

INTELLIGENT CONTROL

by
Panos J. Antsaklis

Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN 46556 USA
antsaklis.1@nd.edu
<http://www.nd.edu/~pantsakl>

Written for the
Encyclopedia of Electrical and Electronics Engineering
John Wiley & Sons, Inc.
1997

Intelligent control describes the discipline where control methods are developed that attempt to emulate important characteristics of human intelligence. These characteristics include adaptation and learning, planning under large uncertainty and coping with large amounts of data. Today, the area of intelligent control tends to encompass everything that is not characterized as conventional control; it has, however, shifting boundaries and what is called "intelligent control" today, will probably be called "control" tomorrow. The main difficulty in specifying exactly what is meant by the term Intelligent control stems from the fact that there is no agreed upon definition of human intelligence and intelligent behavior and the centuries old debate of what constitutes intelligence is still continuing, nowadays among educators, psychologists, computer scientists and engineers. Apparently the term Intelligent control was coined in the 70's by K.S. Fu. Reference 1 is the main source of the several descriptions of intelligent control and its attributes discussed in this article.

There are a number of areas related to the area of Intelligent control. Intelligent control is interdisciplinary as it combines and extends theories and methods from areas such as control, computer science and operations research. It uses theories from mathematics and seeks inspiration and ideas from biological systems. Intelligent control methodologies are being applied to robotics and automation, communications, manufacturing, traffic control, to mention but a few application areas. Neural networks, fuzzy control, genetic algorithms, planning systems, expert systems, hybrid systems are all areas where related work is taking place. The areas of computer science and in particular artificial intelligence provide knowledge representation ideas, methodologies and tools such as semantic networks, frames, reasoning techniques and computer languages such as prolog. Concepts and algorithms developed in the areas of adaptive control and machine learning help intelligent controllers to adapt and learn. Advances in sensors, actuators, computation technology and communication networks help provide the necessary for implementation Intelligent control hardware.

In the following, fundamental ideas of Intelligent control are emphasized, rather than particular methodologies such as fuzzy control; note that several related areas are described at length elsewhere in this encyclopedia. Fundamental ideas and characteristics of intelligent systems are introduced in the section on Foundations of Intelligent Control, and a historical perspective is brought in in the section on Intelligent Learning Control where the role of machine learning is discussed. The quest for machines that exhibit higher autonomy has been the driving force in the development of control systems over the centuries and this is discussed in the section on Intelligent Control for High Autonomy Systems. Hybrid Systems that contain both continuous and digital components are also briefly discussed, as they are central in Intelligent control.

FOUNDATIONS OF INTELLIGENT CONTROL

The term "intelligent control" has come to mean, particularly to those outside the control area, some form of control using fuzzy and/or neural network methodologies. Intelligent control, however does not restrict itself only to those methodologies. In fact, according to some definitions of intelligent control not all neural/fuzzy controllers would be considered intelligent. The fact is that there are problems of control today, that cannot be formulated and studied in the conventional differential/difference equation mathematical framework using "conventional (or traditional) control" methodologies; these methodologies were developed in the past decades to control dynamical systems. To address these problems in a systematic way, a number of methods have been developed in recent years that are collectively known as "intelligent control" methodologies. There are significant differences between conventional and intelligent control and some of them are described below. It is worth remembering at this point that intelligent control uses conventional control methods to solve "lower level" control problems and that conventional control is included in the area of intelligent control. In summary, intelligent control attempts to build upon and enhance the conventional control methodologies to solve new challenging control problems.

Conventional and Intelligent Control

The word control in "intelligent control" has different, more general meaning than the word control in "conventional control". First, the processes of interest are more general and may be described, for example by either discrete event system models or differential/difference equation models or both. This has led to the development of theories for hybrid control systems, which study the control of continuous-state dynamic processes by discrete-state controllers. In addition to the more general processes considered in intelligent control, the control objectives can also be more general. For example, "replace part A in satellite" can be the general task for the controller of a space robot arm; this is then decomposed into a number of subtasks, several of which may include for instance "follow a particular trajectory", which may be a problem that can be solved by conventional control methodologies. To attain such control goals for complex systems over a period of time, the controller has to cope with significant uncertainty that fixed feedback robust controllers or adaptive controllers cannot deal with. Since the goals are to be attained under large uncertainty, fault diagnosis and

control reconfiguration, adaptation and learning are important considerations in intelligent controllers. It is also clear that task planning is an important area in intelligent control design. So the control problem in intelligent control is an enhanced version of the problem in conventional control. It is much more ambitious and general. It is not surprising then that these increased control demands require methods that are not typically used in conventional control. The area of intelligent control is in fact interdisciplinary, and it attempts to combine and extend theories and methods from areas such as control, computer science and operations research to attain demanding control goals in complex systems.

Note that the theories and methodologies from the areas of operations research and computer science cannot, in general be used directly to solve control problems, as they were developed to address different needs; they must first be enhanced and new methodologies need to be developed in combination with conventional control methodologies, before controllers for very complex dynamical systems can be designed in systematic ways. Also traditional control concepts such as stability may have to be redefined when, for example, the process to be controlled is described by discrete event system models; and this issue is being addressed in the literature. Concepts such as reachability and deadlock developed in operations research and computer science are useful in intelligent control, when studying planning systems. Rigorous mathematical frameworks, based for example on predicate calculus are being used to study such questions. However, in order to address control issues, these mathematical frameworks may not be convenient and they must be enhanced or new ones must be developed to appropriately address these problems. This is not surprising as the techniques from computer science and operations research are primarily analysis tools developed for non real-time systems, while in control, synthesis techniques to design real-time feedback control laws for dynamic systems are mainly of interest. In view of this discussion, it should be clear that intelligent control research, which is mainly driven by applications has a very important and challenging theoretical component. Significant theoretical strides must be made to address the open questions. The problems are nontrivial, but the pay-off is very high indeed.

As it was mentioned above, the word control in intelligent control has a more general meaning than in conventional control; in fact it is closer to the way the term control is used in every day language. Because intelligent control addresses more general control problems that also include the problems addressed by conventional control, it is rather difficult to come up with meaningful bench mark examples. Intelligent control can address control problems that cannot be formulated in the language of conventional control. To illustrate, in a rolling steel mill, for example, while conventional controllers may include the speed (rpm) regulators of the steel rollers, in the intelligent control framework one may include in addition, fault diagnosis and alarm systems; and perhaps the problem of deciding on the set points of the regulators, that are based on the sequence of orders processed, selected based on economic decisions, maintenance schedules, availability of machines etc.. All these factors have to be considered as they play a role in controlling the whole production process which is really the overall goal.

Another difference between intelligent and conventional control is in the separation between controller and the system to be controlled. In conventional control the system to be controlled, called the plant, typically is separate and distinct from the controller. The controller is

designed by the control designer, while the plant is in general given and cannot be changed; note that recent attempts to coordinate system design and control have been reported in areas such as space structures and chemical processes, as many times certain design changes lead to systems that are much easier to control. In intelligent control problems, which are most often complex and challenging, there may not be a clear separation of the plant and the controller; the control laws may be imbedded and be part of the system to be controlled. This opens new opportunities and challenges as it may be possible to affect the design of processes in a more systematic way.

Areas relevant to intelligent control, in addition to conventional control include hybrid systems, planning and knowledge based systems, machine learning, search algorithms, fault diagnosis and control reconfiguration, predicate logic, automata, Petri nets, neural nets and fuzzy logic. In addition, in order to control complex systems, one has to deal effectively with the computational complexity issue; this has been in the periphery of the interests of the researchers in conventional control, but it is clear that computational complexity is a central issue whenever one attempts to control complex systems.

Intelligence And Intelligent Control

It is appropriate at this point to briefly comment on the meaning of the word intelligent in "intelligent control". Note that the precise definition of "intelligence" has been eluding mankind for thousands of years. More recently, this issue has been addressed by disciplines such as psychology, philosophy, biology and of course by artificial intelligence (AI); note that AI is defined to be the study of mental faculties through the use of computational models. No consensus has emerged as yet of what constitutes intelligence. The controversy surrounding the widely used IQ tests, also points to the fact that we are well away from having understood these issues. In this article we introduce and discuss several characterizations of intelligent systems that appear to be useful when attempting to address complex control problems.

Intelligent controllers can be seen as machines which emulate human mental faculties such as adaptation and learning, planning under large uncertainty, coping with large amounts of data etc. in order to effectively control complex processes; and this is the justification for the use of the term intelligent in intelligent control, since these mental faculties are considered to be important attributes of human intelligence. An alternative term, that is further discussed below in this article, is "autonomous (intelligent) control"; it emphasizes the fact that an intelligent controller typically aims to attain higher degrees of autonomy in accomplishing and even setting control goals, rather than stressing the (intelligent) methodology that achieves those goals. We should keep in mind that "intelligent control" is only a name that appears to be useful today. In the same way the "modern control" of the 60's has now become "conventional (or traditional) control", as it has become part of the mainstream, what is called intelligent control today may be called just "control" in the not so distant future. What is more important than the terminology used are the concepts and the methodology, and whether or not the control area and intelligent control will be able to meet the ever increasing control needs of our technological society.

Defining Intelligent Control Systems

Intelligent systems can be characterized in a number of ways and along a number of dimensions. There are certain attributes of intelligent systems, that are of particular interest in the control of systems; see reference 1. We begin with a general characterization of intelligent systems: An intelligent system has the ability to act appropriately in an uncertain environment, where an appropriate action is that which increases the probability of success, and success is the achievement of behavioral subgoals that support the system's ultimate goal. In order for a man-made intelligent system to act appropriately, it may emulate functions of living creatures and ultimately human mental faculties.

An intelligent system can be characterized along a number of dimensions. There are degrees or levels of intelligence that can be measured along the various dimensions of intelligence. At a minimum, intelligence requires the ability to sense the environment, to make decisions and to control action. Higher levels of intelligence may include the ability to recognize objects and events, to represent knowledge in a world model, and to reason about and plan for the future. In advanced forms, intelligence provides the capacity to perceive and understand, to choose wisely, and to act successfully under a large variety of circumstances so as to survive and prosper in a complex and often hostile environment. Intelligence can be observed to grow and evolve, both through growth in computational power and through accumulation of knowledge of how to sense, decide and act in a complex and changing world.

The above characterization of an intelligent system is rather general. According to this, a great number of systems can be considered intelligent. In fact, according to this definition even a thermostat may be considered to be an intelligent system, although of low level of intelligence. It is common however to call a system intelligent when in fact it has a rather high level of intelligence. There exist a number of alternative but related definitions of intelligent systems which emphasize systems with high degrees of intelligence. For example, the following definition emphasizes the fact that the system in question processes information, and it focuses on man-made systems and intelligent machines: Machine intelligence is the process of analyzing, organizing and converting data into knowledge; where (machine) knowledge is defined to be the structured information acquired and applied to remove ignorance or uncertainty about a specific task pertaining to the intelligent machine. This definition relates to the principle of increasing precision with decreasing intelligence of Saridis.

Next, an intelligent system can be characterized by its ability to dynamically assign subgoals and control actions in an internal or autonomous fashion: Many adaptive or learning control systems can be thought of as designing a control law to meet well-defined control objectives. This activity represents the system's attempt to organize or order its "knowledge" of its own dynamical behavior, so to meet a control objective. The organization of knowledge can be seen as one important attribute of intelligence. If this organization is done autonomously by the system, then intelligence becomes a property of the system, rather than of the system's designer. This implies that systems which autonomously (self)-organize controllers with respect to an internally realized organizational principle are intelligent control systems.

A procedural characterization of intelligent systems is given next: Intelligence is a property of

the system which emerges when the procedures of focusing attention, combinatorial search, and generalization are applied to the input information in order to produce the output. One can easily deduce that once a string of the above procedures is defined, the other levels of resolution of the structure of intelligence are growing as a result of the recursion. Having only one level structure leads to a rudimentary intelligence that is implicit in the thermostat, or to a variable-structure sliding mode controller.

Control and Intelligent Systems

The concepts of intelligence and control are closely related and the term "Intelligent control" has a unique and distinguishable meaning. An intelligent system must define and use goals. Control is then required to move the system to these goals and to define such goals. Consequently, any intelligent system will be a control system. Conversely, intelligence is necessary to provide desirable functioning of systems under changing conditions, and it is necessary to achieve a high degree of autonomous behavior in a control system. Since control is an essential part of any intelligent system, the term "intelligent control systems" is sometimes used in engineering literature instead of "intelligent systems" or "intelligent machines". The term "intelligent control system" simply stresses the control aspect of the intelligent system.

Below, one more alternative characterization of intelligent (control) systems is included. According to this view, a control system consists of data structures or objects (the plant models and the control goals) and processing units or methods (the control laws): An intelligent control system is designed so that it can autonomously achieve a high level goal, while its components, control goals, plant models and control laws are not completely defined, either because they were not known at the design time or because they changed unexpectedly.

Characteristics or Dimensions of Intelligent Systems.

There are several essential properties present in different degrees in intelligent systems. One can perceive them as intelligent system characteristics or dimensions along which different degrees or levels of intelligence can be measured. Below we discuss three such characteristics that appear to be rather fundamental in intelligent control systems.

Adaptation and Learning: The ability to adapt to changing conditions is necessary in an intelligent system. Although adaptation does not necessarily require the ability to learn, for systems to be able to adapt to a wide variety of unexpected changes learning is essential. So the ability to learn is an important characteristic of (highly) intelligent systems.

Autonomy and Intelligence: Autonomy in setting and achieving goals is an important characteristic of intelligent control systems. When a system has the ability to act appropriately in an uncertain environment for extended periods of time without external intervention it is considered to be highly autonomous. There are degrees of autonomy; an adaptive control system can be considered as a system of higher autonomy than a control system with fixed controllers, as it can cope with greater uncertainty than a fixed feedback controller. Although

for low autonomy no intelligence (or "low" intelligence) is necessary, for high degrees of autonomy, intelligence in the system (or "high" degrees of intelligence) is essential.

Structures and Hierarchies: In order to cope with complexity, an intelligent system must have an appropriate functional architecture or structure for efficient analysis and evaluation of control strategies. This structure should provide a mechanism to build levels of abstraction (resolution, granularity) or at least some form of partial ordering so to reduce complexity. An approach to study intelligent machines involving entropy (of Saridis) emphasizes such efficient computational structures. Hierarchies (that may be approximate, localized or combined in heterarchies) that are able to adapt, may serve as primary vehicles for such structures to cope with complexity. The term "hierarchies" refers to functional hierarchies, or hierarchies of range and resolution along spatial or temporal dimensions, and it does not necessarily imply hierarchical hardware. Some of these structures may be hardwired in part. To cope with changing circumstances the ability to learn is essential so these structures can adapt to significant, unanticipated changes.

In view of the above, a working characterization of intelligent systems (or of (highly) intelligent (control) systems or machines) that captures the essential characteristics present in any such system is: An intelligent system must be highly adaptable to significant unanticipated changes, and so learning is essential. It must exhibit high degree of autonomy in dealing with changes. It must be able to deal with significant complexity, and this leads to certain types of functional architectures such as hierarchies.

Some Examples

Man-made systems that solve complex problems and incorporate some of the above essential characteristics of intelligent control systems do exist today. Here are some examples from reference 1: A hierarchically intelligent control System was designed and built at the NASA CIRSSSE/RPI (Renssellar Polytechnic Institute) laboratories, to do truss construction remotely in deep space for the NASA space station "Freedom". This Intelligent control system had a functional hierarchy that consisted of three levels: the lowest was the Execution level, the highest was the Organization level and the middle was the Coordination level (see Figure 1 and the section on Intelligent Autonomous Control later in this article). The innovation of the project was that a system was directing the flow of data at the execution level located at the site, while only commands were communicated to and from the coordination level on Earth. The following are examples of intelligent control systems in NIST's (National Institute for Standards and Technology) RCS (Real-time Control System) implementations: Robot vision-based object pursuit; robot deburring; composites fabrication; automated manufacturing research facility; robot machine loading/unloading for a milling workstation; multiple autonomous undersea vehicles; NASA space station telerobotics; army field material handling robot; DARPA submarine automation; coal mine automation; and army unmanned land vehicles. Other examples of existing intelligent control systems include mobile robots that exhibit some autonomy at Oak Ridge National Laboratory, and at the Massachusetts and Georgia Institutes of Technology.

For additional information and insight into the foundations of Intelligent control, the

interested reader may refer to references 1-8.

INTELLIGENT LEARNING CONTROL

The term Intelligent control was coined in the 70's. Earlier terms used included Learning Control and Self-organizing Control. A brief description of some of the early developments in the area that is known today as Intelligent control is given.

As discussed previously, learning is an important dimension or attribute of Intelligent control. Highly autonomous behavior is a very desirable characteristic of advanced control systems, so they perform well under changing conditions in the plant and the environment (even in the control goals), without external intervention; note that intelligent autonomous control is discussed at length below in this article. This requires the ability to adapt to changes affecting, in a significant manner, the operating region of the system. Adaptive behavior of this type typically is not offered by conventional control systems. Additional decision making abilities should be added to meet the increased control requirements. The controller's capacity to learn from past experience is an integral part of such highly autonomous controllers. The goal of introducing learning methods in control is to broaden the region of operability of conventional control systems. Therefore the ability to learn is one of the fundamental attributes of autonomous intelligent behavior; see references 1, 2.

The ability of man-made systems to learn from experience and, based on that experience, improve their performance is the focus of machine learning. Learning can be seen as the process whereby a system can alter its actions to perform a task more effectively due to increases in knowledge related to the task. The actions that a system may take depend on the nature of the system. For example, a control system may change the type of controller used, or vary the parameters of the controller, after learning that the current controller does not perform satisfactorily within a changing environment. Similarly, a robot may need to change its visual representation of the surroundings after learning of new obstacles in the environment. The type of action taken by the machine is dependent upon the nature of the system and the type of learning system implemented. The ability to learn entails such issues as knowledge acquisition, knowledge representation, and some level of inference capability. Learning, considered fundamental to intelligent behavior, and in particular the computer modeling of learning processes has been the subject of research in the field of machine learning since the 1960's; see references 9,10.

Learning Control

The problem of learning in automatic control systems has been studied in the past, especially in the late 60's, and it has been the topic of numerous papers and books; see for example references 11-15. References 11, 13, 15 provide surveys on the early learning techniques. All of these approaches involve a process of classification, in which all or part of the prior information required is unknown or incompletely known. The elements or patterns that are presented to the control system are collected into groups that correspond to different pattern classes or regions; see reference 15. Thus learning was viewed as the estimation or successive

approximation of the unknown quantities of a function; see reference 11. The approaches developed for such learning problems can be separated into two categories: deterministic and stochastic. Where can learning be used in the control of systems? As it was already mentioned, learning plays an essential role in the autonomous control of systems. There are many areas in control where learning can be used to advantage and these needs can be briefly classified as follows: 1. Learning about the plant; that is learning how to incorporate changes and then how to derive new plant models. 2. Learning about the environment ; this can be done using methods ranging from passive observation to active experimentation. 3. Learning about the controller; for example, learning how to adjust certain controller parameter to enhance performance. 4. Learning new design goals and constraints. What is the relation between adaptive control and learning control? Learning is achieved, in a certain sense, when an adaptive control algorithm is used to adapt the controller parameters so that for example stability is maintained. In this case the system learns and the knowledge acquired is the new values for the parameters. Note however, that if later the same changes occur again and the system is described by exactly the same parameters identified earlier, the adaptive control algorithm still needs to recalculate the controller and perhaps the plant parameters since nothing was kept in memory. So, in that sense the system has not learned. It has certainly learned what to do when certain type of changes take place. In particular, it has been told exactly what to do, that is it was given the adaptive algorithm, and this is knowledge by rote learning. The knowledge represented by the new values of the controller and the plant parameters and the circumstances under which these values are appropriate, are not retained. So a useful rule of thumb is that a controller to be a learning controller, memory is required where past knowledge is stored in such a way so it can be used to benefit when a similar situation arises.

Some Historical Notes

Regarding terminology it is perhaps beneficial at this point to bring in a bit of history: In the 60's, adaptive control and learning received a lot of attention in the control literature. It was not always clear however what it was meant by those terms. The comment by Y.Tsytkin, in reference 14 describes quite clearly the atmosphere of the period: "It is difficult to find more fashionable and attractive terms in the modern theory of automatic control than the terms of adaptation and learning. At the same time, it is not simple to find any other concepts which are less complex and more vague." Adaptation, learning, self-organizing systems and control were competing terms for similar research areas, and K.S. Fu says characteristically in reference 11: "The use of the word 'adaptive' has been intentionally avoided here... adaptive and learning are behavior-descriptive terms, but feedback and self-organizing are structure, or system configuration-descriptive terms. Nevertheless the terminology war is still going on...It is certainly not the purpose of this paper to get involved with such a war." The term pattern recognition was also appearing together with adaptive, learning and self-organizing systems in the control literature of that era. It is obvious that there was no agreement as to the meaning of these terms and their relation. Pattern recognition is today a research discipline in its own right, developing and using an array of methods ranging from conventional algorithms to artificial intelligence methods implemented via symbolic processing. The term self-organizing system is not being used as much today in the control literature. Adaptive control has gained renewed popularity in the past decades mainly emphasizing studies in the convergence of

adaptive algorithms and in the stability of adaptive systems; the systems considered are primarily systems described by differential (or difference) equations where the coefficients are (partially) unknown. In an attempt to enhance the applicability of adaptive control methods, learning control has been recently reintroduced in the control literature; see for example reference 7 for learning methods in control with emphasis on neural networks.

INTELLIGENT CONTROL FOR HIGH AUTONOMY SYSTEMS

From a control systems point of view the use of Intelligent control methods is a natural next step in the quest for building systems with higher degrees of autonomy. These ideas are discussed below.

In the design of controllers for complex dynamical systems there are needs today that cannot be successfully addressed with the existing conventional control theory. They mainly pertain to the area of uncertainty. Heuristic methods may be needed to tune the parameters of an adaptive control law. New control laws to perform novel control functions to meet new objectives should be designed, while the system is in operation. Learning from past experience and planning control actions may be necessary. Failure detection and identification is needed. Such functions have been performed in the past by human operators. To increase the speed of response, to relieve the operators from mundane tasks, to protect them from hazards, high degree of autonomy is desired. To achieve this, high level decision making techniques for reasoning under uncertainty and taking actions must be utilized. These techniques, if used by humans, may be attributed to intelligent behavior. Hence, one way to achieve high degree of autonomy is to utilize high level decision making techniques, intelligent methods, in the autonomous controller. Autonomy is the objective, and intelligent controllers are one way to achieve it.

Evolution of Control Systems and the Quest for Higher Autonomy

The first feedback device on record was the water clock invented by the Greek Ktesibios in Alexandria Egypt around the 3rd century B.C.. This was certainly a successful device as water clocks of similar design were still being made in Baghdad when the Mongols captured that city in 1258 A.D.. The first mathematical model to describe plant behavior for control purposes is attributed to J.C. Maxwell, of the Maxwell equations' fame, who in 1868 used differential equations to explain instability problems encountered with James Watt's flyball governor; the governor was introduced in 1769 to regulate the speed of steam engine vehicles. When J.C. Maxwell used mathematical modeling and methods to explain instability problems encountered with James Watt's flyball governor, it demonstrated the importance and usefulness of mathematical models and methods in understanding complex phenomena and signaled the beginning of mathematical system and control theory. It also signaled the end of the era of intuitive invention. Control theory made significant strides in the past 120 years, with the use of frequency domain methods and Laplace transforms in the 1930s and 1940s and the development of optimal control methods and state space analysis in the 1950s and 1960s. Optimal control in the 1950s and 1960s, followed by progress in stochastic, robust, adaptive and nonlinear control methods in the 1960s to today, have made it possible to control more accurately significantly more complex dynamical systems than the original flyball

governor.

Conventional control systems are designed today using mathematical models of physical systems. A mathematical model, which captures the dynamical behavior of interest, is chosen and then control design techniques are applied, aided by CAD packages, to design the mathematical model of an appropriate controller. The controller is then realized via hardware or software and it is used to control the physical system. The procedure may take several iterations. The mathematical model of the system must be "simple enough" so that it can be analyzed with available mathematical techniques, and "accurate enough" to describe the important aspects of the relevant dynamical behavior. It approximates the behavior of a plant in the neighborhood of an operating point.

The control methods and the underlying mathematical theory were developed to meet the ever increasing control needs of our technology. The need to achieve the demanding control specifications for increasingly complex dynamical systems has been addressed by using more complex mathematical models such as nonlinear and stochastic ones, and by developing more sophisticated design algorithms for, say, optimal control. The use of highly complex mathematical models however, can seriously inhibit our ability to develop control algorithms. Fortunately, simpler plant models, for example linear models, can be used in the control design; this is possible because of the feedback used in control which can tolerate significant model uncertainties. When the fixed feedback controllers are not adequate, then adaptive controllers are used. Controllers can then be designed to meet the specifications around an operating point, where the linear model is valid and then via a scheduler a controller emerges which can accomplish the control objectives over the whole operating range. This is, for example, the method typically used for aircraft flight control and it is a method to design fixed controllers for certain classes of nonlinear systems. Adaptive control in conventional control theory has a specific and rather narrow meaning. In particular it typically refers to adapting to variations in the constant coefficients in the equations describing the linear plant; these new coefficient values are identified and then used, directly or indirectly, to reassign the values of the constant coefficients in the equations describing the linear controller. Adaptive controllers provide for wider operating ranges than fixed controllers and so conventional adaptive control systems can be considered to have higher degrees of autonomy than control systems employing fixed feedback controllers.

Intelligent Control for High Autonomy Systems

There are cases where we need to significantly increase the operating range of the system. We must be able to deal effectively with significant uncertainties in models of increasingly complex dynamical systems in addition to increasing the validity range of our control methods. We need to cope with significant unmodelled and unanticipated changes in the plant, in the environment and in the control objectives. This will involve the use of intelligent decision making processes to generate control actions so that certain performance level is maintained even though there are drastic changes in the operating conditions. I have found useful to keep in mind an example that helps set goals for the future and also teaches humility, as it shows how difficult, demanding and complex autonomous systems can be: Currently, if there is a problem on the space shuttle, the problem is addressed by the large number of

engineers working in Houston Control, the ground station. When the problem is solved the specific detailed instructions about how to deal with the problem are sent to the shuttle. Imagine the time when we will need the tools and expertise of all Houston Control engineers aboard the space shuttle, or an other space vehicle for extended space travel. What needs to be achieved to accomplish this goal is certainly highly challenging!

In view of the above it is quite clear that in the control of systems there are requirements today that cannot be successfully addressed with the existing conventional control theory. It should be pointed out that several functions proposed in later sections, to be part of the high autonomy control system, have been performed in the past by separate systems; examples include fault trees in chemical process control for failure diagnosis and hazard analysis, and control system design via expert systems.

An Intelligent Control Architecture For High Autonomy Systems

To illustrate the concepts and ideas involved and to provide a more concrete framework to discuss the issues, a hierarchical functional architecture of an intelligent controller that is used to attain high degrees of autonomy in future space vehicles is briefly outlined; full details can be found in reference 16. This hierarchical architecture has three levels, the Execution Level, the Coordination Level, and the Management or Organization Level. The architecture exhibits certain characteristics, which have been shown in the literature to be necessary and desirable in autonomous systems. Based on this architecture we identify the important fundamental issues and concepts that are needed for an autonomous control theory.

Architecture Overview: Structure and Characteristics: The overall functional architecture for an autonomous controller is given by the architectural schematic of the Figure 1, below. This is a functional architecture rather than a hardware processing one; therefore, it does not specify the arrangement and duties of the hardware used to implement the functions described. Note that the processing architecture also depends on the characteristics of the current processing technology; centralized or distributed processing may be chosen for function implementation depending on available computer technology.




Figure 1. Intelligent Autonomous Controller Functional Architecture. The three levels of a

hierarchical Intelligent control architecture are the Execution Level, the Coordination Level, and the Management or Organization Level.

The architecture in Figure 1 has three levels. At the lowest level, the Execution Level, there is the interface to the vehicle and its environment (the process in the figure) via the sensors and actuators. At the highest level, the Management or Organization Level, there is the interface to the pilot and crew, ground station, or onboard systems. The middle level, called the Coordination Level, provides the link between the Execution Level and the Management Level. Note that we follow the somewhat standard viewpoint that there are three major levels in the hierarchy.

Figure 1. Intelligent Autonomous Controller Functional Architecture. The three levels of a hierarchical Intelligent control architecture are the Execution Level, the Coordination Level, and the Management or Organization Level.

It must be stressed that the system may have more or fewer than three levels which however can be conceptually combined into three levels. Some characteristics of the system which dictate the actual number of levels are the extent to which the operator can intervene in the system's operations, the degree of autonomy or level of intelligence in the various subsystems, the dexterity of the subsystems, and the hierarchical characteristics of the plant. Note that the three levels shown here in Figure 1 are applicable to most architectures of intelligent autonomous controllers, by grouping together sublevels of the architecture if necessary. The lowest, Execution Level involves conventional control algorithms, while the highest, Management and Organization Level involves only higher level, intelligent, decision making methods. The Coordination Level is the level which provides the interface between the actions of the other two levels and it uses a combination of conventional and intelligent decision making methods.

The sensors and actuators are implemented mainly with hardware. Software and perhaps hardware are used to implement the Execution Level. Mainly software is used for both the Coordination and Management Levels. There are multiple copies of the control functions at each level, more at the lower and fewer at the higher levels. Note that the autonomous controller is only one of the autonomous systems on the space vehicle. It is responsible for all the functions related to the control of the physical system and allows for continuous on-line development of the autonomous controller and to provide for various phases of mission operations. The tier structure of the architecture allows us to build on existing advanced control theory. Development progresses, creating each time higher level adaptation and a new system which can be operated and tested independently. The autonomous controller performs many of the functions currently performed by the pilot, crew, or ground station. The pilot and crew are thus relieved from mundane tasks and some of the ground station functions are brought aboard the vehicle. In this way the degree of autonomy of the vehicle is increased.

Functional Operation: In Figure 1, commands are issued by higher levels to lower levels and response data flows from lower levels upwards. Parameters of subsystems can be altered by systems one level above them in the hierarchy. There is a delegation and distribution of tasks from higher to lower levels and a layered distribution of decision making authority. At each

level, some preprocessing occurs before information is sent to higher levels. If requested, data can be passed from the lowest subsystem to the highest, e.g., for display. All subsystems provide status and health information to higher levels. Human intervention is allowed even at the control implementation supervisor level, with the commands however passed down from the upper levels of the hierarchy.

Here is a simple illustrative example to clarify the overall operation of the autonomous controller. Suppose that the pilot desires to repair a satellite. After dialogue with the Management Level via the interface, the task is refined to "repair satellite using robot A". This is a decision made using the capability assessing, performance monitoring, and planning functions of the Management Level. The Management Level decides if the repair is possible under the current performance level of the system, and in view of near term other planned functions. Using its planning capabilities, it then sends a sequence of subtasks to the Coordination Level sufficient to achieve the repair. This sequence could be to order robot A to: "go to satellite at coordinates xyz", "open repair hatch", "repair". The Coordination Level, using its planner, divides say the first subtask, "go to satellite at coordinates xyz", into smaller subtasks: "go from start to $x_1y_1z_1$ ", then "maneuver around obstacle", "move to $x_2y_2z_2$ ", ..., "arrive at the repair site and wait". The other subtasks are divided in a similar manner. This information is passed to a control implementation supervisor at the Coordination Level, which recognizes the task, and uses stored control laws to accomplish the objective. The subtask "go from start to $x_1y_1z_1$ ", can for example, be implemented using stored control algorithms to first, proceed forward 10 meters, to the right 15 degrees, etc. These control algorithms are executed in the controller at the Execution Level utilizing sensor information; the control actions are implemented via the actuators.

Characteristics of Hierarchical Intelligent Controllers for High Autonomy Systems

Based on the architecture described above, important fundamental concepts and characteristics that are needed for an autonomous intelligent control theory are now identified. The fundamental issues which must be addressed for a quantitative theory of autonomous intelligent control are discussed.

There is a successive delegation of duties from the higher to lower levels; consequently the number of distinct tasks increases as we go down the hierarchy. Higher levels are concerned with slower aspects of the system's behavior and with its larger portions, or broader aspects. There is then a smaller contextual horizon at lower levels, i.e. the control decisions are made by considering less information. Also notice that higher levels are concerned with longer time horizons than lower levels. Due to the fact that there is the need for high level decision making abilities at the higher levels in the hierarchy, the proposition has been put forth that there is increasing intelligence as one moves from the lower to the higher levels. This is reflected in the use of fewer conventional numeric-algorithmic methods at higher levels as well as the use of more symbolic-decision making methods. This is the "principle of increasing intelligence with decreasing precision" of Saridis; see also reference 5 and the references therein. The decreasing precision is reflected by a decrease in time scale density, decrease in bandwidth or system rate, and a decrease in the decision (control action) rate. (These properties have been studied for a class of hierarchical systems in reference 17.) All

these characteristics lead to a decrease in granularity of models used, or equivalently, to an increase in model abstractness. Model granularity also depends on the dexterity of the autonomous controller.

It is important at this point to discuss briefly the dexterity of the controller. The Execution Level of a highly dexterous controller is very sophisticated and it can accomplish complex control tasks. The Coordination Level can issue commands to the controller such as "move 15 centimeters to the right", and "grip standard, fixed dimension cylinder", in a dexterous controller, or it can completely dictate each mode of each joint (in a manipulator) "move joint 1, 15 degrees", then "move joint 5, 3 degrees", etc. in a less dexterous one. The simplicity, and level of abstractness of macro commands in an autonomous controller depends on its dexterity. The more sophisticated the Execution Level is, the simpler are the commands that the control implementation supervisor needs to issue. Notice that a very dexterous robot arm may itself have a number of autonomous functions. If two such dexterous arms were used to complete a task which required the coordination of their actions then the arms would be considered to be two dexterous actuators and a new supervisory autonomous controller would be placed on top for the supervision and coordination task. In general, this can happen recursively, adding more intelligent autonomous controllers as the lower level tasks, accomplished by autonomous systems, need to be supervised.

There is an ongoing evolution of the intelligent functions of an autonomous controller. It is interesting to observe the following: Although there are characteristics which separate intelligent from non-intelligent systems, as intelligent systems evolve, the distinction becomes less clear. Systems which were originally considered intelligent evolve to gain more character of what are considered to be non-intelligent, numeric-algorithmic systems. An example is a route planner. Although there are AI route planning systems, as problems like route planning become better understood, more conventional numeric-algorithmic solutions are developed. The AI methods which are used in intelligent systems, help us to understand complex problems so we can organize and synthesize new approaches to problem solving, in addition to being problem solving techniques themselves. AI techniques can be viewed as research vehicles for solving very complex problems. As the problem solution develops, purely algorithmic approaches, which have desirable implementation characteristics, substitute AI techniques and play a greater role in the solution of the problem. It is for this reason that we concentrate on achieving autonomy and not on whether the underlying system can be considered "intelligent".

Models for Intelligent Controllers

In highly autonomous control systems, the plant is normally so complex that it is either impossible or inappropriate to describe it with conventional mathematical system models such as differential or difference equations. Even though it might be possible to accurately describe some system with highly complex nonlinear differential equations, it may be inappropriate if this description makes subsequent analysis too difficult or too computationally complex to be useful. The complexity of the plant model needed in design depends on both the complexity of the physical system and on how demanding the design specifications are. There is a tradeoff between model complexity and our ability to perform analysis on the system via the

model. However, if the control performance specifications are not too demanding, a more abstract, higher level, model can be utilized, which will make subsequent analysis simpler. This model intentionally ignores some of the system characteristics, specifically those that need not be considered in attempting to meet the particular performance specifications; see also the discussion on hybrid systems later in this article. For example, a simple temperature controller could ignore almost all dynamics of the house or the office and consider only a temperature threshold model of the system to switch the furnace off or on.

Discrete Event System (DES) models using finite automata, Petri nets, queuing network models, Markov chains, etc. are quite useful for modeling the higher level decision making processes in the intelligent autonomous controller. The choice of whether to use such models will, of course, depend on what properties of the autonomous system need to be studied.

The quantitative, systematic techniques for modeling, analysis, and design of control systems are of central and utmost practical importance in conventional control theory. Similar techniques for intelligent autonomous controllers do not exist. This is mainly due to the hybrid structure (nonuniform, non homogeneous nature) of the dynamical systems under consideration; they include both continuous-state and discrete-state systems. Modeling techniques for intelligent autonomous systems must be able to support a macroscopic view of the dynamical system, hence it is necessary to represent both numeric and symbolic information. The non uniform components of the intelligent controller all take part in the generation of the low level control inputs to the dynamical system, therefore they all must be considered in a complete analysis. Research could begin by using different models for different components of the intelligent autonomous controller since much can be attained by using the best available models for the various components of the architecture and joining them via some appropriate interconnecting structure. For instance, systems that are modeled with a logical discrete event system (DES) model at the higher levels and a difference or differential equation at the lower level should be examined; see the discussion on hybrid systems later in this article. In any case, good understanding of hierarchical models is necessary for the analysis and synthesis of intelligent autonomous controllers.

Research Directions

One can roughly categorize research in the area of intelligent autonomous control into two areas: conventional control theoretic research, addressing the control functions at the Execution and Coordination levels, and the modeling, analysis, and design of higher level decision making systems found in the Management and Coordination levels.

It is important to note that in order to obtain a high degree of autonomy it is necessary to adapt or learn. Neural networks offer methodologies to perform learning functions in the intelligent autonomous controller. In general, there are potential applications of neural networks at all levels of hierarchical intelligent controllers that provide higher degrees of autonomy to systems. Neural networks are useful at the lowest Execution level - where the conventional control algorithms are implemented via hardware and software - through the Coordination level, to the highest Management level, where decisions are being made based on possibly uncertain and/or incomplete information. One may point out that at the Execution

level - conventional control level - neural network properties such the ability for function approximation and the potential for parallel implementation appear to be very important. In contrast, at higher levels abilities such as pattern classification and the ability to store information in a, say, associative memory appear to be of significant interest. Machine learning is of course important at all levels.

We stress that in control systems with high degrees of autonomy we seek to significantly widen the operating range of the system so that significant failures and environmental changes can occur and performance will still be maintained. All of the conventional control techniques are useful in the development of autonomous controllers and they are relevant to the study of autonomous control. It is the case however, that certain techniques are more suitable for interfacing to the autonomous controller and for compensating for significant system failures. For instance the area of "restructurable" or "reconfigurable" control systems studies techniques to reconfigure controllers when significant failures occur.

Conventional modeling, analysis, and design methods should be used, whenever they are applicable, for the components of the intelligent autonomous control system as well as fuzzy controllers. For instance, they should be used at the Execution Level of many autonomous controllers. The symbolic/numeric interface is a very important issue; consequently it should be included in any analysis. There is a need for systematically generating less detailed, more abstract models from differential/difference equation models to be used in higher levels of the autonomous controller; see discussion below on hybrid systems. Tools for the implementation of this information extraction also need to be developed. In this way conventional analysis can be used in conjunction with the developed analysis methods to obtain an overall quantitative, systematic analysis paradigm for intelligent autonomous control systems. In short, we propose to use hybrid modeling, analysis, and design techniques for non uniform systems. This approach is not unlike the approaches used in the study of any complex phenomena by the scientific and engineering communities.

HYBRID SYSTEMS

Hybrid control systems contain two distinct types of systems, systems with continuous dynamics and systems with discrete dynamics, that interact with each other. Their study is central in designing intelligent control systems with high degree of autonomy and it is essential in designing discrete event supervisory controllers for continuous systems; see references 1, 18-23.

Hybrid control systems typically arise when continuous processes interact with, or are supervised by, sequential machines. Examples of hybrid control systems are common in practice and are found in such applications as flexible manufacturing, chemical process control, electric power distribution and computer communication networks. A simple example of a hybrid control system is the heating and cooling system of a typical home. The furnace and air conditioner, along with the heat flow characteristics of the home, form a continuous-time system which is to be controlled. The thermostat is a simple discrete event system which basically handles the symbols {too hot, too cold} and {normal}. The temperature of the room is translated into these representations in the thermostat and the

thermostat's response is translated back to electrical currents which control the furnace, air conditioner, blower, etc.

Since the continuous and discrete dynamics coexist and interact with each other it is important to develop models that accurately describe the dynamic behavior of such hybrid systems. In this way it is possible to develop control strategies that fully take into consideration the relation and interaction of the continuous and discrete parts of the system. In the past, models for the continuous and discrete event subsystems were developed separately; the control law was then derived in a rather empirical fashion, except in special cases such as the case of digital controllers for linear time-invariant systems. The study of hybrid systems provides the backbone for the formulation and implementation of learning control policies. In such policies, the control acquires knowledge (discrete data) to improve the behavior of the system as it evolves in time. Hybrid systems has become a distinctive area of study due to opportunities to improve on traditional control and estimation technologies by providing computationally effective methodologies for the implementation of digital programs that design or modify the control law in response to sensor detected events, or as a result of adaptation and learning. The interested reader should consult references 20-23.

Certain important issues in hybrid systems are now briefly discussed using a paradigm of a continuous systems supervised by a discrete event system (DES) controller from references 18, 19. The hybrid control system of interest here consists of a continuous-state system to be controlled, also called the plant, and a discrete-state controller connected to the plant via an interface; see Figure 2.




Figure 2. Hybrid Supervisory Control Architecture. The interface receives continuous measurements $z(t)$ and issues a sequence of symbols $\{z(i)\}$ which the DES controller processes to issue a sequence of control symbols $\{r(i)\}$. These are translated by the interface to (piecewise) continuous input commands $r(t)$.

The plant contains all continuous-state subsystems of the hybrid control system, such as any conventional continuous-state controllers that may have been developed, a clock if time and synchronous operations are to be modeled, etc.. The controller is an event driven, asynchronous discrete event system (DES), described by a finite state automaton or an ordinary Petri net. The hybrid control system also contains an interface that provides the means for communication between the continuous-state plant and the DES controller. The interface receives information from the plant in the form of a measurement of a continuous variable $z(t)$, such as the continuous state, and issues a sequence of symbols $\{z(i)\}$ to the DES

controller. It also receives a sequence of control symbols $\{r(i)\}$ from the controller and issues (piecewise) continuous input commands $r(t)$ to the plant.

The interface plays a key role in determining the dynamics and the behavior of the hybrid control system. Understanding how the interface affects the properties of the hybrid system is one of the fundamental issues in the theory of hybrid control systems. The interface can be chosen to be simply a partitioning of the state space; see reference 18. If memory is necessary to derive an effective control, it is included in the DES controller and not in the interface. Also the piecewise continuous command signal issued by the interface is simply a staircase signal, not unlike the output of a zero-order hold in a digital control system. Including an appropriate continuous system at (the input of) the plant, signals such as ramps, sinusoids, etc., can be generated if desired. So the simple interface is used without loss of generality. It allows analysis of the hybrid control system with development of properties such as controllability, stability and determinism, in addition to control design methodologies; see references 18, 19. In general the design of the interface depends not only on the plant to be controlled, but also on the control policies available, as well as on the control goals. Depending on the control goals, one may or may not need, for example, detailed state information; this corresponds to small or large regions in the partition of the measured signal space (or greater or lower granularity). This is, of course, not surprising as it is rather well known that to stabilize a system, for example, requires less detailed information about the system's dynamic behavior than to do say tracking. The fewer the distinct regions in the partitioned signal space, the simpler (fewer states) the resulting DES plant model and the simpler the DES controller design. Since the systems to be controlled via hybrid controllers are typically complex, it is important to make every effort to use only the necessary information to attain the control goals; as this leads to simpler interfaces that issue only the necessary number of distinct symbols, and to simpler DES plant models and controllers. The question of systematically determining the minimum amount of information needed from the plant in order to achieve specific control goals via a number of specialized control policies is an important question.

INTELLIGENT CONTROL AS A DISTINCT RESEARCH AREA

There may be the temptation to classify the area of intelligent autonomous systems as simply a collection of methods and ideas already addressed elsewhere, the need only being some kind of intelligent assembly and integration of known techniques. This is of course not true. The theory of control systems is not covered by say the area of applied mathematics, because control has different needs and therefore asks different questions. For example, while in applied mathematics the different solutions of differential equations under different initial conditions and forcing functions are of interest, in control one typically is interested in finding the forcing functions that generate solutions, that is system trajectories, that satisfy certain conditions. This is a different problem, related to the first, but its solution requires the development of quite different methods. In a rather analogous fashion the problems of interest in intelligent systems require development of novel concepts, approaches and methods. In particular while computer science typically deals with static systems and no real-time requirements, control systems typically are dynamic and all control laws, intelligent or not, must be able to control the system in real time. So in most cases one cannot really just directly

apply computer science methods to these problems. Modifications and extensions are typically necessary for example in the quantitative models used to study such systems. And although say Petri nets may be adequate to model and study the autonomous behavior at certain levels of the hierarchy, these models may not be appropriate to address certain questions of importance to control systems such as stability, without further development and modifications. In addition, there are problems in intelligent autonomous control systems that are novel and so they have not studied before at any depth. Such is the case of hybrid systems for example that combine systems of continuous and discrete state. The marriage of all these fields can only be beneficial to all. Computer science and operation research methods are increasingly used in control problems, while control system concepts such as feedback, and methods that are based on rigorous mathematical framework can provide the base for new theories and methods in those areas.

BIBLIOGRAPHY

1. P.J. Antsaklis, Defining Intelligent Control, Report of the IEEE Control Systems Society Task Force on Intelligent Control (P.J. Antsaklis, Chair; J. Albus, K. Baheti, J. D. Birdwell, M. Lemmon, M. Mataric, A. Meystel, K. Narendra, K. Passino, H. Rauch, G. Saridis, H. Stephanou, P. Werbos, Members). In IEEE Control Systems Magazine, 14(3), 4-5, 58-66, June 1994.
2. P.J. Antsaklis and K.M. Passino (eds.), An Introduction to Intelligent and Autonomous Control, Norwell: Kluwer Academic Publishers, 1993.
3. J.S. Albus, Outline for a Theory of Intelligence, IEEE Transactions on Systems, Man and Cybernetics, 21(3): 432-509, May/June 1991. 4. A. Meystel, Autonomous Mobile Robots. River Edge, NJ: World Scientific, 1991.
5. G.N. Saridis and K.P. Valavanis, Analytical Design of Intelligent Machines, Automatica, 24(2), 123-133, March 1988.
6. K.P. Valavanis and G.N. Saridis, Intelligent Robotic System Theory: Design and Applications, Norwell: Kluwer Academic Publishers, 1992.
7. D.A. White, D.A. Sofge (eds.), Handbook of Intelligent Control Neural, Fuzzy, and Adaptive Approaches. New York: Van Nostrand Reinhold, 1992.
8. M.M. Gupta and N.K. Sinha (eds.), Intelligent Control: Theory and Practice, Piscataway: IEEE Press, 1994.
9. R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, Machine Learning-An Artificial Intelligence Approach, Palo Alto: Tioga, 1983. Also R.S. Michalski, and G.Tecuci (eds.), Machine Learning - A Multistrategy Approach. San Fransisco: Morgan-Kaufmann, 1994, Vol IV,.
10. J.W. Shavlik and T.G. Dietterich (eds.) Readings in Machine Learning. San Fransisco: Morgan-Kaufmann, 1990.
11. K.S. Fu, Learning Control Systems--Review and Outlook, IEEE Transactions on Automatic Control, 15, 210-221, April 1970.
12. J.M. Mendel, and K.S. Fu, Adaptive, Learning and Pattern Recognition Systems, New York: Academic Press, 1970.
13. J. Sklansky, Learning Systems for Automatic Control, IEEE Transactions on Automatic Control, 11, 6-19, January 1966.

14. YA.Z. Tsypkin, *Adaptation and Learning in Automatic Systems*, New York: Academic Press, 1971.
15. YA.Z. Tsypkin, Self-Learning-What is it?, *IEEE Trans on Automatic Control*, 13, 608-612, December 1968.
16. P.J.Antsaklis and K.M.Passino, Introduction to Intelligent Control Systems with High Degree of Autonomy, in *An Introduction to Intelligent and Autonomous Control*, P.J.Antsaklis and K.M.Passino (eds.), Norwell: Kluwer Academic Publishers, 1993.
17. K.M. Passino, P.J. Antsaklis, Event Rates and Aggregation in Hierarchical Discrete Event Systems, *Journal of Discrete Event Dynamic Systems: Theory and Applications*, 1(3), 271-287, 1992.
18. P. J. Antsaklis, J. A. Stiver and M. D. Lemmon, Hybrid System Modeling and Autonomous Control Systems. In R.L. Grossman, A. Nerode, A.P. Ravn, H. Rischel (eds.), *Hybrid Systems*. New York: Springer-Verlag, 1993, LNCS 736, 366-392.
19. J.A. Stiver, P.J. Antsaklis and M.D. Lemmon, A Logical DES Approach to the Design of Hybrid Control Systems, *Mathematical and Computer Modeling*, Special issue on Discrete Event Systems, 23(11/12), June 1996.
20. R L Grossman, A Nerode, A P Ravn, H Rischel (eds.), *Hybrid Systems*, New York: Springer-Verlag, 1993, LNCS 736.
21. P. Antsaklis, W. Kohn, A. Nerode, S. Sastry (eds.), *Hybrid Systems II*, New York: Springer-Verlag, 1995, LNCS 999.
22. R. Alur, T. Henzinger and E. Sontag (eds.), *Hybrid Systems III*, New York: Springer-Verlag, 1996, LNCS 1066.
23. P. Antsaklis, W. Kohn, A. Nerode, S. Sastry (eds.), *Hybrid Systems IV*, New York: Springer-Verlag, 1997, LNCS.