow \{\texttt{true}, \texttt{false}\}$. The valuation val assigns to each variable in $\mathcal{V}$ a value, either true or false. The set of valuations of $\mathcal{V}$ is noted $\mathcal{V}al(\mathcal{V})$. Given a boolean expression $B(\mathcal{V})$ and $v \in \mathcal{V}al(\mathcal{V})$, we denote by $B(v)$ the predicate B valuated according to the values of $\mathcal{V}$. Moreover, we denote by B^+ the set of variables that appear as positive elements in the predicate B, i.e. $B^+ = \{\mathcal{V}_i \in \mathcal{V} \mid \mathcal{V}_i \wedge B(\mathcal{V}) = B(\mathcal{V})\}$. Respectively, we denote by B^- the set of variables that appear as negative elements in the predicate B, i.e. $B^- = \{\mathcal{V}_i \in \mathcal{V} \mid \neg\mathcal{V}_i \wedge B(\mathcal{V}) = B(\mathcal{V})\}$.

Now, the system on which control will be applied is modeled by a *Synchronous automaton*:

Definition 1 *A synchronous automaton $\mathcal{A}$ is the tuple $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{init}, \mathcal{C}_{init}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ such that:*

1. *$\mathcal{Q}$ is a finite set of* states*; $\mathcal{Q}_{init} \in \mathcal{Q}$ is a set of* initial states*;*

2. *$\mathcal{V}$ and $\mathcal{O}$ are the sets of Boolean Input and Output variables.*

3. $\mathcal{C}_{init} : \mathcal{Q}_{init} \longrightarrow \mathcal{B}ool(\mathcal{V})$ *is the* initial conditions *attached to initial states;* $\mathcal{B}ool(\mathcal{V})$ *is the set of boolean expressions over the set of variables* $\mathcal{V}$*: it is generated by the grammar:* $b ::= x \| b$ *and* $b \| not\ b$*, where* $x \in \mathcal{I}$*;*

4. $\mathcal{T} \subseteq \mathcal{Q} \times \mathcal{B}ool(\mathcal{V}) \times 2^{\mathcal{O}} \times \mathcal{Q}$ *is the set of* transitions*. For* $t = (q, B(\mathcal{V}), O, q')$*,* q*,* q' *are the source and the target states,* $B(\mathcal{V}) \in \mathcal{B}ool(\mathcal{V})$ *is the triggering condition of the transition, and* $O \subseteq \mathcal{O}$ *is the set of outputs emitted whenever the transition is triggered.*

Now, for a state q and a valuation $\upsilon \in \mathcal{V}al(\mathcal{V})$ of the variables, we say that υ is *admissible* in q whenever there exists a transition $(q, B(\mathcal{V}), O, q') \in \mathcal{T}$ such that $B(\upsilon) = true$; at the same time, the set of variables O is emitted. The transition is also said to be admissible.

Example 1 Figure 2 illustrates this definition, with a graphical syntax that will be used in the remainder of the paper. We see states $\{A, B, C\}$, with initial states $\{A, B\}$, the input variables are $\{a, b, c1\}$ and the set of output variables is reduced to the singleton $\{c\}$, initial conditions $\mathcal{C}_{\text{init}}(A) = \texttt{not c1}$ and $\mathcal{C}_{\text{init}}(B) = \texttt{c1}$, and transitions indicated by arrows: each has a source and a sink state, and a label with a Boolean expression on variables giving its triggering condition, and emissions. Implicit transitions, not represented by arrows, are the ones going from each state to itself when no condition labeling an outgoing transition is true.

A synchronous automaton *deterministic* whenever given an initial valuation υ of the variables $\mathcal{V}$, there is only one valid initial state and given a state and a valuation of the variables, then only one destination state can be reached and the outputs are identical (see [24] for the formal definition).

The semantics of an automaton. $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{\text{init}}, \mathcal{C}_{\text{init}}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ is the following. Assume $\mathcal{A}$ is initialized in state $q_{\text{init}} = q_o \in \mathcal{Q}_{\text{init}}$ with valuation υ_o such that $\mathcal{C}_{\text{init}}(q_{\text{init}})(\upsilon_o) = true$. Further, with the valuation υ_1, $\mathcal{A}$ evolves into state q_1 and emits O_1 such that $(q_o, B_1, O_1, q_1) \in \mathcal{T}$ and $B_1(\upsilon_1) = true$. Assume now that the system has evolved so far into state q_i (i.e. after the reception of i variable valuations), then upon the reception υ_{i+1}, the system will evolve into state q_{i+1} and emit O_{i+1}, such that $(q_i, B_{i+1}, O_{i+1}, q_{i+1}) \in \mathcal{T}$ is an admissible transition in the synchronous automaton $\mathcal{A}$ and $B_{i+1}(\upsilon_{i+1}) = true$ (i.e. B_{i+1} valuates to *true* w.r.t. the valuation υ_{i+1}).

Now given a synchronous automaton $\mathcal{A}$, its behavior is characterized in terms of traces depending on variables. Let $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{\text{init}}, \mathcal{C}_{\text{init}}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ and $\upsilon_0, \upsilon_1, \upsilon_2, ..., \upsilon_n, ...\in \mathcal{V}al(\mathcal{V})$ be an infinite sequence of valuations of the variables $\mathcal{V}$. A *trace* is a sequence of tuples $t = \{(\upsilon_i, q_i, O_i)\}_{i\geq 0}$ where $q_i \in \mathcal{Q}$ are states such that

$$q_o \in \mathcal{Q}_{\text{init}} \wedge \mathcal{C}_{\text{init}}(q_o)(\upsilon_0) = true \wedge O_0 = \emptyset \text{ and } \forall i, \exists (q_i, B_{i+1}, O_{i+1}, q_{i+1}) \in \mathcal{T} \ \wedge \ B_{i+1}(\upsilon_{i+1}) = true$$

We note $Trace(\mathcal{A})$ the set of all traces of $\mathcal{A}$. For $q \in \mathcal{Q}$, we note

$$\textit{Sub-trace}(\mathcal{A}, q) = \big\{t_q = \{(\upsilon_i, q_i, O_i)\}_{i\geq o} \mid q_o = q \ \wedge \forall i \geq 0, \ (q_i, B_{i+1}, O_{i+1}, q_{i+1}) \in \mathcal{T} \\ \wedge \ B_{i+1}(\upsilon_{i+1}) = true\big\}$$

2.1.1 Composition Operators

Let $\mathcal{A}_i = (\mathcal{Q}_i, \ \mathcal{Q}_{\text{init}i}, \mathcal{C}_{\text{init}i}, \mathcal{V}_i, \mathcal{O}_i, \mathcal{T}_i)$, for $i = 1, 2$ be two automata. We want to define the operation that describes the parallel composition, called the synchronous product. This consists in connecting the output variables of $\mathcal{A}_1$ to the input variables of $\mathcal{A}_2$ whenever they have the same name and vice-versa. Our composition will also perform a form of encapsulation, in the sense that the inputs of each automaton which are outputs of the other will be encapsulated i.e., not considered to be inputs of the composition; however they remain visible as outputs.

Note that we must have $\mathcal{O}_1 \cap \mathcal{O}_2 = \emptyset$. The parallel composition will be denoted $\mathcal{A}_1 \parallel \mathcal{A}_2$. All the output variables remain visible in the composition, however, the input variables of $\mathcal{A}_1 \parallel \mathcal{A}_2$ are given by $(\mathcal{V}_1 \cup \mathcal{V}_2) \setminus (\mathcal{O}_1 \cup \mathcal{O}_2)$ because, given an output of $\mathcal{A}_2$, an input of $\mathcal{A}_1$ can be defined in the composition (same holds for the input of $\mathcal{A}_2$). Let us formally define it.

Definition 2 *The* **synchronous product** *of two automata* $\mathcal{A}_i = (\mathcal{Q}_i, \ \mathcal{Q}_{initi}, \mathcal{C}_{initi}, , \mathcal{V}_i, \mathcal{O}_i, \mathcal{T}_i)$*, for* $i = 1, 2$ *is the automaton* $\mathcal{A}_1 \parallel \mathcal{A}_2 = (\mathcal{Q}, \mathcal{Q}_{init}, \mathcal{C}_{init}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ *defined by:*

- $\mathcal{Q} = \mathcal{Q}_1 \times \mathcal{Q}_2$ *(Cartesian product on sets);* $\mathcal{Q}_{init} = \mathcal{Q}_{init1} \times \mathcal{Q}_{init2}$*;*
- $\mathcal{O} = \mathcal{O}_1 \cup \mathcal{O}_2$ *and* $\mathcal{V} = (\mathcal{V}_1 \cup \mathcal{V}_2) \setminus \mathcal{O}$ *and we denote by* $\Gamma = (\mathcal{O}_1 \cup \mathcal{O}_2) \cap (\mathcal{V}_1 \cup \mathcal{V}_2)$ *the set of variables that were both input and output variables;*
- $\mathcal{C}_{init}(\langle q_1, q_2\rangle) = \mathcal{C}_{init1}(q_1) \wedge \mathcal{C}_{init2}(q_2) \ \forall q_i \in \mathcal{Q}_{initi}, \ i = 1, 2;$

- $\mathcal{T}$ *is defined by* $\begin{cases} (q_i, B_i(\mathcal{V}_i), O_i, q'_i) \in \mathcal{T}_i, \ \forall i = 1,2 \\ (B_1 \wedge B_2)^+ \cap \Gamma \subseteq \mathcal{O} \\ (B_1 \wedge B_2)^- \cap \Gamma \cap \mathcal{O} = \emptyset \end{cases}$
  $\Leftrightarrow (\langle q_1, q_2\rangle, \exists\Gamma(B_1(\mathcal{V}_1)\wedge B_2(\mathcal{V}_2)), O_1 \cup O_2, \langle q'_1, q'_2\rangle) \in \mathcal{T}$ $t = (\langle q_1, q_2\rangle, B_1(\mathcal{V}_1)\wedge B_2(\mathcal{V}_2), \langle q'_1, q'_2\rangle) \in$ $\mathcal{T}$ *iff* $\forall i = 1,2 \ . \ (q_i, B_i(\mathcal{V}_i), q'_i) \in \mathcal{T}_i$. •

The conditions on transitions stipulate that shared positive elements should be outputs, while shared negative elements should not be outputs i.e., both parties agree on the presence of shared elements. The synchronous product is both associative and commutative.

Example 2 *Figure 3 illustrates the synchronous composition. The left part shows the graphical syntax : inclusion of two automata in a round-cornered box, with separation by a dashed line. The right part shows the resulting automaton, with notably the fact that synchronous composition makes for transitions simultaneously in both automata.* ◇

Next we define the Hierarchical Composition. Given a synchronous automaton $\mathcal{A}_b$ and a state of $\mathcal{A}_b$, q_r the idea is to refine the behavior of q_r by means of another synchronous automaton $\mathcal{A}_r$. To simplify the definition, we assume that the top-level and the low level do not share variables (i.e. $(\mathcal{V}_b \cup \mathcal{O}_b) \cap (\mathcal{V}_r \cup \mathcal{O}_r) = \emptyset$).

Definition 3 (Hierarchical Composition) *Let* $\mathcal{A}_b = (\mathcal{Q}_b, \mathcal{Q}_{initb}, \mathcal{C}_{initb}, \mathcal{V}_b, \mathcal{O}_b, \mathcal{T}_b)$ *be a* basis *automaton, let* $q_r \in \mathcal{Q}_b$ *be one of its states to be refined, and let* $\mathcal{A}_r = (\mathcal{Q}_r, \mathcal{Q}_{initr}, \mathcal{C}_{initr}, \mathcal{V}_r, \mathcal{O}_r, \mathcal{T}_r)$ *be the* refinement *automaton used to refine* q_r*. The result of this refinement is the automaton* $\mathcal{A}_b \triangleright (q_r, \mathcal{A}_r) = (\mathcal{Q}, \mathcal{Q}_{init}, \mathcal{C}_{init}, \mathcal{T}, \mathcal{V})$ *where*

- $\mathcal{Q} = \mathcal{Q}_b \setminus \{q_r\} \bigcup \{q_r.q \mid q \in \mathcal{Q}_r\}$;
- *If* $q_r \in \mathcal{Q}_{initb}$*, then* $\forall q_{initr} \in \mathcal{Q}_{initr}, \ q_r.q_{initr} \in \mathcal{Q}_{init}$*, else* $\mathcal{Q}_{initb} = \mathcal{Q}_{init}$;
- $\mathcal{V} = \mathcal{V}_b \cup \mathcal{V}_r$ *and* $\mathcal{O} = \mathcal{O}_b \cup \mathcal{O}_r$;
- $t = (q_1, B, O, q_2) \in \mathcal{T}$ *iff*
  (1) $\big(q_1, q_2 \in \mathcal{Q}_b \setminus \{q_r\} \wedge t \in \mathcal{T}_b\big) \bigvee$

(2) $\big(q_1 \in \mathcal{Q}_b \wedge q_2 = q_r.q_{initr} \wedge (q_1, B', O, q_r) \in \mathcal{T}_b \quad \wedge B = B' \wedge \mathcal{C}_{initr}(q_{initr})\big) \bigvee$

(3) $\big(\exists q, q' \in \mathcal{Q}_r \ s.t. \ q_1 = q_r.q \wedge q_2 \in \mathcal{Q}_b \setminus \{q_r\} \ \wedge$

$(q_r, B', O', q_2) \in \mathcal{T}_b \wedge (q, B'', O'', q') \in \mathcal{T}_r \ \wedge \ B = B' \wedge B'' \wedge O = O' \cup O'')\big) \bigvee$

(4) $\big(\exists q \in \mathcal{Q}_r \ s.t. \ q_1 = q_r.q \wedge \exists q' \in \mathcal{Q}_r \ s.t. \ q_2 = q_r.q' \wedge$

$\exists (q, B', O, q') \in \mathcal{T}_r \wedge B = B' \wedge \neg \vee_{(q_r, B'', q'') \in \mathcal{T}_b} B''\big)$. •

The last bullet depicts when a transition exists in the composed automaton: (1) describes the transitions outside of the refined state q_r, (2) considers the transitions entering q_r whereas (3) deals with the transitions leaving q_r. (4) describes the transitions inside q_r: such a transition exists only if no transition leaving q_r in the basis automaton is allowed at the same time.

Applying this to each state gives the whole hierarchical operator: from the basis automaton $\mathcal{A}_b$, each state $q_i \in \mathcal{Q}_b$ is refined into the automaton $\mathcal{A}_i$. The result is the synchronous automaton noted $\mathcal{A}_b \triangleright (\mathcal{A}_1, \mathcal{A}_2, ..., \mathcal{A}_{|\mathcal{Q}_b|})$ and computed by $(... ((\mathcal{A}_b \triangleright (q_1, \mathcal{A}_1)) \triangleright (q_2, \mathcal{A}_2)) ... \triangleright (q_{|\mathcal{Q}_b|}, \mathcal{A}_{|\mathcal{Q}_b|})$.

Example 3 *Figure 4 illustrates the hierarchical composition. The left part shows the graphical syntax : inclusion of the refinement automaton in the round-cornered box of the refined state. The right part shows the resulting automaton.* ◇

2.1.2 Temporal Properties on Automata

Such transitions systems can have properties related to the reachability of some subset of the state space, or to the existence of paths along which a certain sequence of events exists. They can concern invariants on the states themselves (i.e., the variables of which the valuation defines a state), or the paths that can be taken in the transition system from state to state, etc.

Let $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{\text{init}}, \mathcal{C}_{\text{init}}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ be an automaton, let $E \subseteq \mathcal{Q}$ be a set of states and let $q \in \mathcal{Q}$ be a state a $\mathcal{A}$. We now define some temporal properties, useful in the next section, to express the control objectives that will have to ensured on a system modeled as a synchronous automaton.

We say that a state q is *reachable* for $\mathcal{A}$ whenever there exists a trace $t = \{(v_i, q_i, O_i)\}_{i \geq 0} \in Trace(\mathcal{A})$ that traverses q, i.e. $\exists i \geq 0, q_i = q$.

Definition 4 *Let $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{init}, \mathcal{C}_{init}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ be a synchronous automaton and $E \subseteq \mathcal{Q}$, then*

- *$\mathcal{A}$ satisfies the* Reachability of E *whenever there exists $q \in E$ that is reachable for $\mathcal{A}$.*

- *$\mathcal{A}$ satisfies the* Strong Reachability of E*, whenever for all reachable states q, there exists a trajectory initialized in q that reaches E.*

- *$\mathcal{A}$ satisfies the* Invariance of E *whenever* $\forall t \in Trace(\mathcal{A}), \quad s.t.\ t = \{(\upsilon_i, q_i, O_i)\}_{i \geq 0}, \quad q_i \in E,\ \forall i \geq 0$, *i.e.* $\forall q_{init} \in \mathcal{Q}_{init}$*, whatever the behavior of $\mathcal{A}$ initialized in q_{init}, all the traversed states belong to E.*

- *$\mathcal{A}$ satisfies the* Potential Invariance of E*, whenever for all reachable states $q \in E$,* $\forall t \in$ *Sub-trace*$(\mathcal{A}, q), \quad s.t.\ t = \{(\upsilon_i, q_i, O_i)\}_{i \geq 0}, \quad q_i \in E,\ \forall i \geq 0$ *.] i.e. once $\mathcal{A}$ reached one of the states of E, then it will remain inside forever.*

- *$\mathcal{A}$ satisfies the* Persistence *of E, whenever A satisfies both the Strong Reachability of E and the* Potential Invariance of E

One can also consider more intricate safety properties that can be modeled by means of particular synchronous automata, called *Observer*

Definition 5 *An* observer *for a synchronous automaton $\mathcal{A}$ is a deterministic and reactive*[1] *synchronous automaton $\omega = (\mathcal{Q}_\omega, \mathcal{Q}_{init\omega}, \mathcal{C}_{init\omega}, \mathcal{V}_\omega, \mathcal{O}_\omega, \mathcal{T}_\omega)$, such that*

- $\mathcal{V}_\omega = \mathcal{V}_A \cup \mathcal{O}_A$, $\mathcal{O}_\omega = \emptyset$

- *$Error \in \mathcal{Q}_\omega$ is a trap state such that $(Error, true, Error) \in \mathcal{T}_\omega$.*

An observer expresses the negation of a safety property of a system, hence the state *Error* can be seen as a "bad" location which is entered when the system violates the property. The verification consists in checking whether the composition of the two automata (the plant and the observer) reaches the state *Error* on the observer's side.

[1] A synchronous automaton is reactive whenever for all states, whatever the valuation of the variables, a transition can be triggered.

2.2 Controller Synthesis

As usual in the controller synthesis setting, the events labeling the transitions can be partitioned into those that can not be controlled (e.g. inputs received from sensors, failures) and those of which the value can be determined or constrained, typically by a discrete controller (typically the starting of some task). The former are called *uncontrollable*, and the latter *controllable*. Note that, in the reactive automaton framework, events can occur simultaneously. Hence a transition between two consecutive states can be labeled by a vector or conjunction of events (some controllable, some others uncontrollable). This constitutes one of the main differences with [20]. In our case, transitions are partially controllable, whereas in the Ramadge & Wonham formulation, the events and transitions are either controllable or uncontrollable. See also [1, 22, 26] for works related to the control of synchronous automata.

2.2.1 Controllers

Let $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{\text{init}}, \mathcal{C}_{\text{init}}, \mathcal{V}, \mathcal{O}, \mathcal{T})$ be an automaton. We partition the set $\mathcal{V}$ of variables into the set of *uncontrollable* variables $\mathcal{V}_U$ and the set of *controllable* variables, $\mathcal{V}_C$. We note the set of valuations on controllable (resp. uncontrollable) variables $\mathcal{V}al(\mathcal{V}_C)$ (resp. $\mathcal{V}al(\mathcal{V}_U)$). Each valuation of the variables (υ^u, υ^c) contains an uncontrollable component υ^u and a controllable one υ^c. We have no direct influence on the υ^u part which depends only on the state q, but we can observe it. On the other hand, we have full control over υ^c and we can choose any value of υ^c provided it is admissible. Various strategies can be chosen to determine the value of the controls υ_i^c's. We will here consider control policies where the value of the controllable variables υ^c is statically computed from the current state and the valuation of the uncontrollable variables. Such a controller is called a *static controller*. It is given by:

Definition 6 *A* controller $\mathcal{C}_{\mathcal{A}}$ *of* $\mathcal{A}$ *is given by a pair* $(\mathcal{C}_i, \mathcal{C})$*, where* $\mathcal{C}_i \subseteq \mathcal{Q}_{init}$ *is the restricted set of initial states and* $\mathcal{C}$ *is a function* $\mathcal{C} : \mathcal{Q} \times \mathcal{V}al(\mathcal{V}_U) \longrightarrow \mathcal{B}ool(\mathcal{V}_C)$*.* •

For a given state $q \in \mathcal{Q}$ and a given valuation of the uncontrollable variables $v_u \in \mathcal{V}al(\mathcal{V}_U)$, $\mathcal{C}(s, v_u)$ is a boolean predicate over the variables $\mathcal{V}_C$, such that $\mathcal{C}(q, v_u)(v_c) = true$ means that the controller allows the controllable variables to valuate to v_c.

In the framework of Figure 5, the control strategy is the following: given a state q and a set of uncontrollable events that occurs, the set of controllable events is given by the controller according to the restrictions on transitions computed during the synthesis phase.

The automaton $\mathcal{A}$ controlled by the controller $\mathcal{C}_{\mathcal{A}}$ is another automaton $(\mathcal{Q}^{\mathcal{C}}, {\mathcal{Q}_{\text{init}}}^{\mathcal{C}}, \mathcal{V}, calO, \mathcal{T}^{\mathcal{C}})$, noted $(\mathcal{A}, \mathcal{C}_{\mathcal{A}})$, such that $\mathcal{Q}^{\mathcal{C}} = \mathcal{Q}$, ${\mathcal{Q}_{\text{init}}}^{\mathcal{C}} = \mathcal{C}_i$, and $\mathcal{T}^{\mathcal{C}} \subseteq \mathcal{Q} \times Bool(v) \times \mathcal{Q}$ is such that

$$t = (q, B(\mathcal{V}), O, q') \in \mathcal{T} \Leftrightarrow \tag{1}$$

$$(q, B(\mathcal{V}) \wedge \mathcal{C}(q, \mathcal{V}_u)(\mathcal{V}_c), O, q') \in \mathcal{T}^{\mathcal{C}} \tag{2}$$

And no other transitions are allowed.

There can be several controllers satisfying the control objectives; actually, sometimes forbidding any move is a control which avoids the states not satisfying the property, but this is less than satisfactory w.r.t. the activity of the control system. The notion of maximally permissive controller is the controller which insures the properties satisfaction while keeping the greatest subset of behaviors of the original, uncontrolled, system.

Using symbolic methods [15] (based on BDD techniques, avoiding state space enumeration), we can compute controllers $\mathcal{C}_{\mathcal{A}}$ which ensure either the *(Potential) invariance* of a set of states, or the *(Strong) reachability* of a set of states from the initial states of the system, or the *persistence* of a set of states [15], etc. For more details on the way controllers are synthesized, the reader is referred to [17]. Note that if the property is expressed by means of an observer ω, with a sink state $Error$, then it is sufficient to perform the synchronous product between the automaton $\mathcal{A}$ modeling the plant and the observer and to compute a controller that avoids states of the form $(q, Error)$ to be reachable in $\mathcal{A} \parallel \omega$.

2.3 Optimal Controller Synthesis

On the bases of the notions introduced in Section 2.2, we now introduce the optimal discrete control synthesis problem that is based on the notion of cost function.

2.3.1 Automata Extension

In order to take into account the notion of levels of e.g., quality, time or energy consumption, in the control objectives, let us first extend the synchronous automata definition. Formally, an automaton with costs $(\mathcal{A}, \mathit{Costs})$ consists of an automaton $\mathcal{A} = (\mathcal{Q}, \mathcal{Q}_{\text{init}}, \mathcal{C}_{\text{init}}, \mathcal{V}, \mathcal{T})$ and of a *state cost function Costs*, defined as follows: $\mathit{Costs} : \mathcal{Q} \longrightarrow I\!N$, where the set $I\!N$ represents the set of naturals. Figure 6 illustrates an automaton with costs.

Composition Operators The operators of composition are naturally extended to automata with costs. The extensions need to define how costs are composed on states: the operators composing costs on state do not require any hypothesis *a priori*, except that their result domain has to be $I\!N$. The most common operators are among $+$, $*$, max... Let $(\mathcal{A}, \mathit{Costs})$ be a cost automaton We note OP the composition operator on the state cost.

Synchronous Product. Let $(\mathcal{A}_1, \mathit{Costs}_1)$ and $(\mathcal{A}_2, \mathit{Costs}_2)$ be two cost automata. Then, the *synchronous product of* $(\mathcal{A}_1, \mathit{Costs}_1)$ *and* $(\mathcal{A}_2, \mathit{Costs}_2)$ is the cost automaton $(\mathcal{A}_1, \mathit{Costs}_1) \parallel (\mathcal{A}_2, \mathit{Costs}_2) = (\mathcal{A} \parallel \mathcal{A}', \mathit{Costs})$, where $\forall \langle q_1, q_2 \rangle \in Q_{\mathcal{A} \parallel \mathcal{A}'},\ \mathit{Costs}(\langle q_1, q_2 \rangle) = \mathit{Costs}_1(q_1)\ \mathit{OP}\ \mathit{Costs}_2(q_2)$.

Hierarchical composition. Let $(\mathcal{A}_b, \mathit{Costs}_b)$ and $(\mathcal{A}_r, \mathit{Costs}_r)$ be two cost automata. Let $q_r \in \mathcal{Q}_b$ be a state of $\mathcal{A}_b$, to be refined by $\mathcal{A}_r$. The result of this refinement is the cost automaton $(\mathcal{A}_b, \mathit{Costs}_b) \triangleright (q_r, (\mathcal{A}_r, \mathit{Costs}_r)) = (\mathcal{A}_b \triangleright (q_r, \mathcal{A}_r), \mathit{Costs})$, where we have $\forall q, \mathit{Costs}(q) = \mathit{Costs}_b(q)$ and $\forall q_r \,.\, q,\ \mathit{Costs}(q_r \,.\, q) = \mathit{Costs}_b(q_r)\, \mathit{OP}_r\, \mathit{Costs}_r(q)$, where OP_r is a state cost function making the link between the high and low level of the automaton.

Bounding Properties Let us now go through some properties that may be checked on a cost automaton. They will be used as the basis to express control objectives in the next section.

Local Bounding Property This property ensures that every reachable state (from the initial ones) has bounded costs. Let $(\mathcal{A}, \mathit{Costs})$ be a cost automaton, let $\mathit{Low}, \mathit{Up} \in I\!N$. The cost automaton $(\mathcal{A}, \mathit{Costs})$ is locally bounded iff $\forall t = \{(v_i, q_i)\}_i \in \mathit{Trace}(\mathcal{A}), \forall i > 0,\ \mathit{Low} \leq \mathit{Costs}(q_i) \leq \mathit{Up}$. This property can be easily extended to finite traces as follows:

Bounding Property on Traces The cost of a trace $t = \{(v_i, q_i)\}_i \in \mathit{Sub\text{-}trace}(\mathcal{A}, q)$ within $K \geq 0$ steps is defined by $\mathit{Costs}(t, K, q) = \sum_{i=1..K} \mathit{Costs}(q_i)$. Let $\mathit{Low}, \mathit{Up}$ be integers. The cost bound on

traces within K steps is verified iff $\forall q$, $\forall t \in \mathit{Sub\text{-}trace}(\mathcal{A})$, $\mathit{Low} \leq \mathit{Costs}(t, K, q) \leq \mathit{Up}$.

2.3.2 Optimal Controller Synthesis Problem

Ensuring bounding properties Let $(\mathcal{A}, \mathit{Costs})$ be a cost automaton and let $\mathcal{V} = \mathcal{V}_U \cup \mathcal{V}_C$ be the partition of variables of $\mathcal{A}$ into uncontrollable and controllable. Let $\mathit{Low}, \mathit{Up} \in I\!N$ be integers. The controller synthesis problem for local bounding (resp. bounding on traces) consists in finding a controller $\mathcal{C}$ of $\mathcal{A}$ such that the controlled system of $\mathcal{A}$ by $\mathcal{C}$ is locally bounded (resp. bounded on traces) by Low, Up. The above properties are still logical properties expressed on costs. Indeed, to ensure this property it is sufficient to first compute the set of states $E = \{q \in \mathcal{Q} \mid \mathit{Low} \leq \mathit{Costs}(q) \leq \mathit{Up}\}$ and to make it invariant (See Section 2.2.1).

Optimization properties are also defined using costs. The idea is to make the system evolve into state with the highest (resp. lowest) cost w.r.t. its current position and the valuation of the uncontrollable variables $\mathcal{V}_U$. Intuitively speaking, the cost function is used to express priority between the different states that a system can reach in one transition.

Formally, if it is maximization under consideration, let $(\mathcal{A}, \mathit{Costs})$ be a cost automaton, given a state q and an admissible uncontrollable valuation v^u, then the valuation v_1^c is said to be better compared to the valuation v_2^c whenever, the states q_1 and q_2 s.t. $(q, B, O, q_1) \in \mathcal{T} \ \wedge \ B((v^u, v_1^c)) = \mathit{true}$ (resp. for q_2) are such that $\mathit{Costs}(q_1) \geq \mathit{Costs}(q_2)$. The controller has thus to choose, for a pair (q, v^u), a control compatible with v^u in q that allows the system to evolve into one of the states that has a maximal cost. Hence, if the system is in state q and the valuation of $\mathcal{V}_u$ is received, the set of suitable valuations of $\mathcal{V}_c$ is

$$\begin{aligned}\mathcal{I}_{max}(q, v^u) = \{v^c \in \mathit{val}(\mathcal{V}_C) \ /\exists (q, B, O, q') \in \mathcal{T}, \text{ s.t. } B((v^u, v^c)) = \mathit{true} \wedge \\ \forall v^{c\prime} \text{ s.t. } (q, B', O', q'') \in \mathcal{T}, \text{ s.t. } B'((v^u, v^{c\prime})) = \mathit{true}, \ \mathit{Costs}(q') \geq \mathit{Costs}(q'')\}\end{aligned}$$

Based on this policy, one can easily derive a controller $\mathcal{C} = (\mathcal{C}_{init}, \mathcal{C}_{\mathcal{A}})$, such that $\mathcal{C}_{init} = \mathcal{Q}_{init}$ and $\mathcal{C}_{\mathcal{A}}$ is such that $\forall q \in \mathcal{Q}$, $\forall v^u$ admissible in q $\mathcal{C}_{\mathcal{A}}(q, v^u)(v^c) = \mathit{true} \Leftrightarrow v^c \in \mathcal{I}_{max}(q, v^u)$. The minimization may be similarly defined by replacing "max" with "min" in the above equations.

Cost function composition. When considering two criteria, e.g., quality and energy, you may want to first maximize quality w.r.t. the cost function C_1, and then minimize energy w.r.t. C_2. Given a current state q and a valuation of the uncontrollable variables v^u, for the first objective, you follow the previous methodology, and obtain a set $\mathcal{I}_{max}(q, v^u)$ that are solutions of $\mathcal{C}^1_{\mathcal{A}}(q, v^u)(\mathcal{V}_c)$. Then, the set of suitable valuations for the controllable variables are given by:

$$\mathcal{I} = \{v^c \in \mathcal{I}_{max}(q, v^u) \ /\exists(q, B, O, q') \in \mathcal{T}, \text{ s.t. } B((v^u, v^c)) = true \ \wedge \forall v^{c\prime} \in \mathcal{I}_{max}(q, v^u) \text{ s.t.} \\ (q, B', O', q'') \in \mathcal{T}, \text{ with } B'((v^u, v^{c\prime})) = true, C_2(q') \leq C_1(q'')\}$$

i.e. among the remaining valuations of the controllable variables, we only keep the ones that minimize C_2. The same kind of techniques can be used if one wants to mix event and state costs. These optimal synthesis functionalities are also implemented efficiently in Sigali, and can be used for experiment, as outlined in the next section.

2.4 Tools implementing the models and synthesis

The current implementation of the method, which has been used for the example, relies on the chain of Figure 7. Centrally, we use Sigali [15], which is a tool that performs model-checking, controller synthesis for logical goals, and optimal controller synthesis. The components behaviors, task activations schemes and properties observers, can be describes using a synchronous formalism. The equational language Signal is the synchronous language originally connected with the synthesis tool Sigali [15]. Another such formalism is Mode Automata [14], for which the tool MATOU provides for compilation and has been adapted to generate the `z3z` format, the input format of *Sigali*. The global properties and the weights are expressed into `z3z` by the means of Sigali macros. The result of the synthesis in Sigali is a controller, in the form of a logical relation, which can be interpreted by a resolver module: for a given state and uncontrollable input, the constraints on controllable signals are solved, for example in an incremental, interactive way following the manual valuation of signals. The resolver can be coupled with the original specification of the uncontrolled system, using either Signal, or the tool *SigalSimu* in the case of Mode Automata.

3 Modeling tasks for a safe handling

3.1 Informal motivation of the model

We consider reactive systems with the classical global cyclic scheme, as in StateCharts [10], or in the synchronous approach [9, 4]. For a set of tasks in parallel, each global reaction or cycle will see the performance of the computations associated with each of the active tasks. If one considers tasks with a variety of possible active modes, such as proposed in the Mode Automata [14], then one can consider that these modes are differentiated by some characteristics, such as e.g., time cost, quality, the use of communication bandwidth, energy consumption. Switching between them provides for degrees of freedom in the control of configuration e.g., bounding consumption, while maximizing quality. We want to obtain this control automatically, through discrete control synthesis.

3.2 Task Activation Schemes

3.2.1 Simple task activation schemes

As shown in Fig. 8(a), a *standard task*, with application request, is initially `Idle`. It goes from `Idle` to `Act` when there is a request (through the uncontrollable event `r`) and the controller accepts it (through the controllable event `go`), i.e. the control constraints allow it. With the intention of "installing" controllability in the model, a `Wait` state has been incorporated to enable the recording of a request when the activity of another task prevents the controller from starting it. The controller may choose to make it active once the conditions are favorable. Termination of a task is signaled by the event `stop`. Under this model, only the event `go` is assumed to be *controllable*; the others are *uncontrollable*. This model features emitted events, signaling actual `start` and `end` of the activity, which will be used by the observers as described in Section 3.3.

Robots or control systems often require to be always under control, even for rest configurations, because of gravity or other external forces. This motivates the introduction of a pattern for *default tasks*, fully controllable, shown in Fig. 8(b). It is similar to the standard task except that it is not necessary to have a request in order for a default task to become active.

Another interesting pattern is that of *sustainable* tasks illustrated in Figure 8(c), where the

possible stopping of the activity (signaled by the event `stop`) can be memorized, but the activity sustained. The behavior is quite symmetrical to the memorization and control of requests seen above. An example is the class of tasks in control systems where a control loop has met its objective (e.g., reaching a position), but no other control task has been scheduled in sequence. Sustaining the current task makes that the control objective is regulated further (e.g., the device is kept at the position reached), and the corresponding actuators a still under control.

The preceding patterns are simple particular cases. Others (see Figure 9) would be e.g.,

- to allow the controller to reject requests (i.e., no wait/memorization), with emission of a rejection warning (such a pattern imposes quite strong constraints on a multi-task system);
- the possibility for the environment or application to cancel a request (e.g., in case of abortion or preemption on the requesting process),
- an initialization phase, to be fulfilled before control can be taken over by the task;

3.2.2 Multi-mode task schemes

A particularly interesting case of hierarchical structure is that of tasks that can be executed according to different modes or versions. The notion of modes and their switching is approached in various ways, some related to operating system mechanisms, or language constructs [21]. We adopt the notion of modes as in Mode Automata [14].

Once a set of tasks with modes is defined, one has to consider the switching between them. Within a multi-task environment, starting a new task can occur in several ways: when enough computing resource is available, it can be started right away; when some computing resource has to be made available, it can be achieved by lowering quality (hence time cost) of other already active tasks, through mode switching, while keeping global quality maximal; when no sufficient resource can be released, then we have to go to a waiting state. This kind of control, managing mode switches according to criteria of time cost and quality level, is what we want to obtain automatically, through discrete control synthesis.

Fig. 10 gives the hierarchical automaton for a task, say T_i, for the case with three modes. The pattern formalizes the aspects described above. We have a first level with states, naturally similar

to the previous ones: $\texttt{Idle}_i$ (or $\texttt{I}_i$), where the task is deactivated; $\texttt{Wait}_i$ (or $\texttt{W}_i$), where it has been requested, by an event $\texttt{req}_i$, but is not launched yet, because of lack of authorization $\texttt{Go}_i$ by the controller (due to constraints with the environment or other tasks); $\texttt{Act}_i$ (or $\texttt{A}_i$), where it has been requested, and authorized through $\texttt{Go}_i$, and hence is active.

Within the active state $\texttt{Act}_i$, several modes can be defined. Each task T_i has m_i modes $M_{ij}, 1 \leq j \leq m_i$. For each task T_i, transitions between modes are labeled by conditions c_{k_i}, managed by the controller in order to authorize the switch or not. The mode which becomes active when entering $\texttt{Act}_i$ is chosen according to the one of these conditions which is true. The choice of the mode when starting is determined by the control values.

In the case of three modes, illustrated in Figure 10, we can think of applications such that we have: $\texttt{H}_i$ (highest-quality and time); $\texttt{M}_i$ (medium); $\texttt{L}_i$ (low). These three modes can be switched between according to the transitions and their conditions c_{k_i}. In particular, you have to go through the $\texttt{M}_i$ mode between $\texttt{H}_i$ and $\texttt{L}_i$.

3.2.3 Multi-task systems

The tasks that were modeled previously can be assembled into multi-task systems. Their individual interface basically consists of the inputs and outputs illustrated in Figure 11(a).

Tasks server. In order to model a multi-task system, we consider just the parallel composition of the n tasks, as in synchronous languages [9, 14], in accordance with the definition in section 2.1.1. Hence, a global state or configuration S is described by a vector of state values, one for each task, giving the current active mode. Figure 11(b) illustrates such a parallel composition of 4 tasks, the first and third from the left without modes. Each receives its requests and stop events, as well as the controls `go` and possibly `Ci`, from the controller which has access to all state and event information. Each task also sends out its starts and ends to the environment, as control instructions on the side of the tasks computation processes, or as informations towards an upper level.

For an example with two instances T_1 and T_2 of the three-mode pattern of Figure 10, we can be in a mode where T_1 is in state $\texttt{Idle}_1$ while T_2 is in state $\texttt{Act}_2$, in mode $\texttt{H}_2$: the configuration can be noted in short $S_1 = (\texttt{I}_1, \texttt{H}_2)$. It will be useful in the following to define the number m of

active modes in a configuration S. One can observe that $m \leq n$, as at most one mode is active per task. As illustrated in Figure 12, a global step is taken when one or several mode changes or a task start and/or stop event occurs, going from one global configuration to the next. Such a transition is labeled with a vector of conditions/events of the local transitions, or ϵ on the side where no transition is taken. In the example, when in the configuration S_1, upon the occurrence of ($\texttt{Req}_1$ `and` $\texttt{Go}_1$) and in the absence of $\texttt{Stop}_2$, the configurations shown in Figure 12 can be reached in one step. They are reached through transitions which are labeled by, for T_1: either one of $\texttt{c0}_1$, $\texttt{c1}_1$, $\texttt{c2}_1$, and for T_2: either $\texttt{c1}_2$ or ϵ (empty label, no movement of T_2). E.g., going from S_1 to $S_2 = (\texttt{H}_1, \texttt{M}_2)$ would be through a transition labeled with a combination of the two involved local transitions: $((\texttt{Req}_1 \texttt{ and } \texttt{Go}_1 \texttt{ and } \texttt{c0}_1), \texttt{c1}_2)$.

Synchronizations and applications. Complete systems with multiple tasks can have an application program on top of the tasks server. One way of combining tasks is to have a language of "loose" synchronizations, where part of the scheduling is expressed by the user, in the form of a script or a scenario, and controller synthesis serves as a completion, computing the part relevant for constraints management. A task-level language can be used to describe sequence, parallelism, reaction to events, resulting in a sequence of requests, themselves controlled by the automatically synthesized controller, ensuring satisfaction of the appropriate properties. Automata for the sequence and parallel of two tasks are illustrated in Figure 13 where, upon reception, respectively, of `req-seq` and `req-par`, requests are emitted accordingly towards the tasks, the ends of which then being awaited for. Requests to tasks are the local events, and not inputs any more.

The resulting system is illustrated in Figure 11, where, on top of the server, an application component, when receiving a request from its own environment, emits requests towards the tasks, and reacts on their ends, until reaching its own end. The controller can rely on the states and events from the application, in addition to that of the tasks.

3.3 Associated logical properties

3.3.1 How to specify properties

Formulation Properties can concern configurations of the system, defining states that are defined to be consistent or dangerous. They can also concern sequences of events, states or tasks, that can be either forbidden, or required. The objectives for which controllers will be synthesized will consist of making a set of states verifying a property e.g., invariant, or reachable.

Their formulation can involve writing logic formulæ which can be quite technical, especially in the case of temporal logics. There are works contributing to making such specifications more accessible, on the basis of domain-specific knowledge and notion, e.g., in robotics (e.g. proposals for a specialized language of properties in ORCCAD).

Observers can be used as an alternative. Instead of a predicate on state variables, or a temporal logic formula, they are defined by an automaton, recognizing the sequence, with a terminal state. The global system is the parallel composition of the observer and the pre-existing system. It can be submitted to a control objective of safety, keeping out of the terminal error state. This method has the advantage of making it possible to use the same specification language as for behaviors. As an example, to ensure that the system avoids a certain sequence of task activations, one can define a transition system recognizing that sequence, with a final state designating the error. The reachability of this state is then a simple property that can be used for verification or synthesis.

Generic, architecture-specific, and application-specific properties In a way similar to the proposal of typical task control schemes, we propose property schemes typical to the class of systems under consideration. The fact that we consider a specific domain favors the identification of such schemes, which can be encoded once and for all, and used by engineers without the need to re-formulate them, hence avoiding a risk of misinterpretation and discrepancy with the intended specification. Of course, the possibility remains to encounter the need for more particular properties, specific either to the considered architecture (resources, sensors and actuators, communication network, ...), or even to the application executed on this platform.

We see a gradation beginning, at the most general level, with generic, architecture-independent

properties. These properties should be verified, or controlled for, for any particular set of tasks and resources. An example for control systems is that every actuator should always be under control, and that all tasks controlling an actuator should be exclusive. Another example is that access to other exclusive (single-user) resources should be controlled appropriately.

At a more particular level, we see architecture-specific, application-independent properties. They concern aspects that have to be satisfied for all uses of an architecture, i.e., whatever the application executed on the platform, and the sequence of requests and inputs. At this level, one can state constraints holding between particular tasks, whatever the application. An example is tasks of a chemical plant, not using the same pipes or tanks, but required to be exclusive because of potential reaction through fumes. Another example is the interdiction of combinations of tasks involving manual interaction, because of safety regulations w.r.t. to the level of attention required.

Finally, application-specific constraints can always be part of a design. These properties, having to be determined for each application, might involve all the expressiveness of a temporal logic.

3.3.2 Properties on states

Generic objectives on any set of tasks We identified a few properties and objectives that can be defined on the task patterns proposed above. They are domain-specific in that they concern a particular abstraction of control tasks, and the constraints are related to requirements of control systems. At the same time, they are generic w.r.t. that domain, meaning that they are relevant to a wide range of applications, and that they can be significantly used for any instantiation of the given patterns. They can also be systematically derived for robotic missions built from task-level components as we introduced. In that sense, we have a framework where models as well as objectives can be automatically compiled from high-level specifications in a domain-specific language.

They are based on the notion that the system to be controlled is composed of a set of devices or resources (e.g. actuators of a robot) $d \in \mathcal{D}$. The system has a set of tasks $t \in \mathcal{T}$. These tasks are instantiations of the patterns described previously, i.e., for each of them we have the states $idle_t$, $wait_t$, act_t, and we have a predicate active(t) telling us whether or not a task is active, i.e., for a standard or default task, in state Act, or for a sustainable task, in either state $\mathsf{Act\ or\ Sust}$. Their composition, together with observer automata in a set $\mathcal{O}$, defines a global automaton. A function

u defines, for a task t, the set $u(t)$ of devices (or resources) used by t: $u : \mathcal{T} \longrightarrow 2^{\mathcal{D}}$. Reciprocally, another function a defines, for a device d, the set of tasks that can make access to d: $a : \mathcal{D} \longrightarrow 2^{\mathcal{T}}$.

Unicity of control for an actuator. A basic property of such a system is that there always is, for each actuator, at least one control law controlling a given actuator (otherwise the actuator is not under control, and may e.g., fall down), and at most one (otherwise the actuator receives possibly contradictory commands). This can be written as an invariant to be satisfied by all behaviors: $\forall d \in \mathcal{D}, \exists! t \in \mathcal{T} \ . \ d \in u(t) \wedge \text{active}(t)$. In order to distinguish between resources requiring to be under control, and others accepting at most one user, we will decompose this property.

For synthesis purposes, invariance objectives can be expressed on the state information of the control. For this, we define a Boolean condition in terms of the model state variables (i.e., an observer of the situation), and then define the objective as a simple expression, like: achieving invariance of the truth of the value of that Boolean.

Existence at all instants of a control law can be described as: $\bigwedge_{d \in \mathcal{D}} (\bigvee_{t \in a(d)} \text{active}(t))$. The objective is to *make this invariantly true.*

Exclusivity of resource access, more specifically of control laws on the same actuator, is another basic property to be maintained in a control system, i.e., at most one control law at a time can be controlling a given actuator. It can be achieved in a way quite close to the previous one: it is a property on states, where the situation *to be avoided* is that for any actuator there is more than one active task, in other words: $\bigwedge_{d \in \mathcal{D}} (\bigvee_{(t,t') \in a(d)^2, t \neq t'} (\text{active}(t) \wedge \text{active}(t')))$. The objective is to *make it invariantly false.*

Architecture-specific properties.

Exclusivity between two tasks or modes can be related to a common resource as seen above, or to more external issues, specific to the system under consideration. For example, side-effects of different tasks (in terms of temperature, or dispersion of chemicals in the ambient, ...), or safety-related considerations (the fact that it can be possible but dangerous to perform two particular tasks at the same time) can be known to make them incompatible. This can be captured by a simple expression on the corresponding states; for an exclusivity between task t_i and mode m_{jk}: $\text{active}(t_i) \wedge \text{active}(m_{jk})$. The objective is to *make it invariantly false.*

Reachability of a rest configuration. Some systems may require a procedure to be stopped, more elaborate than simply switching them off. Or similarly, they can have a defined rest position or configuration, into which they are considered safe, and that configuration should always be reachable. Such a configuration can be defined w.r.t. some of the resources and devices in presence; some of them may have to be under the control of a task, which should then be active, the others being idle.

For such a rest configuration defined by an expression on state variables, hence defining a subset of states, the objective is to *make it reachable from all the reachable state space.*

Another use of reachability is to verify that, if a controller was found w.r.t. invariance constraints, it is not so conservative that all tasks are accepting requests only to be blocked in a waiting state. This can be excluded by defining a set of states where the functionality is fulfilled (e.g., a terminal state), and requiring it to be always reachable. An example is in fault-tolerant systems, where for all faults (within the fault model considered), the functionality can still be fulfilled [8].

3.3.3 Task sequences and transitory modes.

Until now properties considered are static in the sense that they concern situations rather than series of transitions. More dynamical properties are related to allowed task sequences and requested transitory modes. As was said before, this can involve defining observers recognizing these sequences, and distinguishing the terminal state (or set of them), in order for it to be used in objectives of reachability or invariance.

Avoiding t_3 between t_1 and t_2 is an example of property. More precisely, we want to have, between the ending of t_1 and the next activation of t_2, no activation of t_3. Cases with simultaneity are acceptable. For this, an observer can be proposed as in Figure 14.

This automaton, initially in state `Clear`, observes the start and end events emitted by the tasks. An occurrence of the starting of task t_1 causes a transition to state `InT1`.

There, in case t_3 is started, a transition goes to `Error`. Otherwise, upon the task deactivation `endT1`: in the absence of `startT2` the transition goes to `afterT1`; in the simultaneous presence of `startT2` the transition goes directly to `InT2`.

From `afterT1`, if t_3 starts, we go to `Error`, otherwise, when t_2 starts, we go to `InT2`.

From `InT2`, `startT3` before t_2's end brings to `Error`, otherwise `endT2` brings back to the initial state (where t_3's activity may be observed) or, if we have simultaneously `startT1`, to `InT1` directly.

The desired property is that a sequence reaching `Error` can not be followed. For this, the synthesis objective is: *making invariant the value false for the state variable* `Error`.

Imposing t_2 between t_1 and t_3 can then be seen as avoiding t_3 between t_1 and t_2, re-using the same observer. In particular, it can be useful in order to have automatically a transitory mode through a default task t_2 between two tasks t_1 and t_3 on the same actuator. This corresponds to the existence of control laws which needs to be separated by a special mode handling the transitory situation between the two (e.g., velocity control followed by position control, or cooling between two tasks soliciting a device which heats up).

3.4 Weights, costs and quality of service

Until now, only strictly logical aspects of the system behaviors and constraints have been considered. We will now assign quantitative characteristics to states and modes, in order to encompass some aspects relating to resource usage and sharing, characterization of quality of the actions performed, and the like. They are typically specified locally to these states, and we define how to deduce them for tasks and applications. They can be used for the synthesis of controllers for which the objective is to keep bounded, minimize or maximize the global weights.

3.4.1 Modes, tasks and applications characteristics

We define a cost function representing quality C_q, and another one representing computing time C_t (i.e., the time to perform computation within one step of the reactive system). Such cost functions must be manually defined, on the basis of e.g., WCET analysis.

Modes The functions are defined for each mode of each task: $q_{ij} = C_q(M_{ij})$ and $c_{ij} = C_t(M_{ij})$. One can also think to associate costs to the events in order to take into account the time necessary for mode changes. In this paper, we consider *a priori* that a task in `Idle` or `Wait` state has a null cost and quality. However, a non-null cost of a waiting task could e.g. be used to account for the possible operating system overhead. We also consider that the cost and quality of a mode are

related in such a way that: $\forall j, k : q_{ij} \geq q_{ik} \Rightarrow c_{ij} \geq c_{ik}$. They need not be considered proportional, though (e.g., in robot movement control, computation of inertia involves a big time cost for a small precision gain, when acceleration is small).

Example 4 *In the example of the three-mode task pattern of Figure 10, we can define costs for two instances T_1 and T_2 as shown in Figure 15.*

Tasks For a task T_i, we can define its current time cost: $t_i = \sum_j C_t(M_{ij}) * \delta_{ij}$, where δ_{ij} is equal to 1 whenever the task i is in mode M_{ij}, 0 otherwise. Accordingly, its current quality is: $q_i = \sum_j C_q(M_{ij}) * \delta_{ij}$. Let us recall that, in our framework, only one mode is active at each time. So the current task cost is the cost of the currently active mode.

Applications They are composed of a set of tasks in parallel. When implemented on a single processor architecture, the computations for each of them are executed in sequence, and the duration of a reaction is the sum of individual durations for active task cycles. Indeed, the so-called "zero-time" hypothesis of synchronous languages is a metaphor of the fact that interactions inside a cycle can be compiled away, and hence costless; actual implementation has a worst-case execution time (WCET) to be measured w.r.t. the environment dynamics.

For a composition of the n tasks, the current global time cost (within a cycle) is: $T = \sum_i t_i$. The current global quality of an application is: $Q = \sum_i q_i$. In the example of the two instances of the three-mode task pattern, and for the particular configurations illustrated in Figure 12, they are characterized as shown in Figure 16.

3.4.2 Associated properties and optimization

Bounding the sum of costs. Sharing the processor means that at each cycle, the time to compute each of the tasks (one step of each) must be contained within a global period T_{max}, i.e.: *the sum of local costs should always be less*: $T < T_{max}$. The control of mode switching must *go only to global configurations where this property is true.* This is a logical invariance objective.

In the example, we take $T_{max} = 11$. Then, we have to exclude configurations S_5 and S_6, i.e., the controller has to *forbid* the transitions from S_1 to S_5 and S_6.

Maximizing quality. We want to deliver the functionality at the best possible global quality (i.e., least degradation): maximal, and evenly distributed. From a configuration S with successors $S' \in Succ(S)$, we want to keep only transitions going to a configuration where the property holds.

First objective: maximizing global quality. *Amongst the remaining possible next global configurations S', go only to those where the sum of local qualities is maximal* i.e., $Q = Q_{max} = max_{S' \in Succ(S)}(Q(S'))$. In the example, on the remaining successor states, we have $Q_{max} = 7$: we keep only configurations S_2 and S_7. There can be several successors with equal, maximal, global quality, but with different local distributions amongst modes.

Second objective: minimize time. For an equivalent quality, we want to pay the least cost: *Amongst the remaining possible next global configurations S', go only to those where the time cost T is minimum, i.e., $T = T_{min} = min_{S' \in Succ(S)}(T(S))$.* In the example, for the same quality $Q_{max} = 7$, S_2 has a time cost of 10, S_7 has a time cost of 8. Hence, we keep only configuration S_7. That is to say, the controller must authorize only ($\texttt{c2}_1$,ϵ) from S_1.

Alternate objective: having a homogeneous quality. Another objective could concern homogeneous quality, defined as: amongst configurations with equal maximal quality, choose those where the values are closest to the average. That is to say, the property we want to be satisfied is: *amongst the remaining possible next global configurations S', go only to those where the difference D between local qualities and the average is minimum, i.e., $D = D_{min} = min_{S' \in Succ(S)}(D(S))$* . The average quality of modes *of active tasks* (which are m in number) is: $Q_{avg} = \frac{Q}{m}$ We define the difference $D(S) = \sum_i (|q_i - Q_{avg}| \times a_i)$. More elaborate notions of distance could be meaningful here, like variance (the average of distances to the average) or standard deviation. In the example, on the states remaining after the first objective, we have $Q_{avg} = 3 . 5$, and $D(S_2) = 1, D(S_7) = 5$. Hence, we keep only S_2 i.e., the controller must authorize only ($\texttt{c0}_1$ `and` $\texttt{c1}_2$) from S_1.

4 Application: case study of a robotic system

4.1 The robotic system

The system decomposes into sub-systems, according to the actuators: the articulated arm, and the gripper held at the end of the arm. This system could itself become a sub-system, e.g. in

an excavation system, completed with the rotating cabin on which the arm is mounted, and the mobile base, carrying the whole, itself composed of two tracks [25]. Each subsystem i.e., actuator, is equipped with a library of control tasks, corresponding to different functionalities of the device, and different ways of achieving them, according to different criteria. The complete system is simply constituted by the composition of all its actuators, each with its control tasks.

The gripper is equipped with 3 tasks. The manipulation task `Gmanu` (standard) allows for the operator to directly decide on the movements of the grip through teleoperation. The task `Gmaint` (default) maintains the gripper at the current opening position. The task `Gauto` (two-modes task) corresponds to an automated movement control task, according to two versions, differing by the algorithms: the high mode `H` makes more accurate computations of control, involving e.g., mechanical models of friction, whereas the low mode makes an approximation. Hence, the low mode `L` has a lower quality of control, and a lower cost in computation time.

The articulated arm is equipped with 4 tasks. The task `Amanu` (standard) offers manual control to the operator. The task `Ahome` (sustainable) brings the arm towards a predefined resting position, where it can be sustained, thereby keeping the device under control. The automated movement task `Auto` (two-modes task) is defined by a control law, for example, trajectory following, with two versions: the high mode `H` is more accurate, whereas the low mode `L` makes an approximation, and uses less elaborate sensing. The task `Amaint` (default) maintains the current position.

An application for this system can be a sequence of bringing the arm in its home position (in order to reinitialize all relevant parameters), then having an automated movement towards some position, then performing a manual arm movement e.g., to finely approach something to be gripped, then manually gripping the object there, and then, finally, bringing the arm back in its home position again.

4.2 Properties and objectives

4.2.1 Generic properties

The most basic ones concern presence and uniqueness of control. Hence, for each actuator:

(a) *It must always be under control (otherwise unpredictable movements can occur).*

(b) *There must be at most one active control law (otherwise control can incoherent).*

This translates into a synthesis objective as seen in Section 3.3.2.

4.2.2 Architecture-specific properties

(a) *No manual manipulation of the gripper when the arm is manually teleoperated.* This corresponds to a safety requirement related to attention span of the human operator, and the avoidance of human failure. This translates into the synthesis obejctive of making invariantly false the conjunction of the activity of `Amanu` and `Gmanu`.

(b) *Between a manual arm movement and an automated one, the default task maintaining the current position should be activated.* It corresponds to a fine recalibration of sensors and actuators between these two movements. This translates into using the observer of Section 3.3.3 shown in Figure 14. It is a case where imposing t_2 (`Amaint`) between t_1 (`Amanu`) and t_3 (`Aauto`) can then be seen as avoiding t_3 between t_1 and t_2.

They can also be quantitative properties: for a given architecture, bounds must be respected regarding consumption, e.g., of energy or of computing power. Costs of computing power for each task are assigned as in the table shown in Figure 17.

These costs add up when tasks are active in parallel, thereby defining a global cost. Giving the processor capacity, *the upper bound must be equal to 12.* It entails that the subset of states where the global cost is strictly higher than the bound has to be forbidden by control.

4.2.3 Application-specific properties

In the example of sequential application above, *the termination of the sequence should always be reachable.* A counter-example would be that other properties, quantitative or logical, would involve control sequences going into loops or waiting configurations with no possibility to proceed.

4.3 Model specification in Mode Automata

An example of task is given in Figure 18, which is the concrete encoding of the corresponding automata of Figure 8(a). The Mode Automata [14] textual syntax reads as follows. Each automaton

is given a name (e.g., `Ahome`), and begins with a declaration of states. The initial state is indicated by the keyword `init`. With each of them, equations can be associated, defining values of variables (here, we consider Boolean variables), at each instant when the system is in the corresponding state (e.g., in state `S_Ahome`, variable `end_Ahome` takes the value of `go_Ahome`, and in state `I_Ahome` it takes the value `false`). Then, a list of transitions is given, from source state to sink state, with a condition on variables. For example, from state `S_Ahome`, a transition is taken to state `I_Ahome` if the condition `go_Ahome` is true. One can note that the equations defining `end_Ahome` make that it is true on the instant this transition is taken, and false afterwards; this is a way of encoding in Mode Automata the emission of `end` upon reception of `go` (i.e., `go/end`).

The observer from Figure 14, observing that no two activations of automatic arm movement (tasks 1 and 2) occur without a position maintaining task (task 3) in between, can also be encoded in Mode Automata.

The application example is encoded as a sequential automaton as well: it emits requests to the tasks, and proceeds into sequence upon their termination, sending a request to the next, until the `End` state.

The complete model is obtained by assembling task models and observers, by synchronous compositionThis complete model is then compiled, according to Figure 7, using Matou, into an executable format, on the one hand, and on the other hand into the `z3z` format, an equational encoding of the global transition system, taken as input by the verification and synthesis tool Sigali.

4.4 Controller synthesis with Sigali

Using the Sigali tool begins with loading the `z3z` file, encoding the transition system `S`, where variables receive automatically generated names: a typical example is the `Error` state of the observer `obs`, recognizable in variable `Error_de_obs_0`.

Observer. We want to *make invariant* the sub-states of states where the observer is not in state `Error`, i.e., where `Error_de_obs_0` is false. This is concretely encoded by first defining the subset of states called `PROP_obs`, where `B_False` designates the set of states where the variable

is false. The synthesis operation for invariance, called `S_Security` is then appliedto the original system `S` with the set `PROP_obs`, the result redefining `S`.

```
S : S_Security(S, B_False(S, Error_de_obs_0) );
```

Unicity of control. First, we give an expression for one active task, for each actuator.

```
One_G : A_Gmanu_de_Gmanu_1
 or A_Gmaint_de_Gmaint_2 or AH_Gauto_de_Gauto_3
 or AL_Gauto_de_Gauto_3 ;
One_A : AH_Aauto_de_Aauto_4
 or AL_Aauto_de_Aauto_4 or A_Ahome_de_Ahome_5
 or S_Ahome_de_Ahome_5 or A_Amanu_de_Amanu_6
 or A_Amaint_de_Amaint_7 ;
One : One_G and One_A  ;
```

The same can be done for the second part, by defining a variable `Several` true when several tasks are active on either actuator. The controller should then make invariant the set where `One` is true and `Several` is false, which is encoded as follows:

```
PROP_one:B_True(S, One);
PROP_only:B_False(S, Several) ;
PROP_one_only:intersection(PROP_one,PROP_only);
S : S_Security(S, PROP_one_only);
```

Insuring exclusion between manual control for the gripper and for the arm goes along the same lines:

```
PROP_excl : B_False(S, A_Gmanu_de_Gmanu_1
                       and A_Amanu_de_Amanu_6);
S : S_Security(S, PROP_excl);
```

where an expression on state variables is made invariantly false.

Reachability. A resting position of the system is defined, where both actuators are controlled by position maintaining tasks. This configuration must be always reachable, from any other configuration of the system.

```
Rest : I_Gmanu_de_Gmanu_1 and
 A_Gmaint_de_Gmaint_2 and I_Gauto_de_Gauto_3
 and I_Aauto_de_Aauto_4 and I_Ahome_de_Ahome_5
and I_Amanu_de_Amanu_6 and A_Amaint_de_Amaint_7;
S : S_Reachable(S,B_True(S, Rest));
```

This objective is to be applied after all the invariance ones, as the removal of transitions might break the reachability. In the case of the systems we consider, actually, reachability can be treated by verification, at the end of the synthesis process, as if it is not satisfied, then it is unlikely that it can be controlled.

Another reachability objective can be associated with the sequential application: the reachability of its terminal state `End`:

```
S : S_Reachable(S, B_True(S, End_de_appli_8));
```

Weight-related properties. These properties involve declaring values of the computing power consumption, and specifying the bound.

Declaring weights. This is done as follows for each state variables corresponding to an active state: using the SIGALI instruction `a_var`, a value is associated with the variable having the value true; for all other cases, it is given the value 0, which is neutral for addition. This way, combining values of composed task controllers is done simply by adding their costs into `C_G` for the gripper, and `C` for the global cost.

```
C_GHauto : a_var(AH_Gauto_de_Gauto_3,0,4,0);
C_GLauto : a_var(AL_Gauto_de_Gauto_3,0,2,0);
C_G      : C_Gmanu+C_Gmaint+C_GHauto+C_GLauto;
C        : C_G + C_A ;
```

Respecting bounds. For this, a bound is declared, and the instruction `a_inf` is used to obtain the set of states `InBound_power` where the cost in lower than the bound (e.g. 12 in the example). The latter is then being made invariant.

```
InBound_power : a_inf(C, 12) ;
S : S_Security(S, InBound_power) ;
```

4.5 Interactive simulation

At every phase in the above construction of the model and control, it is possible to obtain a simulator of the behaviors of the controlled system. This way, a user can observe how the behaviors change when adding one objective, and verify whether the constraints added correspond to the problem to be solved. A simulation consists of iterating, step by step, the following three operations:

1. *simulating the environment* is done through the uncontrollable inputs panel (see Fig. 19), where one can enter the requests from the operator, and the events signaling termination of tasks.

2. *choosing among correct controls* is done through the controllable events panel (see Fig. 19). Values ruled out by constraints are represented by non-selectable buttons. There can be possibly several allowed values, if the constraints do not completely determine the control from the inputs. In order to obtain a control function, i.e., an input-deterministic controller, the specification has to be strengthened, or optimization criteria (w.r.t. costs on states or events) have to be applied, or a random choice can be applied. In Figure 19, the situation shown is related to our example, where the request for `Aauto` is refused, because there has been no `Amaint` since its last activation.

3. *the dynamical evolution* is observed with the task states display . Changes in behaviors obtained after adding a control objective can show in, e.g., observing that a requested task goes into wait state when another, incompatible one is active, until termination.

This interactive simulation scheme is a special case of the general execution scheme: acquiring external inputs and system state, choosing the control (which can be done by online choice between correct values), and emitting commands towards the system.

5 Conclusions and perspectives

We have proposed modelling patterns for multi-tasks real-time control systems, such as in robotic, automotive or avionic systems, as well as for some of their relevant safety properties, that have to

be insured through all possible switchings in control activities. This formal modelling in terms of transition systems allows for the application of the discrete controller synthesis technique. This provides us with a systematic methodology, for the automatic generation of safe task handlers which enforce the properties given as objectives. We experimented the approach using models, languages and tools from the synchronous approach to reactive systems.

Possible extensions based on these results are to define an end-user-friendly, domain-specific programming language on these bases, where languages constructs are in terms of tasks, modes, start and end control, resources and their shareability, and a specification of sequences of tasks constituting an application; this will alleviate all need for direct formal technical expertise for the end-user. It would also be interesting to use more advanced synthesis techniques, like optimization on bounded-length paths, especially when applied after a reachability objective, or hierarchical structure in synthesis [16, 7], ... We also plan to explore execution mechanisms, and to replace simulation by coupling the resulting controller with an exploitation system hosting the actual tasks; finally, one may also apply these techniques to other application domains, where a specific use fo the same kind of techniques can be made, e.g., fault tolerance [8].

Acknowledgments. We gratefuly acknowledge helpful discussions with M. Abdennebi, H. Alla, K. Altisen, A. Clodic, S. Iddir, P. Manon, F. Maraninchi, A. Medina-Rodriguez, T. Neveu, P. P. Parida.

References

[1] A. Aziz, F. Balarin, R. K. Brayton, M. D. Dibenedetto, A. Sladanha, and A. L. Sangiovanni-Vincentelli. Supervisory control of finite state machines. In P. Wolper, editor, *Proceedings of the 7th International Conference On Computer Aided Verification*, volume 939 of *Lecture Notes in Computer Science*, pages 279–292, Liege, Belgium, July 1995. Springer Verlag.

[2] K. Altisen, A. Clodic, F. Maraninchi, and E. Rutten. Using controller-synthesis techniques to build property-enforcing layers. In *Proceedings of the European Symposium on Programming, ESOP'03,* April 7 - 11, 2003,Warsaw, Poland, 2003.

[3] K. Altisen, G. Gößler, and J. Sifakis. Scheduler modelling based on the controller synthesis paradigm. *Journal of Real-Time Systems*, 23(1), 2002.

[4] A. Benveniste, P. Caspi, S. Edwards, N. Halbwachs, P. Le Guernic, and R. de Simone. The synchronous languages twelve years later. *Proc. of the IEEE, Special issue on embedded systems*, 91(1):64–83, January 2003.

[5] J-J. Borrelly, E. Coste-Manière, B. Espiau, K. Kapellos, R. Pissard-Gibollet, D. Simon, and N. Turro. The ORCCAD architecture. *Int. J. of Robotics Research*, 17(4), 1998.

[6] A. Chakrabarti, L. de Alfaro, T.A. Henzinger, and F.Y.C. Mang. Synchronous and bidirectional component interfaces. In *CAV'02*, 2002.

[7] B. Gaudin and H Marchand. Supervisory control of product and hierarchical discrete event systems. *European Journal of Control*, 10(2), 2004.

[8] A. Girault and E. Rutten. Automating the addition of fault tolerance with discrete controller synthesis. *Int. journal on Formal Methods in System Design*, 2009. *(to appear).*

[9] N. Halbwachs. *Synchronous programming of reactive systems.* Kluwer, 1993.

[10] D. Harel and A. Naamad. The STATEMATE semantics of STATECHARTS. *ACM Trans. on Software Engineering and Methodology*, 5(4):293–333, Oct. 1996.

[11] Ch. Kloukinas and S. Yovine. Synthesis of safe, QoS extendible, application specific schedulers for heterogeneous real-time systems. In *5th Euromicro Conference on Real-Time Systems (ECRTS'03), Porto, Portugal, July*, 2003.

[12] O. Kushnarenko and S. S. Pinchinat. Intensional approachs for symbolic methods. In *Electronic Notes in Theoretical Computer Science*, volume 18, 1998.

[13] O. Maler, A. Pnueli, and J. Sifakis. On the synthesis of discrete controllers for timed systems. In E. W. Mayr and C. Puech, editors, *Proceedings STACS'95*, volume 900 of *Lecture Notes in Computer Science*, pages 229–242. Springer-Verlag, 1995.

[14] F. Maraninchi and Y. Rémond. Mode-automata: a new domain-specific construct for the development of safe critical systems. *Science of Computer Programming*, 46(3), 2003.

[15] H. Marchand, P. Bournai, M. Le Borgne, and P. Le Guernic. Synthesis of discrete-event controllers based on the Signal environment. *Discrete Event Dynamical System: Theory and Applications*, 10(4), October 2000.

[16] H. Marchand and B. Gaudin. Supervisory control problems of hierarchical finite state machines. In *Proc. of the 41th IEEE Conference on Decision and Contro,* USA, 2002.

[17] H. Marchand and M. Le Borgne. The supervisory control problem of discrete event systems using polynomial methods. Research Report 1271, Irisa, October 1999.

[18] H. Marchand and E. Rutten. Managing multi-mode tasks with time cost and quality levels using optimal discrete control synthesis. In *Proceedings of the 14th Euromicro Conference on Real-Time Systems, ECRTS'02,* June 2002, Austria, pages 241–248, 2002.

[19] S. Pinchinat and H. Marchand. Symbolic abstractions of automata. In *Proc of 5th Workshop on Discrete Event Systems, WODES 2000*, pages 39–48, Ghent, Belgium, August 2000.

[20] P. J. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of the IEEE; Special issue on Dynamics of Discrete Event Systems*, 77(1):81–98, 1989.

[21] J. Real and A. Wellings. Implementing mode changes with shared resources in ada. In *Proceedings of the* 11th Euromicro Conference on Real-Time Systems, ECRTS'99*, York, England, UK, June 9th - 11th*, 1999.

[22] I. Romanovski and P. Caines. Multi-agent product systems: analysis, synthesis and control. In *Proc of 6th Workshop on Discrete Event Systems*, Spain, October 2002.

[23] E. Rutten. A framework for using discrete control synthesis in safe robotic programming and teleoperation. In *Proc. IEEE Int. Conf. on Robotics and Automation, ICRA'2001,* Seoul, Korea, may, 2001.

[24] Eric Rutten and Hervé Marchand. Automatic generation of safe handlers for multi-task systems. Rapport de Recherche 5345, INRIA, October 2004. www.inria.fr/rrrt/rr-5345.html.

[25] D. Simon, M. Personnaz, and R. Horaud. TELEDIMOS: telepresence simulation platform for civil work machines: real-time simulation and 3d vision reconstruction. In *Proc. Workshop on Advances in Robotics for Mining and Underground Applications, Australia,*, 2000.

[26] S. Takai and T. Ushio. Supervisory control of a class of concurrent discrete event systems. *IEICE Transactions on Fundamentals*, E87-A(4):850–855, 2004.

A Figure captions

List of Figures

1 The considered general system architecture. 41

2 An example automaton. 42

3 An example of synchronous product (left), with the resulting automaton (right). . . 43

4 An example of hierarchical composition (left), with the resulting automaton (right). 44

5 Discrete control synthesis: from uncontrolled system (left) to closed-loop (right). . . 45

6 An example of automaton with costs. 46

7 Implementation of the approach: the tools involved. 47

8 Discrete models of a task controls. 48

9 Other patterns of tasks: *rejection and warning*, *request cancel*, *initialization phase*. . 49

10 Task pattern with multiple modes. 50

11 Tasks interfaces, and tasks server, with a controller and an application on top. . . . 51

12 Configurations reachable in one step from S_1, upon reception of ($\texttt{Req}_1$ `and` $\texttt{Go}_1$) and in the absence of $\texttt{Stop}_2$, according to controllables $C_{k_1}, C_{k'_2}$, with weights of cost and quality. 52

13 Sequence and parallel of two tasks. 53

14 An observer for an activation of t_3 between t_1 and t_2. 54

15 Cost and quality weights for modes. 55

16 Global cost and quality weights for the application. 56

17 Cost and quality weights for the tasks on the example. 57

18 Standard task `Amanu` in Mode Automata. 58

19 Panels for interactive simulation. 59

B Figures

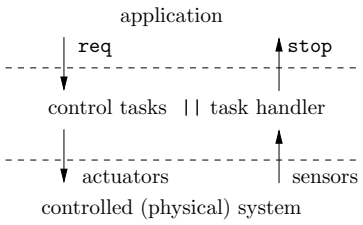

Figure 1: The considered general system architecture.

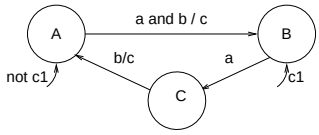

Figure 2: An example automaton.

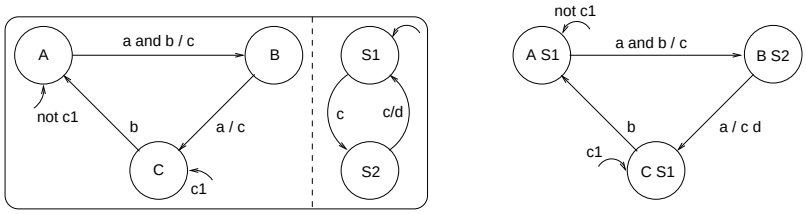

Figure 3: An example of synchronous product (left), with the resulting automaton (right).

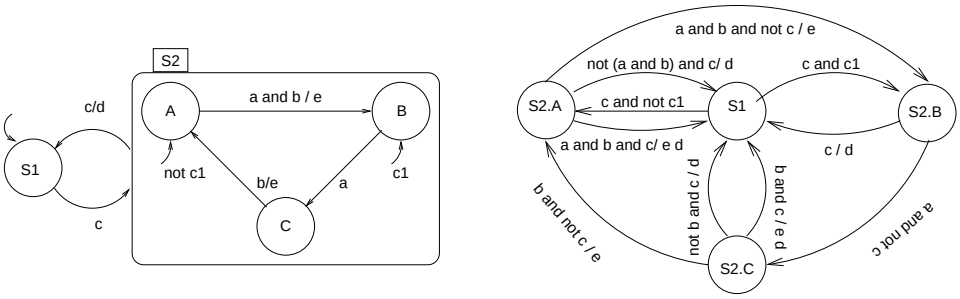

Figure 4: An example of hierarchical composition (left), with the resulting automaton (right).

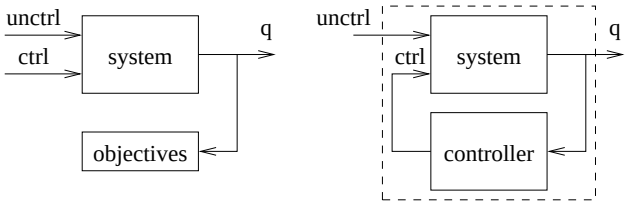

Figure 5: Discrete control synthesis: from uncontrolled system (left) to closed-loop (right).

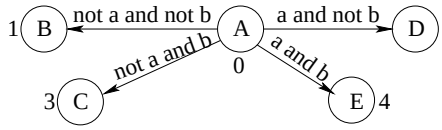

Figure 6: An example of automaton with costs.

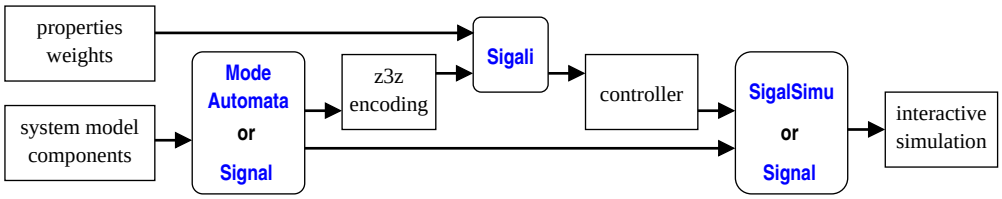

Figure 7: Implementation of the approach: the tools involved.

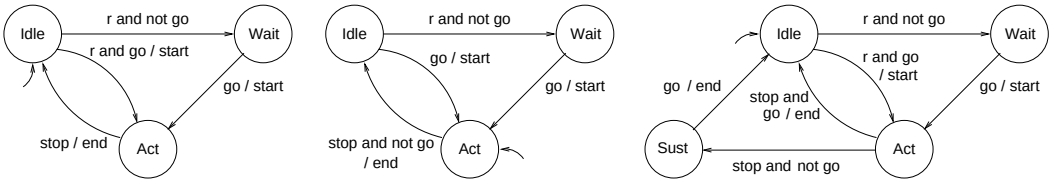

(a) *standard* task (b) *default* task (c) *sustainable* task

Figure 8: Discrete models of a task controls.

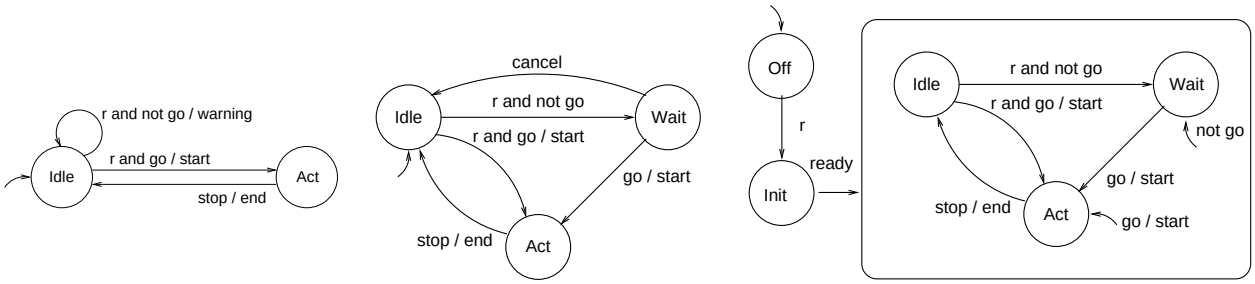

Figure 9: Other patterns of tasks: *rejection and warning*, *request cancel*, *initialization phase*.

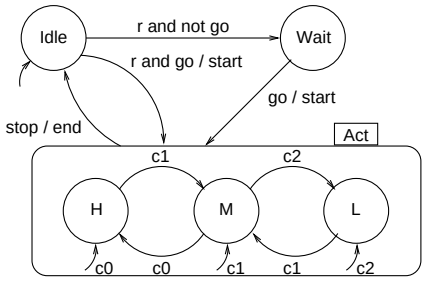

Figure 10: Task pattern with multiple modes.

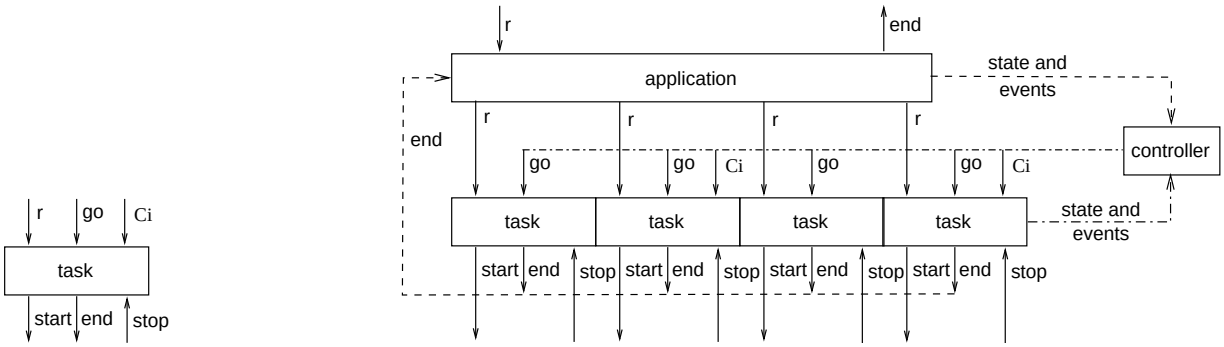

Figure 11: Tasks interfaces, and tasks server, with a controller and an application on top.

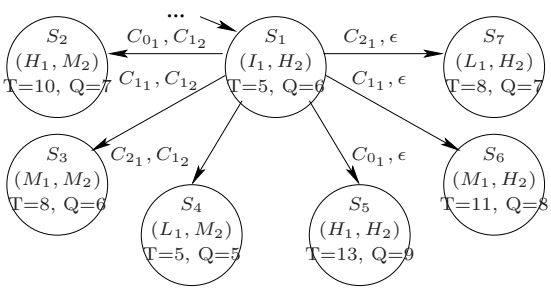

Figure 12: Configurations reachable in one step from S_1, upon reception of ($\texttt{Req}_1$ `and` $\texttt{Go}_1$) and in the absence of $\texttt{Stop}_2$, according to controllables $C_{k_1}, C_{k'_2}$, with weights of cost and quality.

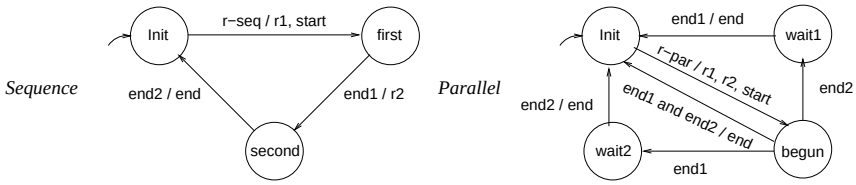

Figure 13: Sequence and parallel of two tasks.

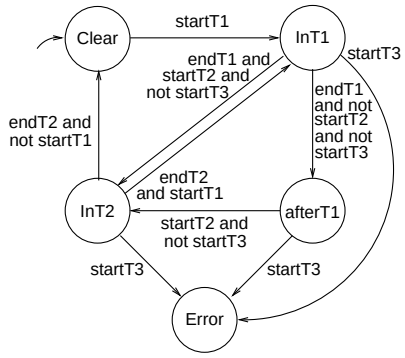

Figure 14: An observer for an activation of t_3 between t_1 and t_2.

		$\mathtt{I}_i$	$\mathtt{W}_i$	$\mathtt{H}_i$	$\mathtt{M}_i$	$\mathtt{L}_i$
T_1	C_t	0	0	7	5	2
	C_q	0	0	3	2	1
T_2	C_t	0	0	6	3	1
	C_q	0	0	6	4	3

Figure 15: Cost and quality weights for modes.

	S_1	S_2	S_3	S_4	S_5	S_6	S_7
T	5	10	8	5	13	11	8
Q	6	7	6	5	9	8	7

Figure 16: Global cost and quality weights for the application.

Gripper	cost	quality	Arm	cost	quality
`Gmanu`	5	4	`Aauto H`	8	8
`Gmaint`	3	3	`Aauto L`	6	6
`Gauto H`	4	5	`Ahome`	4	4
`Gauto L`	2	2	`Amanu`	7	7
			`Amaint`	4	4

Figure 17: Cost and quality weights for the tasks on the example.

```
AUTOMATON Amanu

STATES
 I_Amanu        init   [ start_Amanu=req_Amanu and go_Amanu; end_Amanu=false; ]
 W_Amanu               [ start_Amanu=go_Amanu; ]
 A_Amanu               [ end_Amanu=stop_A; start_Amanu=false; ]

TRANS
 FROM I_Amanu TO W_Amanu     [ req_Amanu and not go_Amanu ]
 FROM I_Amanu TO A_Amanu     [ req_Amanu and go_Amanu ]
 FROM W_Amanu TO A_Amanu     [ go_Amanu ]
 FROM A_Amanu TO I_Amanu     [ stop_A ]
```

Figure 18: Standard task `Amanu` in Mode Automata.

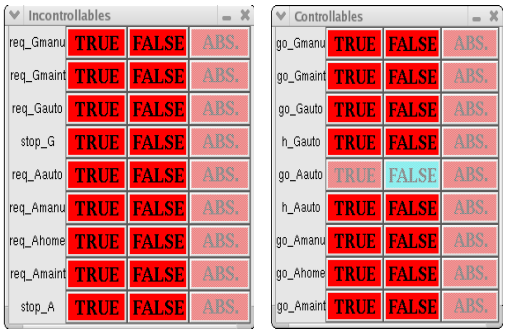

Figure 19: Panels for interactive simulation.