



A Formal Approach to Agent Design: An Overview of Constraint-Based Agents

ALAN K. MACKWORTH

mack@cs.ubc.ca

Lab for Computational Intelligence, Department of Computer Science, University of British Columbia, Vancouver, B.C. V6T 1Z4, Canada

YING ZHANG

yzhang@parc.com

Systems and Practices Lab, Palo Alto Research Center, Palo Alto, CA 94304, USA

Abstract. Formal models for agent design are important for both practical and theoretical reasons. The Constraint-Based Agent (CBA) design approach includes two formal models: Constraint Nets and Timed \forall -automata. A constraint net models the agents and the environment symmetrically as, possibly hybrid, dynamical systems; a timed \forall -automaton specifies the desired real-time dynamic behaviors of the situated agents. Given a constraint-based specification of the desired behavior, a constraint-based agent can be synthesized as a constraint solver. Using formal modeling and specification, it is also possible to verify complex agents as obeying real-time temporal constraint specifications. This overview paper presents a summary of the development and application of the CBA framework.

Keywords: formal specification, agent design, artificial intelligence, robot architectures, constraint-based requirements, system verification, automata

1. Introduction and Motivation

Constraint programming has evolved several powerful frameworks for building problem-solvers as constraint-satisfying devices. Primarily, these devices are offline problem-solvers. For example, the Constraint Satisfaction Problem (CSP) paradigm has evolved and matured over the last twenty-five years. The algorithms developed in the CSP paradigm were made more available and more useful when they were incorporated into the Constraint Programming (CP) language paradigms. Despite this success, however, a major challenge still facing the constraint research community is to develop useful theoretical and practical tools for the constraint-based design of embedded intelligent systems. Many applications require us to develop online constraint-satisfying systems that function in a dynamic, coupled environment [16]. An archetypal example of an application in this class is the design of controllers for sensory-based robots [17, 24, 31]. If we examine this problem we see that almost all the tools developed to date in the CSP and CP paradigms are inadequate for the task, despite the superficial attraction of the constraint-based approach. The fundamental difficulty is that, for the most part, the CSP and CP paradigms still presume a disembodied, offline model of computation.

Consider an agent coupled to its active environment as shown in Figure 1. Each is an open dynamic system in its own right, acting on, and reacting to, the other. The coupled pair forms a closed system that evolves over time.

To deal with such embedded applications, we must radically shift our perspective on constraint satisfaction from the offline model in which a solution is a function of pre-given




Figure 1. An agent interacting with its environment.

static inputs to an online model where a solution is a temporal trace of values, resulting from a transduction of the input trace over time. Values in the agent's input trace may depend on earlier values in its output trace. In fact, the input trace for the agent is itself a causal mapping of its output trace, representing the dynamics of the environment, as shown in Figure 1.

In this paper we present an overview of our formal approach to specifying, designing, building and evaluating embedded intelligent systems: the Constraint-Based Agent (CBA) framework.

Agent-based systems embedded as controllers in real or virtual environments must be designed in an online model based on various time structures: continuous, discrete and event-based. The requisite online computations, or transductions, are to be performed over various type structures including continuous and discrete domains. Such time and domain structures are modeled formally in Constraint Nets (CN) [32], a model for hybrid dynamic systems. CN provides a new view of computation, constraint satisfaction and constraint programming. To this end, we have defined constraint satisfaction as a dynamic system process that approaches asymptotically the solution set of the given, possibly time-varying, constraints [33]. Under this view, constraint programming is the creation of a dynamic system with the required property.

A formal model for the agent systems alone does not complete the whole framework. In order to design agent systems that guarantee some desired behaviors, a formal behavior specification is required. We have developed timed \forall -automata for specifying behavioral requirements for constraint-based agents. The timed \forall -automata model was developed in [28, 35, 40] as an extension of discrete time \forall -automata [19] to continuous time, with annotations for real-time. Timed \forall -automata are simple, yet able to represent many important features of dynamic systems such as safety, stability, reachability and real-time response.

Given a constraint-based specification of the desired behavior, a constraint-based agent can be synthesized as a constraint solver. Using formal modeling and specification, it is also possible to verify complex agents as obeying real-time temporal constraint specifications.

The agent-based view of design has been taken in the research communities of multi-agent systems and distributed AI [9, 10, 27]. There, however, most of the work has been focused on discrete, rather than discrete/continuous hybrid, systems. Furthermore, even though online real-time intelligent behaviors have been explored, few have emphasized the importance of offline verification of the overall system.

Over the last decade, there has been a growing interest in hybrid dynamic systems. In this community, much work has also been done in formal approaches to system specification and verification [2, 3, 11, 19]. In general, there are two schools. One is to develop a uniform specification for both systems and their requirements; the other is to use two different specifications, one for systems and one for requirements. The advantage of the former is that the same formal approach can apply to both system synthesis and system verification. However, in most cases, if the specification language is powerful for both systems and requirements, the synthesis or verification tasks become hard. We advocate the latter approach, that is, using timed \forall -automata for requirements specification and using Constraint Nets [28, 32, 35] for system modeling. Control synthesis [28, 34] and verification [28, 30, 31, 35, 39] have also been studied in this framework.

There have also recently been studies on heterogeneous modeling and design for concurrent and embedded systems, such as UC Berkeley's Ptolemy Project [21]. The system uses a mix of discrete, concurrent process models with continuous differential equations. The result is a very large and complex system. Choosing the right models and determining the interactions among the mix of models are nontrivial.

The idea of using constraint solvers as controllers has also been explored in the control community, where it comes under the rubric of Model Predictive Control (MPC) [6, 13]. There, the control objective is represented as an optimization function over an infinite time horizon. At each control step, the controller solves a constrained optimization problem to find the best control value based on the model and the objective over a finite horizon. This type of computation, in general, is either for linear systems, or so time and resource expensive that it is hard to implement for real-time embedded systems.

In this paper we present an overview of our approach, called Constraint-Based Agents. Even though our framework is formal, the presentation in this paper is intended to be informal, emphasizing the concepts and methods. For details of the formal modeling, specification and verification, please read the cited references. In the rest of this paper, we first discuss the design of embedded agents. The essence of the CBA model consists of a symmetrical coupling of an agent and its active environment. We then describe the Constraint Net model and introduce timed \forall -automata. We say that an agent is constraint-based if its behavior satisfies a specification in a well-defined constraint-based temporal logic or timed \forall -automaton. We then discuss general synthesis and verification procedures for this type of requirements specification. We finally describe the use of the CBA framework in the application of robot soccer, and finish with our conclusions.

2. Agents in the World

The most obvious artificial agents in the world are robots. However, the CBA approach applies equally to embedded devices, pure software agents and natural animate agents. There are many ways of using a CBA model, including the embedded mode, simulation mode, verification mode, optimization mode, learning mode and design mode [22, p. 449]. The agent design problem is formidable, regardless of whether the agent is designed or modified by a human (programming), by nature (evolution), by another

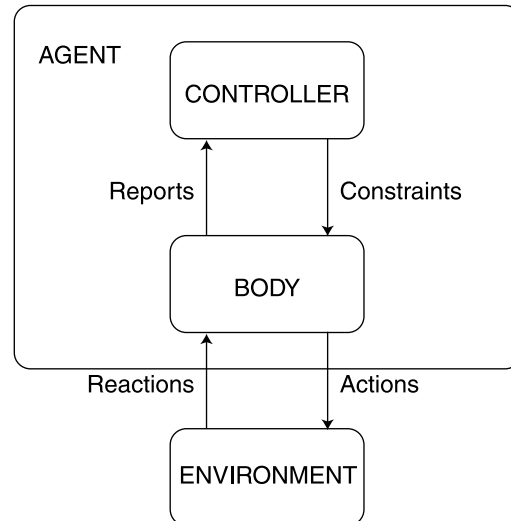


Figure 2. The structure of a constraint-based agent system.

agent (bootstrapping), or by itself (learning). An agent is, typically, an intelligent hybrid dynamic system, consisting of a controller coupled to its body as shown in Figure 2.

The controller and the body each consist of discrete-time, continuous-time or event-driven components operating over discrete or continuous domains. The controller has perceptual subsystems that can (partially) observe the state of the body, and, through it, the state of the environment.

Parenthetically, the ‘body’ of an agent is simply the direct interface of the agent to its environment. The body executes actions in the environment, senses state in the environment (which may well cause state changes in the body) and reports to the controller. In the case of a robotic agent, the body consists of one or more physical systems but in the case of an embedded software agent, the body is simply the software module that directly interfaces to the virtual or physical environment. Control theorists typically call the body the ‘plant’. Some models do not differentiate between the body and the environment; we prefer to make that differentiation, based on the distinction between what is directly, and what is indirectly, controlled.

Agent design methodologies are evolving dialectically [14]. The symbolic methods of ‘Good Old Fashioned Artificial Intelligence and Robotics’ (GOFAIR) constitute the original thesis. The antithesis is reactive ‘Insect AI’ and control theory. The emerging synthesis, Situated Agents, has promise, but needs formal rigor and practical tools [8, 12, 14, 17, 18, 23, 31].

In 1992, Mackworth proposed robot soccer as a grand challenge problem [14] since it has the task characteristics that force us to confront the fundamental issues of agent design in a practical way for a perceptual, collaborative, real-time task with clear performance criteria. At the same time, Mackworth described the first system for playing robot soccer. Since then it has been a very productive testbed, stimulating research toward the goal of building perceptual agents, both in our laboratory [5, 15, 24, 34, 36, 37, 39, 41] and in many other groups around the world.

3. The Constraint Net Model

The Constraint Net (CN) model [32] was developed by Ying Zhang and Mackworth as a model for building intelligent hybrid systems as Situated Agents. In CN, an agent system is modelled formally as a symmetrical coupling of an agent with its environment. Even though an agent system is, typically, a hybrid dynamic system, its CN model is unitary. Most other agent and robot design methodologies use hybrid models of hybrid systems, awkwardly combining offline computational models of high-level perception, reasoning and planning with online models of low-level sensing and control.

CN is a model for agent systems software implemented as modules with I/O ports. A module performs a transduction from its input traces to its output traces, subject to the principle of causality: an output value at any time can depend only on the input values before, or at, that time. The model has a formal semantics based on the least fixpoint of sets of equations [32]. In applying it to an agent operating in a given environment, one separately specifies the system of the agent body, the agent control program, and the environment. The total system can then be simulated and/or shown to have various global properties, such as safety and liveness. This approach allows one to specify, synthesize, and verify behaviors of embedded control systems. Our goal is to develop it as a practical tool for building real, complex, sensor-based agents. It can be seen as a formal development of Brooks' subsumption architecture [8] that enhances its reliability, modularity and scalability while avoiding the limitations of the augmented finite state machine approach, combining proactivity with reactivity.

An agent situated in an environment is modelled as three machines: the agent body, the agent controller and the environment, as shown above in Figure 2. Each is modelled separately as a dynamical system by specifying a CN with identified input and output ports. The agent is modelled as a CN consisting of a coupling of its body CN and its controller CN by identifying corresponding input and output ports. Similarly the agent CN is coupled to the environment CN to form a closed agent-environment CN, as shown above in Figure 1.

The CN model is realized as an online dataflow-like distributed programming language with a formal algebraic denotational semantics. A visual programming environment called CNJ, Constraint Nets in Java, is described in [25], along with CNML, an XML-based interchange format.

4. Agents in Design

The framework of Constraint-Based Agents consists of a modeling language such as CN and a specification language such as a real-time temporal logic, that allows the designer to specify and prove properties of the situated agent by proving them of the agent-environment CN. We have shown how to specify, design, verify and implement systems for a robot that can track other robots [29], a robot that can escape from mazes and a two-handed robot that assembles objects [31], an elevator system [38] and a car-like robot that can plan and execute paths under non-holonomic constraints [34].

Although CN can carry out traditional symbolic computation online, such as solving Constraint Satisfaction Problems and path planning, notice that much of the symbolic

reasoning and theorem-proving may be outside the agent, in the mind of the designer, for controller synthesis and verification. GOFAIR does not make this distinction, assuming that such symbolic reasoning occurs explicitly in, and only in, the mind of the agent.

The question “Will the agent do the right thing?” [31] is answered positively if we can:

1. model the coupled agent system at a suitable level of abstraction,
2. specify the required global properties of the system’s evolution, and
3. verify that the model satisfies the specification.

In CBA, the modelling language and the specification language are totally distinct since they have very different requirements. The modelling language is a generalized dynamical system language. Two versions of the specification language, Timed Linear Temporal Logic [34] and Timed \forall -automata [30], have been developed with appropriate theorem-proving and model-checking techniques for verifying systems. In [22, Chapter 12] we describe how to build a situated robot controller using CN as realized in a logic program. In this paper, we will present an overview of Timed \forall -automata.

5. Timed \forall -Automata

In general, there are two uses of automata: first, to describe computations, for example, input/output state automata, and second, to characterize a set of sequences, for example, regular grammars/languages. Examples of the first category are mostly deterministic and examples of the second category are mostly non-deterministic. However, all the original automata work is based on discrete time steps/sequences. Approaches to extending automata to continuous time have been explored in the hybrid systems community [2, 3, 11]. The timed \forall -automata model that we developed belongs to the second category, i.e., *non-deterministic finite* state automata specifying behaviors over *continuous* time. The discrete time version of \forall -automata was originally proposed as a formalism for the specification and verification of temporal properties of concurrent programs [19].

Intuitively, a \forall -automaton A consists of a set of states Q and a set of *initial* assertions e , each on every state, and a set of *transitional* assertions c , each between any two states. The set of states can be characterized into three categories: *stable* S , *recurrent* R and *bad* $B = Q - (S \cup R)$. Each assertion denotes a constraint defined on a domain of interest. Let D be a domain of interest; D can be finite, discrete, or continuous, or a cross product of a finite number of domains. Physically, D can represent, for example, velocities, distances, torques, sentences, commands or a combination of the above. A *constraint* C defined on D is a subset of D , $C \subseteq D$. Physically, a constraint represents a certain relation on a domain, such as a relation between external environment stimuli and an agent’s internal knowledge representation, or, a relation between the agent’s internal states and its actions, or, the relation between the current and next state. An element d in domain D *satisfies* constraint C , if and only if $d \in C$.

Let \mathcal{T} be a time domain, which can be continuous, for example, \mathcal{R}^+ . Let $v : \mathcal{T} \rightarrow D$ be a function of time to a given domain of interest, i.e., a *trace* of the agent-environment system. A *run* of A over v is a function $r : \mathcal{T} \rightarrow Q$ satisfying: (1) $v(0)$ satisfies the entry assertion, i.e., $v(0) \in e(r(0))$ and (2) For any $t \in \mathcal{T}$, $v(t^-)$, $v(t)$ and $v(t^+)$ satisfy the transitional assertions in $q(t^-)$, $q(t)$ and $q(t^+)$. When \mathcal{T} is discrete, the second condition is

clear, $\forall t > 0, v(t) \in c(r(pre(t)), r(t))$ where $pre(t)$ is the previous time point of t . When \mathcal{T} is continuous, the second condition becomes more complex [35].

If r is a run, let $Inf(r)$ be the set of automaton-states appearing infinitely many times in r , i.e., $Inf(r) = \{q | \forall t \exists t' \geq t, r(t') = q\}$. A run is called *accepting* if and only if

1. $Inf(r) \cup R \neq \emptyset$, that is, some of states appearing infinitely many times in r belong to R , or
2. $Inf(r) \subseteq S$, that is, all the states appearing infinitely many times in r belong to S .

A \forall -automaton A accepts a trace v , if and only if all runs of A over v are accepting for A .

One of the engineering advantages of using automata as a specification language is its graphical representation. It is useful and illuminating to represent \forall -automata by diagrams. A \forall -automaton can be depicted by a labeled directed graph, where automaton-states are depicted by circle nodes and transition relations by directed arcs. In addition, each automaton-state may have an entry arc pointing to it. Each recurrent state is depicted by a diamond and each stable state is depicted by a square, inscribed within a circle. Nodes and arcs are labeled by assertions as follows. A node or an arc that is left unlabeled is considered to be labeled with **true**. Furthermore,

1. if an automaton-state q is labeled by ψ and its entry arc is labeled by ϕ , the entry condition $e(q)$ is given by $e(q) = \psi \wedge \phi$; if there is no entry arc, $e(q) = \mathbf{false}$, and
2. if arcs from q to q' are labeled by $\phi_i, i = 1..n$, and q' is labeled by ψ , the transition condition $c(q, q')$ is given by $c(q, q') = (\phi_1 \vee \dots \vee \phi_n) \wedge \psi$; if there is no arc from q to q' , $c(q, q') = \mathbf{false}$.

We can infer that, for the \forall -automata in Figure 3, (a) specifies the behavior of reachability, i.e., eventually the system should satisfy constraint G , (b) specifies the behavior of safety, i.e. constraint G is never satisfied, (c) specifies the behavior of bounded response, i.e., whenever constraint E is satisfied, constraint F will be satisfied within bounded time.

In order to specify real-time behaviors, in addition to bounded time, timed \forall -automata are defined.

A timed \forall -automaton TA is a \forall -automaton A augmented with a global time bound and a set of local time bounds. Let $T \subseteq Q$ be a set of local time bounded states, $\tau : T \rightarrow \mathcal{R}^+$ be a local time function, and μ be a global time bound that can be any value in $\mathcal{R}^+ \cup \infty$.

For a timed \forall -automaton, in addition for a run to be accepting, it has to satisfy time constraints. Let $I \subseteq \mathcal{T}$ be a time interval and $|I|$ be the time measurement, and let $r|I$ be a segment of r over time interval I . A run satisfies time constraints if and only if:

1. *Local*: For any $q \in T$ and any time interval I , if $r|I$ is a segment of consecutive states of q , then $|I| \leq \tau(q)$;




Figure 3. Examples of timed \forall -automata.

2. *Global*: For any time interval I , if $r|I$ is a segment of consecutive states of $B \cup S$, then $\int_I \chi_B(r(t))dt \leq \mu$ where $\chi_B : Q \rightarrow \{0, 1\}$ is the characteristic function for the set B .

A timed \forall -automaton TA accepts a trace v , if and only if

1. all runs are accepting for A ; and
2. all runs satisfy the time constraints.

In a graphical representation, a local time state is denoted by a number indicating the time bound. For example, (d) in Figure 3 specifies the behavior of real-time response, i.e., whenever constraint E is satisfied, constraint F will be satisfied within 5 time units.

6. Examples of Behavioral Specification

Timed \forall -automata are a simple yet powerful model for the specification of behaviors of dynamic systems, since it integrates constraint specification with timed dynamic behavior specification.

6.1. Examples of Constraint Specification

Constraint specification alone can specify many performance metrics. As we said, constraints can specify relations between external environment stimuli and an agent's internal knowledge representation, or between an agent's internal states and its actions, or between the current and next states. Constraints can be finite, discrete or continuous, or any combination of the above. Constraints can be linear, nonlinear, equalities or inequalities. Moreover, constraints can also specify optimal conditions or optimality with extra constraints, or combinations of multiple optimal criteria and additional constraints.

Consider the following examples for specifying constraints:

1. *Inequality*: $f(x) \leq 0$ where x is a vector of variables and f is a vector of functions.
2. *Optimality*: $\min|f(x)|$ where $|x|$ is a norm for x .
3. *Negation*: $x \neq y$.
4. *Constrained Optimality*: $\min|f(x)|$ given $g(x) \leq 0$.
5. *Robustness*: Let $f(x)$ be a set of output functions with x as inputs. The robustness can be represented by its Jacobian J . There are many ways to state an optimal condition for robustness. One method is to minimize $|w|$ where w is the vector of diagonal elements of W in the singular value decomposition of $J = UWW^T$.

6.2. Examples of \forall -Automata

With automata, timed dynamic behaviors can be specified. Here is a set of examples for specifying performance using timed \forall -automata, as shown in Figure 3:

1. Let G be a constraint that the distance between the robot and its desired position is less than some constant value. Then Figure 3(a) specifies that the robot will eventually arrive at its desired position.

2. Let G be a constraint that the error of a learning algorithm is less than a desired tolerance. Then Figure 3(a) specifies that the learning will eventually converge. If we let the state of $\neg G$ in Figure 3(a) be a timed state with time bound t , it further specifies that the learning will be done within time t .
3. Let G be a constraint that the distance between the robot and obstacles is less than some constant value. Then Figure 3(b) specifies that the robot will never hit any obstacle. If G denotes that the current memory usage is over the limit, Figure 3(b) specifies that the memory usage at any time is within its limit.
4. Let E be an external stimulus and F be a response. Then Figure 3(c) specifies that there is a response after the stimulus within bounded time. Figure 3(d) specifies that such a response occurs within 5 time units.

Even though timed \forall -automata are powerful, still they are not able to represent all forms of performance metrics. For example, optimal performance over time $\min \int f(t)dt$ is not specifiable with timed \forall -automata. This form is mostly used for characterizing energy, efficiency or overall errors. Furthermore, specifications with probability behaviors are not included either. However, it is possible to add probability; for example, instead of “all runs” must be accepting and satisfy time constraints, we can say “ $x\%$ of the runs” must be accepting and satisfy the timing constraints.

7. Verification and Synthesis of Constraint-Satisfying Controllers

For most dynamic systems, stability or convergence is the most important property that needs to be verified. For example, we can verify that equation $dx/dt = 0$ satisfies the property of the \forall -automaton in Figure 3(a) with G as $|x| \leq \epsilon$ for any positive number ϵ . The most commonly used method for the verification of such properties is the use of Lyapunov functions. We developed a formal method based on model-checking, that generalizes Lyapunov functions [28, 35]. This method is automatic if the domain of interest is finite discrete and time is discrete [28].

The basic verification strategy is, first, to find a set of invariants, each associated with an automaton-state in the timed \forall -automaton. Then, find a set of Lyapunov functions, which are non-increasing in stable states and decreasing in bad states. Finally, find a set of local and global timing functions, where local timing functions are decreasing in timed states and global timing functions, like Lyapunov functions, are non-increasing in stable states and decreasing in bad states, in addition to having bounded values.

Many agents can be designed as online constraint-satisfying devices [30, 33, 34]. A robot in this restricted scheme can be verified more easily. Moreover, given a constraint-based specification and a model of the body and the environment, automatic synthesis of a correct constraint-satisfying controller sometimes becomes feasible, as shown for a simple goal-scoring robot in [34].

As a simple example, in Figure 2 suppose the CONTROLLER is a thermostat turning on or off a furnace, the BODY, that is heating the ENVIRONMENT. The goal of the system is to make the temperature of the ENVIRONMENT, $T_E(t)$, equal to a desired temperature, $T_D(t)$. In other words the CONTROLLER of the AGENT is trying to solve the constraint $T_E = T_D$. One version of CONTROLLER correctness is established if we

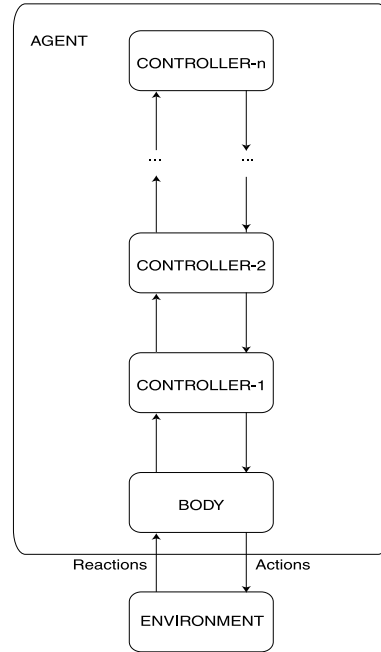


Figure 4. A hierarchical agent controller.

can prove that the (thermal) dynamics of the coupled AGENT-ENVIRONMENT system satisfy the temporal logic formula $\diamond\Box|T_E - T_D| < \epsilon$ where \diamond can be read as eventually and \Box can be read as ‘always’. In other words, the system will, no matter how disturbed, eventually enter, and remain within, an ϵ -neighborhood of the solution manifold of the constraint. As a timed \forall -automaton, this requirement is represented by Figure 3(a). A less restrictive form of correctness corresponds to the specification $\Box\diamond|T_E - T_D| < \epsilon$ which is to say that the system will always return, asymptotically, to the constraint solution manifold if it should happen to leave it. As a timed \forall -automaton, this specification is represented by Figure 3(c). Assuming the on/off thermostat’s dynamics be $dT_E/dt = k$ if on and $dT_E/dt = -m$ if off, a simple controller can be *on* if $T_D - T_E > C$ and *off* if $T_E - T_D > C$ where $k, m, C > 0$. In order to satisfy the “strong” condition shown in Figure 3(a), $C < \epsilon$.

In general, a constraint is simply a relation on the phase space of the agent system, which is the product of the controller, body and environment spaces. A controller is defined to be *constraint-satisfying* if it, repeatedly, eventually drives the system into an ϵ -neighborhood of the constraint using a constraint satisfaction method such as gradient descent or a symbolic technique.

A constraint-satisfying controller may be *hierarchical* with several layers of controller above the body, as shown in Figure 4. In this case, each layer must satisfy the constraints, defined on its state variables, appropriate to the layer, as, typically, set by the layer above. The layers below each layer present to that layer as a virtual agent body in a suitably abstract state space [34, 36]. The lower layers are, typically, reactive and synchronous (or in continuous time) on continuous state spaces; the upper layers are more deliberative and asynchronous (or event-triggered) in symbolic, discrete spaces.




Figure 5. A layer in a constraint-based controller.

A typical layer in a hierarchical controller is shown in Figure 5.

Each layer has two external inputs: the trace of constraint requests coming from above *ConstraintsIn* (*CI*) and the reports coming from below *ReportsIn* (*RI*). Its two outputs are its reports to the level above *ReportsOut* (*RO*) and its constraint requests to the level below *ConstraintsOut* (*CO*). These traces arise from *causal transductions* of the external inputs:

$$CO = C_t(CI, RI) \tag{1}$$

$$RO = R_t(CI, RI) \tag{2}$$

If the constraint-solver can be represented as a state-based solver then the layer may be represented as shown in Figure 6.

Here, for simplicity, a discrete-time state-based layer is shown. It produces an extra output *StateOut* (*SO*) that is consumed as an extra input *StateIn* (*SI*) after a unit delay (Δ). In this case the behavior of the layer may be represented by computing the values of each of the three outputs as a *transliteration* (function) of the current values of the three inputs:

$$CO(t) = C_f(CI(t), RI(t), SI(t)) \tag{3}$$

$$RO(t) = R_f(CI(t), RI(t), SI(t)) \tag{4}$$

$$SO(t) = S_f(CI(t), RI(t), SI(t)) \tag{5}$$

$$SI(t + 1) = SO(t) \tag{6}$$

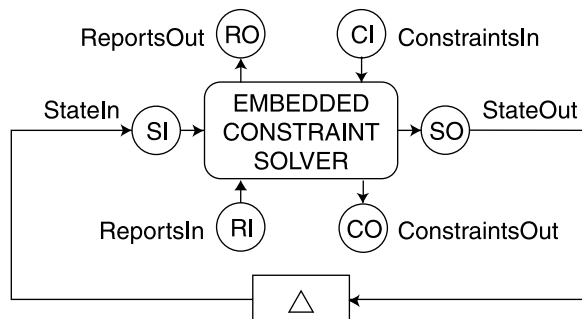


Figure 6. A layer with state in a constraint-based controller.

8. Robot Soccer Players

The CBA framework has also been motivated, developed and tested by application to the experimental challenge of designing, building and verifying controllers with perceptual systems for robot soccer players with both off-board and on-board vision systems.

In the Dynamo (Dynamics and Mobile Robots) project in our laboratory, we have experimented, since 1991, with multiple mobile robots under visual control. The Dynamite testbed consists of a fleet of radio-controlled vehicles that receive commands from a remote computer. Using our custom hardware and a distributed MIMD environment, vision programs are able to monitor the position and orientation of each robot at 60 Hz; planning and control programs generate and send motor commands at the same rate. This approach allows umbilical-free behavior and very rapid, lightweight fully autonomous robots. Using this testbed we have demonstrated various robot tasks [5], including playing soccer [24] using a 2-layer deliberative/reactive controller architecture.

One of the Dynamo robots, Spinoza, is a self-contained robot consisting of an RWI base with an RGB camera on a pan-tilt platform mounted as its head and a trinocular stereo camera in its base. As an illustration of these ideas, consider the task for Spinoza of repeatedly finding, tracking, chasing and kicking a soccer ball, using the pan-tilt camera. After locating the moving ball Spinoza is required to track it, move to within striking distance of the ball and strike it. The available motor commands control the orientation of the base, the forward movement of the base, and the pan and tilt angles of the camera. The parameters can be controlled in various relative/absolute position modes or rate mode. The available rate of pan substantially exceeds the rate of base rotation. A hierarchical constraint-based active-vision controller, using prioritized constraints and constraint arbiters, can be specified for Spinoza that will, repeatedly, achieve and maintain (or re-achieve) the desired goal subject to safety conditions such as staying inside the soccer field, avoiding obstacles and not accelerating too quickly. If the dynamics of Spinoza and the ball are adequately modelled by the designer then this constraint-based vision system will be guaranteed to achieve its specification.

Yu Zhang and Mackworth have extended these ideas to build 3-layer constraint-satisfying controllers for a complete soccer team [39]. The controllers for our softbot soccer team, UBC Dynamo98, are modelled in CN and implemented in Java, using the Java Beans architecture [36, 41]. They control the soccer players' bodies in the Soccer Server developed by Noda Itsuki for RoboCup. These experiments provide evidence that the constraint-based CN approach is a clean and practical design framework for perceptual robots.

9. Conclusion

The Constraint-Based Agent approach is a framework for the specification, design, analysis, implementation and validation of artificial and natural agent systems. It requires a new model of online and embedded computation for Constraint Programming such as Constraint Nets, and a general behavior specification language like timed \forall -automata.

In this overview paper, we have shown how to use formal methods to specify the dynamic system using Constraint Nets and the desired behaviors with timed \forall -automata.

The advantage of formal methods over other methods lies in their precision and generality. Constraint Nets, with its unified semantic model for hybrid systems, provides a clear component structure for complex systems. Timed \forall -automata, with its graphical depiction and constraint specification, is a simple yet powerful formalism for specifying many properties of dynamic systems.

Acknowledgments

We are grateful to Rod Barman, Cullen Jennings, Stewart Kingdon, Jim Little, Valerie McRae, Don Murray, Dinesh Pai, David Poole, Michael Sahota, Fengguang Song, Robert St-Aubin, Vlad Tucakov and Yu Zhang for help with this. This work is supported, in part, by the Natural Sciences and Engineering Research Council of Canada and the Institute for Robotics and Intelligent Systems Network of Centres of Excellence. Alan Mackworth holds a Canada Research Chair in Artificial Intelligence.

References

1. Albus, J. S. (1991). Outline for a theory of intelligence. In *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 21, pages 473–509. IEEE.
2. Alur, R., Courcoubetis, C., Henzinger, T. A., & Ho, P. (1993). Hybrid automata: an algorithmic approach to the specification and verification of hybrid systems. In Grossman, R. L., Nerode, A., Ravn, A. P., & Rischel, H., eds., *Hybrid Systems*, No. LNCS 736, pages 209–229. Springer-Verlag.
3. Alur, R., & Dill, D. (1990). Automata for modeling real-time systems. In Peterson, M. S., ed., *ICALP90: Automata, Languages and Programming*, No. LNCS443, pages 322–335. Springer-Verlag.
4. Arkin, R. C. (1998). *Behavior-Based Robotics*. Cambridge, MA: The MIT Press.
5. Barman, R. A., Kingdon, S. J., Little, J. J., Mackworth, A. K., Pai, D. K., Sahota, M., Wilkinson, H., & Zhang, Y. (1993). Dynamo: real-time experiments with multiple mobile robots. In *Intelligent Vehicles Symposium*, pages 261–266. Tokyo.
6. Bemporad, A., Morari, M., Dua, V., & Pistikopoulos, E. N. (2000). The explicit solution of model predictive control via multiparametric quadratic programming. In *Proc. of American Control Conference*, pages 872–876. Chicago, Illinois.
7. Bowen, J. P. (2001). Formal methods web site, <http://www.afm.sbu.ac.uk/>.
8. Brooks, R. A. (1991). Intelligence without reason. In *IJCAI-91*, pages 569–595. Sydney, Australia.
9. Excalibur: Adaptive Constraint-Based Agents in Artificial Environments, <http://www.ai-center.com/projects/excalibur/>.
10. Freuder, E. C. (1997). Constraints and agents. AAAI Workshop, <http://www.aaai.org/Press/Reports/Workshops/ws-97-05.html>.
11. Henzinger, T. A., Manna, Z., & Pnueli, A. (1991). Timed transition systems. In deBakker, J. W., Huizing, C., dePoeper, W. P., & Rozenberg, G., eds., *Real-Time: Theory in Practice*, No. LNCS 600, pages 226–251. Springer-Verlag.
12. Lavignion, J., & Shoham, Y. (1990). Temporal automata. Technical Report STAN-CS-90-1325. Stanford University, Stanford, CA.
13. Lavretsky, E. (2000). Greedy optimal control. In *Proc. of American Control Conference*, pages 3888–3892. Chicago, Illinois.
14. Mackworth, A. K. (1993). On seeing robots. In Basu, A., & Li, X., eds., *Computer Vision: Systems, Theory, and Applications*, pages 1–13. Singapore: World Scientific Press.
15. Mackworth, A. K. (1996). Quick and clean: constraint-based vision for situated robots. In *IEEE Int'l. Conf. on Image Processing*, pages 789–792. Lausanne, Switzerland.
16. Mackworth, A. K. (1997). Constraint-based design of embedded intelligent systems. *Constraints*, 2(1): 83–86.

17. Mackworth, A. K. (1999). The dynamics of intelligence: constraint-satisfying hybrid systems for perceptual agents. In *Hybrid Systems and AI: Modeling, Analysis and Control of Discrete and Continuous Systems*, pages 210–214. Stanford, CA.
18. Mackworth, A. K. (2000). Constraint-based agents: the ABCs of CBAs. In *Proc. 6th Int. Conf. on Principles and Practice of Constraint Programming – CP2000*, Springer LNCS 1894, pages 1–10. Singapore.
19. Manna, Z., & Pnueli, A. (1987). Specification and verification of concurrent programs by \forall -automata. In *Proc. 14th Annual ACM Symposium on Principles of Programming Languages*, pages 1–12.
20. Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18(1): 87–127.
21. The Ptolemy Project, <http://ptolemy.eecs.berkeley.edu/>.
22. Poole, D. L., Mackworth, A. K., & Goebel, R. G. (1998). *Computational Intelligence: A Logical Approach*. New York: Oxford University Press.
23. Rosenschein, S. J., & Kaelbling, L. P. (1986). The synthesis of machines with provable epistemic properties. In Halpern, Joseph, ed., *Proc. Conf. on Theoretical Aspects of Reasoning about Knowledge*, pages 83–98. Los Altos, CA: Morgan Kaufmann.
24. Sahota, M., & Mackworth, A. K. (1994). Can situated robots play soccer? In *Proc. Artificial Intelligence 94*, pages 249–254. Banff, AB.
25. Song, F., & Mackworth, A. K. (2002). CNJ: a visual programming environment for constraint nets. In *Proc. AI, Simulation and Planning in High Autonomy Systems - AIS-2002*, page 6. Lisbon, to appear.
26. Turing, A. M. (1963). Computing machinery and intelligence. In Feigenbaum, E., & Feldman, J., eds., *Computers and Thought*, pages 11–35. New York, NY: McGraw-Hill.
27. Weiss, G., ed., (1999). *Multiagent Systems*. The MIT Press.
28. Zhang, Y. (1994). A foundation for the design and analysis of robotic systems and behaviors. Ph.D. thesis, University of British Columbia, Vancouver, British Columbia.
29. Zhang, Y., & Mackworth, A. K. (1992). Modeling behavioral dynamics in discrete robotic systems with logical concurrent objects. In Tzafestas, S. G., & Gentina, J. C., eds., *Robotics and Flexible Manufacturing Systems*, pages 187–196. Elsevier Science Publishers B.V.
30. Zhang, Y., & Mackworth, A. K. (1994a). Specification and verification of constraint-based dynamic systems. In Borning, A., ed., *Principles and Practice of Constraint Programming*, No. 874 in Lecture Notes in Computer Science, pages 229–242. Springer-Verlag.
31. Zhang, Y., & Mackworth, A. K. (1994b). Will the robot do the right thing? In *Proc. Artificial Intelligence 94*, pages 255–262. Banff, AB.
32. Zhang, Y., & Mackworth, A. K. (1995b). Constraint nets: a semantic model for hybrid dynamic systems. *Theoretical Computer Science*, 138: 211–239.
33. Zhang, Y., & Mackworth, A. K. (1995c). Constraint programming in constraint nets. In Saraswat, V., & Van Hentenryck, P., eds., *Principles and Practice of Constraint Programming*, Chapt. 3, pages 49–68. Cambridge, MA: The MIT Press.
34. Zhang, Y., & Mackworth, A. K. (1995d). Synthesis of hybrid constraint-based controllers. In Antsaklis, P., Kohn, W., Nerode, A., & Sastry, S., eds., *Hybrid Systems II*, Lecture Notes in Computer Science 999, pages 552–567. Springer Verlag.
35. Zhang, Y., & Mackworth, A. K. (1996). Specification and verification of hybrid dynamic systems with timed \forall -automata. In Alur, R., Henzinger, T. A., & Sontag, E. D., eds., *Hybrid Systems III, Verification and Control*, pages 587–603. Springer.
36. Zhang, Y., & Mackworth, A. K. (1998a). A constraint-based controller for soccer-playing robots. In *Proceedings of IROS '98*, pages 1290–1295. Victoria, BC, Canada.
37. Zhang, Y., & Mackworth, A. K. (1998b). Using reactive deliberation for real-time control of soccer-playing robots. In Kitano, H., ed., *RoboCup-97: Robot Soccer World Cup 1*, pages 508–512.
38. Zhang, Y., & Mackworth, A. K. (1999a). Modelling and analysis of hybrid systems: an elevator case study. In Levesque, H., & Pirri, F., eds., *Logical Foundations for Cognitive Agents*, pages 370–396. Berlin: Springer.
39. Zhang, Y., & Mackworth, A. K. (1999b). A multi-level constraint-based controller for the Dynamo98 robot soccer team. In Asada, Minoru, & Kitano, Hiroaki, eds., *RoboCup-98: Robot Soccer World Cup II*, pages 402–409. Springer.
40. Zhang, Y., & Mackworth, A. K. (2000). Formal specification of performance metrics for intelligent systems. In *Proc. Workshop on Performance Metrics for Intelligent Systems, Aug. 14-16, 2000*. National Institute of Standards and Technology, Washington, DC, Section Sixty.
41. Zhang, Y., & Mackworth, A. K. (2002). A constraint-based robotic soccer team. *Constraints*, 7(1): 7–28.