



HAL
open science

Automatic generation of discrete handlers of real-time continuous control tasks

Ahmed Soufyane Aboubekr, Gwenaël Delaval, Roger Pissard-Gibollet, Eric Rutten,
Daniel Simon

► **To cite this version:**

Ahmed Soufyane Aboubekr, Gwenaël Delaval, Roger Pissard-Gibollet, Eric Rutten, Daniel Simon. Automatic generation of discrete handlers of real-time continuous control tasks. IFAC WC 2011 - 18th IFAC World Congress, Aug 2011, Milan, Italy. <hal-00640406>

HAL Id: hal-00640406

<https://inria.hal.science/hal-00640406v1>

Submitted on 11 Nov 2011

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.



HAL Authorization

Automatic generation of discrete handlers of real-time continuous control tasks

Soufyane Aboubekr* Gwenaël Delaval* Roger Pissard-Gibollet*
Eric Rutten* Daniel Simon*

* INRIA Grenoble Rhône-Alpes, France
{firstname.lastname}@inria.fr

Abstract: We present a novel technique for designing discrete, logical control loops, on top of continuous control tasks, ensuring logical safety properties of the tasks sequencings and mode changes. We define this new handler on top of the real-time executives built with the ORCCAD design environment for control systems, which is applied, e.g. to robotics and real-time networked control. It features structures of control tasks, each equipped with a local automaton, used for the reactive, event-based management of its activity and modes. The additional discrete handler manages the interactions between tasks, concerning, e.g., mutual exclusions, forbidden or imposed sequences. We use a new reactive programming language, with constructs for finite-state machines and data-flow nodes, and a mechanism of behavioral contracts, which involves discrete controller synthesis. The result is a discrete control loop, on top of the continuous control loops, all integrated in a coherent real-time architecture. Our approach is illustrated and validated experimentally with the case study of a robot arm.

Keywords: real-time control (TC 3.1), discrete controller synthesis (TC 1.3), reactive systems

1. MOTIVATION: RTOS AND REACTIVE CONTROL

Control systems programming and real-time operating systems.

From the computing architecture point of view, a control system is made of a heterogeneous collection of physical devices, in continuous time, and information sub-systems, with discrete time scales. The physical devices, e.g. mechanical, electrical or chemical devices, are governed by the laws of physics and mechanics. Their input/output transfer characteristics exhibit a complex dynamic behavior (e.g. due to inertia) described by differential equations where time is a continuous variable. For their control, their state is measured or estimated using various sensors. Control theory provides a large set of methods and algorithms to govern their behavior through closed-loop control, ensuring the respect of required performance and crucial properties like stability.

Control systems are often implemented as a set of tasks running on top of a real-time operating system (RTOS). Closed-loop digital control systems use computers to cyclically sample sensors, compute a control law and send control signals to the actuators of the physical process. The performance of a control loop, e.g. measured by the tracking error, and more importantly its stability, strongly relies on the values of the sampling rates and sensor-to-actuator latencies (Åström and Wittenmark (1997)). A general rule states that smaller are the periods and latencies, better is the control performance. Thus it is essential that the implementation respects a specified timing behavior to meet the expected performance, i.e. the actual sampling periods and latencies must be fit in ranges which are consistent with the digital controller specification. ORCCAD is a design environment for such systems (Borrelly et al. (1998)).

Discrete, reactive controllers. Another level of control systems is more related to events and states, which define execution modes of the control system, typically with changes of control law. Reactive languages based on finite state automata, like StateCharts (Harel and Naamad (1996)), or StateFlow in Matlab/Simulink (Scaife et al. (2004)), are widely used for these aspects. Their underlying fundamental model, transition systems, is the basic formalism for discrete control theory. Amongst reactive languages, the synchronous languages (Benveniste et al. (2003)): Lustre, Esterel or Lucid Synchrone are used industrially in avionics and safety-critical embedded applications design (Esterel technologies (2010)). They offer a coherent framework for specification, compilation, and distributed code generation, test generation and verification.

In the framework of discrete control theory, a basic technique used for the design of control loops is supervisory control, with Discrete Controller Synthesis (DCS) algorithms (Ramadge and Wonham (1987); Cassandras and Lafortune (1999)). It can be applied on a controllable system, for a given behavioral property. It consists in computing a constraint on this system, so that the composition of the system and this constraint satisfies the property. There is a tool able of automated DCS (Marchand et al. (2000)), which is concretely connected to reactive languages and has been applied to the modeling of automatic generation of task handlers (Marchand and Rutten (2002)). The BZR language has been defined with a contract mechanism, which is a language-level integration of DCS (Aboubekr et al. (2009); Delaval et al. (2010)): the user specifies possible behaviors of a component, as well as safety constraints, and the compiler synthesizes the necessary control to enforce them. The programmer does not need to design it explicitly, neither to know about the formal technicalities of the encapsulated DCS.

¹ Work partially funded by the FeedNetBack FP7-ICT European project.

At missions design level, the RT automaton gives an abstract view which facilities their composition into more complex actions: the **Robot-Procedures** (RPs). The RP paradigm is used to logically and hierarchically compose RTs and RPs, designed to fulfill a basic goal through several possible modes, e.g. a mobile robot can follow a wall using predefined motion planning, visual servoing, or acoustic servoing according to sensory data availability. RPs design is hierarchical so that common structures and programming tools can be used from basic actions up to a full mission specification.

For specification and validation the original ORCCAD framework uses Esterel (Esterel technologies (2010)) for each RT and RP logical behavior design, verification and code generation. The global behavior is defined by the parallel composition of the automata. The synchronous technology enables the use of formal techniques for automatic verification of the behavior, for liveness and safety properties. For example, a safety property specifically related with control systems states that every physical actuator must be always under control, by one and only one control law. More specific properties can also be defined and validated for various case studies.

For the execution machine for the automata, besides the user-defined signals (pre and post-conditions, exceptions), ORCCAD also defines many signals used at run time to spawn and manage all the real-time threads necessary for the execution of the tasks and procedures. The current ORCCAD Esterel automata are compiled into a transition function in C. Input and output functions are associated to received and emitted signals, which are used to interface the synchronous reactive program with the asynchronous execution environment, i.e. the RTOS. Numerical computations can be called in linked libraries. The execution machine is in charge of feeding the automaton with signals synthesized from collected input events, running the automaton transition and exporting the output actions to the system. The automaton and execution machine are further compiled into a real-time task and event queue glued with the rest of the system, as depicted in section 4.4.

The position of the contribution in this paper is that, until now, in ORCCAD, the discrete events control code is designed as a computer programming work, written manually, then formally verified. One drawback is the difficulty for control engineers users of specifying the discrete control without a methodology related to control theory, and the intrication of verification techniques. Another is that static manual programming of all cases fails to encompass adaptive behavior, with regulation w.r.t. the system's state and available resources. This papers addresses these issues by considering discrete control loops on top of the continuous control loops.

3. PROGRAMMING REACTIVE SYSTEMS IN BZR

We introduce first the basics of the Heptagon language, to program data-flow nodes and hierarchical parallel automata (Colaço et al. (2005)). As for the reactive languages introduced in Section 1, the basic execution scheme is that at each reaction a step function is called, taking input flows as parameters, computing the transition to be taken, updating the state, triggering the appropriate actions, and emitting the output flows. We then define the BZR language which extends Heptagon with a new contract construct (Delaval et al. (2010)).

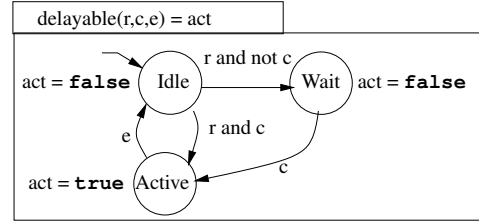


Fig. 2. Example of a node in graphical syntax.

3.1 Data-flow nodes and mode automata

Figure 2 shows a simple example of a Heptagon node, for the control of a task that can be activated by a request r , and according to a control flow c , put in a waiting state; input e signals the end of the task. Its signature is defined first, with a name, a list of input flows (here, simple events which can be seen as Boolean flows), and outputs (here: the Boolean act), which is true when the task is active. In the body of this node we have a mode automaton : upon occurrence of inputs, each step consists of a transition according to their values; when no transition condition is satisfied, the state remains the same. In the example, `Idle` is the initial state. From there transitions can be taken towards further states, upon the condition given by the expression on inputs in the label. Here: when r and c are true then the control goes to state `Active`, until e becomes true, upon which it goes back to `Idle`; if c is false it goes towards state `Wait`, until c becomes true. This is a mode automaton (Colaço et al. (2005)) in the sense that to each state we associate equations to define the output flows. In the example, the output act is defined by different equation in each of the states.

We can build hierarchical and parallel automata, as will be seen in the case study e.g., in Figure 12 In the parallel automaton, the global behavior is defined from the local ones: a global step is performed synchronously, by having each automaton making a local step, within the same logical instant. In the case of hierarchy, the sub-automata define the behavior of the node as long as the upper-level automaton remains in its state.

3.2 Contracts in the BZR language

With this new construct, the management of dynamical adaptivity can be considered as a control loop, on continuous or discrete criteria. It is illustrated in Figure 3(a): on the basis of monitor information and of an internal representation of the system, a control component enforces the adaptation policy or strategy, by taking decisions w.r.t. the adaptation or reconfiguration actions to be executed, forming a closed control loop. The design of control loops with known behavior and properties is

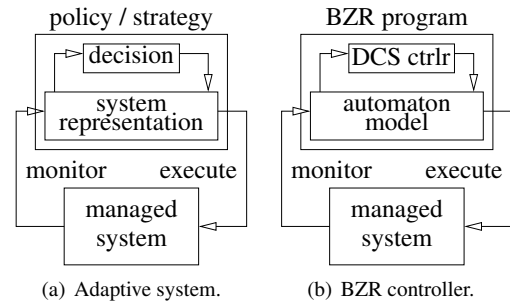


Fig. 3. BZR programming of adaptation control.

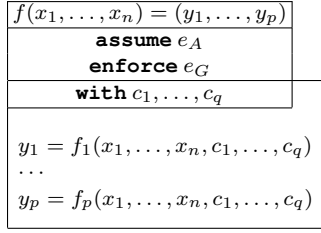


Fig. 4. BZR contract node graphical syntax

the classical object of control theory. Applications of continuous control theory to computing systems have been explored quite broadly. In contrast, qualitative or logical aspects, as addressed by discrete control theory, have been considered only recently for adaptive computing systems (Wang et al. (2010)). In our new approach, DCS is encapsulated in the compilation of BZR (Delaval et al. (2010)). Models of the possible behaviors of the managed system are specified in terms of mode automata, and adaptation policies are specified in terms of contracts, on invariance properties to be enforced. Compiling BZR yields a correct-by-construction controller, produced by DCS, as illustrated in Figure 3(b), in a user-friendly way: the programmer does not need to know technicalities of DCS.

As in the upper box of Figure 4, we associate a *contract* to a node. It is itself a program, with its internal state, e.g., automata, observing traces (for example an error state where e_G is false, to be kept outside an invariant subspace). It has two outputs: e_A , *assumption* on the node environment, and e_G , to be guaranteed or *enforced* by the node. A set $C = \{c_1, \dots, c_q\}$ of local controllable variables is used for ensuring this objective. This contract means that the node will be controlled, i.e., that values will be given to c_1, \dots, c_q such that, given any input trace yielding e_A , the output trace will yield e_G . This is obtained automatically, at compilation, using DCS.

Briefly (see details in Delaval et al. (2010)), we compile such a BZR contract node into a DCS problem as in Figure 5, which shows the transition systems as equations on inputs, outputs and states. The body and the contract are each encoded into a state machine with transition function (resp. TrC and TrB), state (resp. StC and StB) and output function (resp. $OutC$ and $OutB$). The contract inputs XC come from the node's input X and the body's outputs Y , and it outputs e_A, e_G . Assuming e_A produced by the contract program, DCS will obtain a controller $Ctrlr$ for the objective of enforcing e_G . The concrete type of property which is enforced is making invariant the sub-set of states where $e_A \Rightarrow e_G$ is true, by constraining controllable variables c_1, \dots, c_q . The controller then takes the states of the body and the contract, the node inputs X and the contract

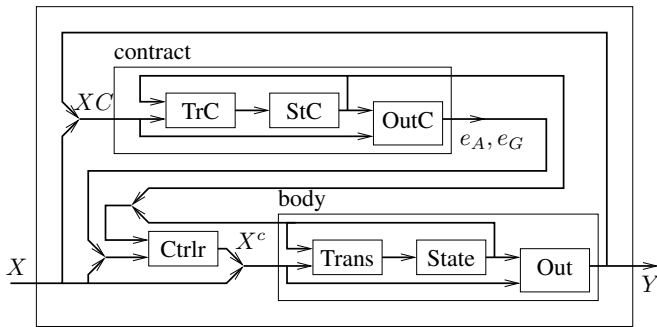


Fig. 5. BZR contract node as DCS problem

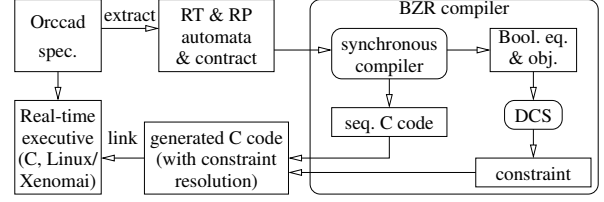


Fig. 6. Development process for BZR with ORCCAD.

outputs e_A, e_G , and it computes the controllables X_c such that the resulting behavior satisfies the objective.

4. DISCRETE CONTROL HANDLERS OF CONTINUOUS CONTROL TASKS

4.1 Integration in a development process

As announced in Section 1, our first contribution is the integration of BZR reactive controllers, using DCS, into the ORCCAD runtime. The general scheme for using BZR consists of a treatment of the control part, using our target-independent language and compiler, in derivation of the main system development process. In its instantiation for the case of ORCCAD, illustrated in Figure 6, one can see phases of:

- extraction of control part from the adaptive system, in the form of a BZR program;
 - BZR compilation: synchronous compilation to:
 - a Boolean equations form, with contracts compiled into DCS objectives; given to DCS to produce the constraint on controllables;
 - a sequential C code for the automata;
- both are then assembled into an executable involving a resolution of the synthesized constraint;
- re-linking of the latter into the global executive.

4.2 General architecture

The contribution of this paper is a novel method for designing discrete, logical control handlers, on top of continuous control tasks. The goal is to ensure, by a discrete control loop, logical safety properties of the tasks sequencings and mode changes. We contribute this new layer on top of the real-time executives built with the ORCCAD design environment for control systems, by establishing the connection with the BZR language and compiler, which is relying upon DCS techniques.

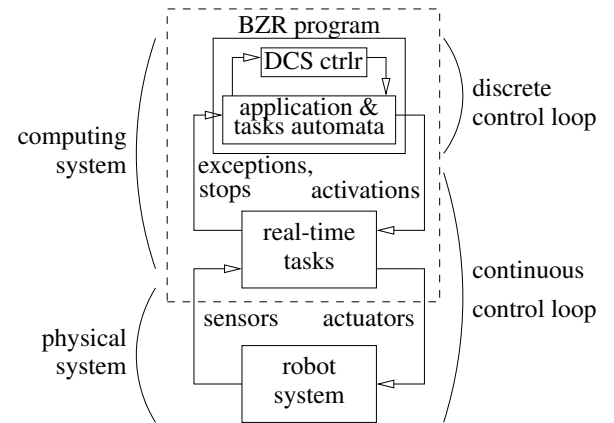


Fig. 7. Discrete control handlers of continuous control tasks.

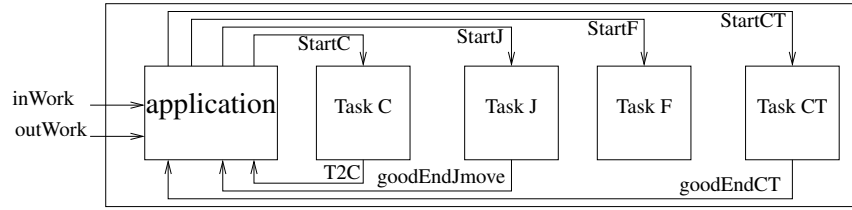


Fig. 9. Complete BZR program (simplified).

sition of the tasks automata, and of the application automaton. Figure 9 illustrates this for the case study.

4.4 Executive-level integration

At this level, we have to interface the code generated by the BZR compiler, as shown in Figure 6, with the ORCCAD-generated real-time executive mentioned in Section 2.2. It implements the transition step function, to be called at the appropriate pace, with appropriate input parameters, and handling of outputs. The implementation of this execution machine (i.e. of the dotted box in Figure 7) is sketched as shown in Figure 10.

A main task sets up the whole system. It spawns all the real-time tasks and associated communication and synchronization objects. In particular it generates the needed clocks used to trigger the cyclic calculation modules. Real-time threads are made cyclic by blocking their first input port on a semaphore which is released by clock ticks. Otherwise they can be triggered by any other event, such as a data production from another thread or a signal sent by a driver.

The automaton is the highest priority task : it is awakened by the occurrence of input signals related to the execution of the controllers, e.g. preconditions, exceptions, and post-conditions issued by the feedback controllers. All events are serialized and received on a FIFO input events queue. In reaction, the automaton tells the RTOS what action must be taken by releasing the corresponding semaphore. Using a model based approach all the glue code is automatically generated, while using only basic features of operating systems make easier porting the tools for targets such as Linux/Posix threads and Xenomai.

Although this automaton is crucial for a safe and successful behavior of the application, it spends most of time doing nothing, just waiting for input events during the cyclic execution of the control algorithms managed by the RTOS. Moreover its transitions take very short times (typically some μsecs) so that the overhead due to discrete events control is negligible.

5. CASE STUDY OF A ROBOT ARM

5.1 Description of the case study

We define a robot arm, called *ArmX*, which is a two-link manipulator with rotational joints (q_1, q_2) shown on Figure 11. Each link i ($\{1,2\}$) has a point masses M_i ($\{1,2\}$) at the end of links. The dynamic model of the manipulator can be written in the form: $\tau = M(q)\ddot{q} + V(q, \dot{q}) + G(q)$ where $M(q)$ is the 2×2 mass matrix of the manipulator, $V(q, \dot{q})$ is an 2×1 vector of centrifugal and Coriolis terms, $G(q)$ is an 2×1 vector of gravity terms and τ the input joint torque. For this simple manipulator all details of calculation can be found in Graig (1989).

ArmX is equipped with a robotic tool changer which allows the robot to switch end effector. Two tools are manipulated by the

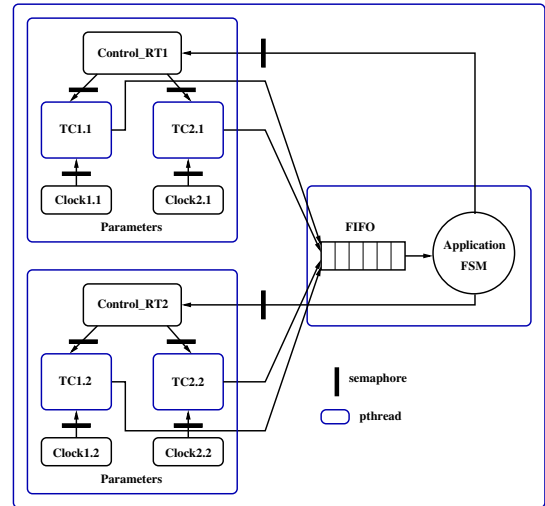


Fig. 10. Implementation of the execution machine.

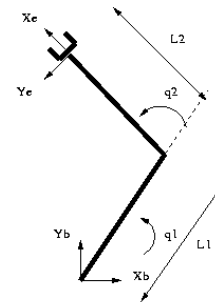


Fig. 11. The ArmX model.

arm, one when the target is inside the robot workspace (for example a gripper) and the second outside of this space (for example a proximity sensor to point the target).

The ORCCAD Robot-Tasks for this application are four control-laws, embedded in four RTs:

- the joint space control task *ArmXjmove* controls the move in the joint space of the manipulator i.e., in terms of values of angles at the joints;
- the Cartesian space control task *ArmXcmove* controls the move in the Cartesian space of the manipulator, in terms of 3d coordinates; it is appropriate for aiming at targets *inside* the workspace.
- the target aiming task *ArmXfmove* controls the pointing towards a point by trajectory following; it is appropriate for aiming at targets *outside* the workspace.
- the tool change task CT first brings the robot to its initial position ($q_1 = 0, q_2 = 0$), in order to then switch the end effector tool.

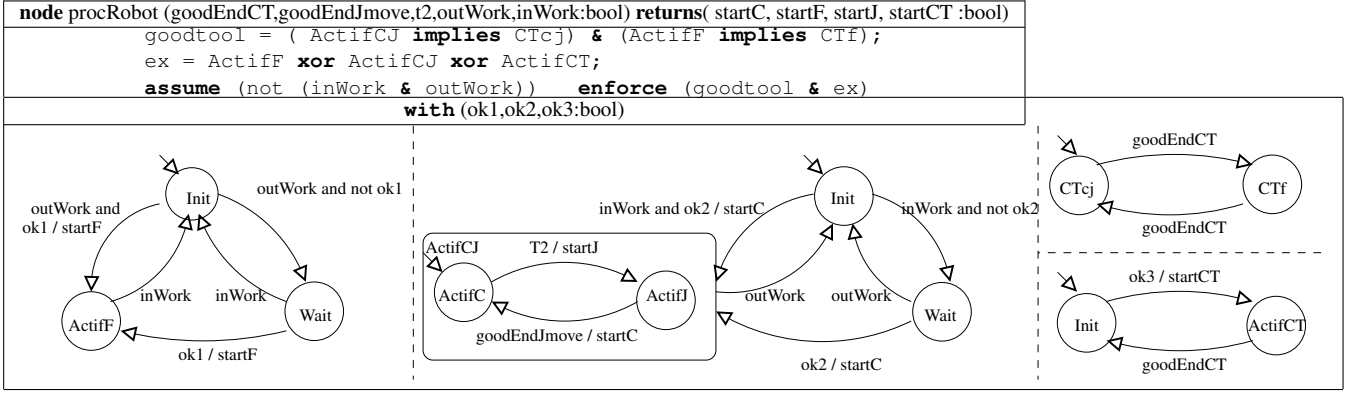


Fig. 12. Global BZR node, with contract.

The simulation environment for our case study represents the dynamics on the two-link manipulator *ArmX* modeled previously. We use its inverse dynamic model to compute joint accelerations: $\ddot{q} = M^{-1}(q)(\tau - V(q, \dot{q}) - G(q))$ and we obtain the current q and \dot{q} by a double Euler integration. So, from ORCCAD or another application, this simulator is perceived like a real robot; we have functions to initialize it, to put torque, to get joint position, etc.

The application designed is a target following task. When the target is inside the robot workspace, the effector follows the target. When it is outside of the robot workspace the manipulator point towards this target. This application must be safe and so it is performed taking into account exceptions like the tracking error is too high, joints limit are reached, or reconfiguration arm is required.

The objective is that the arm automatically changes to the appropriate tool, according to the target being inside or outside the workspace. The fact that the tool change task is inserted automatically in function of the current situation makes it an adaptive system.

To each task corresponds an instance of the generic task control automaton; for the case of the *ArmXcmove* task the automaton is shown in Figure 8. Each of the three other tasks is associated with a similar one. All are featured in the global controller as shown in Figure 9.

5.2 The application *RP* and its global control

We apply the BZR programming methodology: first describe possible behaviors, then specify control objectives in the contract. The application must launch robot tasks corresponding to the current state of the target (inside or outside the workspace) and change the tool arm to get the right tool for each task. So the control objective is first to ensure we have the right tool, and second, to check the smooth running of the application, i.e., allowing at most task to be active at a time, and also at least one, as mentioned in Section 2.2. A set $C = \{ok_1, ok_2, ok_3\}$ of local controllable variable will be used for ensuring this objective. The contract specifies that the node will be controlled, i.e., that values will be given to ok_1, ok_2, ok_3 such that, given any uncontrollable input trace, the output trace will satisfy the two objectives i.e., yield the true value for variables *goodtool* and *ex*. This is done by DCS.

It is named `procRobot`, and illustrated in Figure 12. The node is connected with the implementation shown in Figure 10 by

the inputs and outputs declared in the top part. The body in the lower part is composed of 4 parallel automata, described from left to right:

- the automaton for the F task: it can start the *ArmXfmove* task, by emitting `startF`, when it receives the signal `outWork` and obtains the permission of the controller by the flow `ok1`; if `ok1` is false, then it goes to state `Wait`, until `ok1` becomes true. It models the choice to delay the starting of F, and corresponds to the delayable tasks pattern illustrated in Figure 2.
- the automaton for the C and J tasks: it is hierarchical with two levels. The upper level is also an instance of the delayable task pattern; the Boolean `ok2` is used to mark the choice point.

The sub-automaton is in the `ActifCJ` state manages the alternation between C and J tasks. Upon occurrence of an exception of type T2 in task C, it gives control to the task J. This is a way of handling singularities, which are points that can't be reached by using the control laws of task C: in this case control is given to task J, by sending `startJ`, to reposition the arm to reach this point. At its end a signal `goodEndJmove` is received from the RT, then task C is started again.
- the automaton observing the current tool state (top) is used to memorize the current tool of the arm. It has two states corresponding to two tools manipulated by the arm, the first one is used in the workspace accessible by the arm, and the other in outside. Every change of tool this automaton receives a `goodEndCT` signal from the RT automaton to indicate that the task ended well.
- the automaton for the CT task (bottom) is modeling the fact that it can be triggered by the controller that will be synthesized. Using controllable variable `ok3`, the controller can force the tool change by sending `startCT`, if the arm does not have the tool corresponding to the task for the current position of the target.

This parallel automaton describes the possible sequencings of the tasks. It can be noted that it does not explicitly care for their exclusion, or for managing the appropriateness of the tool. This is shown in the declarative contract, and compiled with DCS.

The contract is in the upper part of Figure 12: it is a program, with equations. Three controllable variables, defined in the `with` part, will be used for ensuring two objectives:

- the right tool for the right task: a Boolean variable `goodtool` is defined, as the conjunction of two impli-

cations: they state that when a task is active (`ActifCJ`, respectively `ActifF`), it implies that the arm carries the right tool (`CTcj`, respectively `CTf`).

- Mutual exclusion and default control: an equation defines `ex`, which is the exclusive disjunction of active states for the tasks. it means actually two things: that there is at most one active task, and also at least one, so that the arm is always controlled, as mentioned in Section 2.2.

The contract is that, assuming that the target can not be inside and outside of the workspace at the same time, control enforces that the two Boolean are true.

5.3 Simulation and typical scenario

Here is a typical scenario showing the intervention of the controller on the system, so that control objectives are preserved. At some point the task `CJmove` is active, and the target inside the workspace, and the tool carried by the arm corresponds to state `CTcj`. Then, the user clicks outside of the workspace, so the application receives the `outWork` input. This causes the automaton for `CJ` to move by a transition to its initial state.

It also causes the automaton for task `F` to quit its initial state; here, we have a choice point conditioned by `ok1`. Due to the first contract property, `goodtool` must be kept true, so given that the current tool state is `CTcj` the controller can not allow the transition to `ActifF`, and must give the value `false` to `ok1`. Hence task `F` goes into `Wait` state. Due to the other contract property, `ex` must be kept true, which forces the controller to maintain at least one active state. Therefore it launches the task `CT` using the controllable variable `ok3`, which will change the tool. At the end of the task `CT`, the `goodEndCT` event allows the automaton observing the current tool to pass in the state `CTf`. Thus we have the right tool for task `F`, and the controller can release `F` from `Wait` to `ActifF`, by giving value `true` to controllable variable `ok1`. This shows how mutual exclusion, and insertion of reconfiguration tasks can be obtained declaratively.

6. CONCLUSION AND PERSPECTIVES

We propose the integration of a set of technique to design discrete control loops on top of continuous control tasks. They ensure logical safety properties of the tasks sequencings and mode changes. The implementation integrates ORCCAD, a real-time control executives design environment, and the BZR reactive language, encapsulating the formal DCS technique in its compilation. A case of a robot arm is studied. It constitutes a concrete approach to implementing hybrid systems.

The DCS technique manages discrete control loops, where objectives are not quantitative or numerical, but logical properties. Therefore its advantage is not evaluated in terms of numerical performance, but rather in that it ensures correctness of the design of the discrete loop. Here, correctness means that the DCS algorithm computes a controller which satisfies the logical property. In particular, the concrete type of property which is enforced is making invariant the sub-set of states where the contract of the program is true, by constraining controllable variables.

Further work includes consolidating the integration of ORCAD and BZR beyond this case study, enriching the models with more quantitative aspects Marchand and Rutten (2002),

defining libraries of control models and contracts, and considering the more involving example of a Mars rover.

REFERENCES

- S. Aboubekr, G. Delaval, and E. Rutten. A programming language for adaptation control: Case study. In *Proc. of the 2nd Workshop on Adaptive and Reconfigurable Embedded Systems, APRES'09*, 2009.
- K. J. Åström and B. Wittenmark. *Computer-Controlled Systems*. Information and System Sciences Series. Prentice Hall, third edition, 1997.
- C. Aubrun, D. Simon, and Y.-Q. Song, editors. *Co-design approaches for dependable networked control systems*. ISTE-Wiley, 2010.
- 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*, 91(1), January 2003.
- 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.
- C. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Kluwer Acad. Publ., 1999.
- A. Cervin, J. Eker, B. Bernhardsson, and K.-E. Årzén. Feedback-feedforward scheduling of control tasks. *Real-Time Systems*, 23(1-2):25-53, July 2002.
- J.-L. Colaço, B. Pagano, and M. Pouzet. A Conservative Extension of Synchronous Data-flow with State Machines. In *ACM Int. Conf. on Embedded Software (EMSOFT'05)*, September 2005.
- G. Delaval, H. Marchand, and E. Rutten. Contracts for modular discrete controller synthesis. In *Proc. of the ACM Conf. on Languages, Compilers and Tools for Embedded Systems, LCTES 2010*, 2010.
- Esterel technologies. Scade: model-based development environment dedicated to safety-critical embedded software, 2010. URL <http://www.esterel-technologies.com>.
- J. Graig. *Introduction to Robotics, mechanics and Control*. Addison-Wesley Publishing Company, 1989.
- D. Harel and A. Naamad. The statemate semantics of statecharts. *ACM Trans. Softw. Eng. Methodol.*, 5(4):293-333, 1996. ISSN 1049-331X.
- H. Marchand and E. Rutten. Managing multi-mode tasks with time cost and quality levels using optimal discrete control synthesis. In *Proc. of the 14th Euromicro Conf. on Real-Time Systems, ECRTS'02*, 2002.
- H. Marchand, P. Bournai, M. Le Borgne, and P. Le Guernic. Synthesis of discrete-event controllers based on the Signal environment. *Discrete Event Dynamic System: Theory and Applications*, 10(4), October 2000.
- P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete event processes. *SIAM J. on Control and Optimization*, 25(1):206-230, January 1987.
- N. Scaife, C. Sofronis, P. Caspi, S. Tripakis, and F. Maraninchi. Defining and translating a "safe" subset of simulink/stateflow into Lustre. In *EMSOFT '04: 4th ACM Int. Conf. on Embedded software*, 2004.
- Y. Wang, H.K. Cho, H. Liao, A. Nazeem, T. Kelly, S. Lafortune, S. Mahlke, and S.A. Reveliotis. Supervisory control of software execution for failure avoidance: Experience from the gadara project. In *Proc. of WODES (Workshop on Discrete Event Systems)*, 2010.