

THE QUEST FOR AUTONOMY REVISITED

Panos Antsaklis, University of Notre Dame

This is a concept paper and it is an update of earlier papers on autonomous intelligent control systems by the author and collaborators. In the present paper the emphasis is on simplifying and explaining important concepts on autonomy and neither completeness of exposition is sought nor supporting documentation is included beyond the sources of this conceptual framework. . It is important to note that for additional details and especially for references the reader should consult the original papers (e.g. [1,5]).

I. Introduction to Autonomous Systems and Autonomous Controllers

Systems with ever increasing degrees of autonomy are more prevalent and important today than ever before. Well known examples include Unmanned Aerial Vehicles (UAV), Autonomous Underwater Vehicles (AUV), office and residential buildings that regulate their energy consumption while adapting to the needs of their inhabitants (smart buildings), safety systems and environmentally friendly energy systems in automobiles (smart cars, smart highways). The trend towards increased autonomy has been around for centuries. The recent surge is fueled primarily by technological leaps in hardware and software and successes in integrating tightly the physical and computer worlds, as in the Cyber-Physical Systems (CPS), where it is not obvious where system properties are attributed to, the physical or the cyber components of the system. The recent successes in increasing autonomy in engineering systems are only the beginning of many more to come.

This article emphasizes characteristics that are necessary for high degrees or high levels of autonomy-yes there are levels or degrees of autonomy. As it will be noted later, adaptation and learning, failure diagnosis and identification, control reconfiguration and planning are some of these characteristics.

When one considers humans collaborating with engineered systems, then the overall system that includes humans in the loop may be considered (fully) autonomous with respect to a set of goals. Depending on the role of the humans in the loop and the level of control authority humans exert, the remaining system will have different degrees or levels of autonomy. These ideas are explained in this article.

It is important to point out that in our discussion of autonomy, the system under consideration always has a set of goals to be achieved and a control mechanism to achieve them. This implies that *every autonomous system is a control system*.

It is useful to think of a system as being surrounded by a boundary separating it from its environment. The system acts upon its environment through its outputs and receives inputs in the form of disturbances or additional information. What the system includes within its boundary, expressed via the particular system model used, depends of course

on the goals and the characteristics/properties used to achieve its goals. It is useful to assume as a starting point that the system may also include a human operator who acts as a highly able controller. So in an automobile, if the goal is for example to keep the vehicle inside a lane while travelling with constant speed, the system may consist of the vehicle and the driver where the system attains its goals in the presence of uncertainties/disturbances, such as gust of wind and road incline.

One can envision an autonomous system consisting of two subsystems, a (sub-) system to be controlled (the plant, as it is called in the control literature) and a controller to be designed. Note that this separation of the plant and the controller is common in the field of control systems theory, although it may be somewhat restrictive in autonomous systems as the assumption is that one can separate the plant from the controller which may not necessarily be easy to do in some cases. However it is a useful concept and it will be used in the current paper. The controller may include a human in the loop in which case it may achieve full autonomy and we will call such controller *autonomous controller* (meaning that such controller is in itself an autonomous system with respect to a new set of goals, that of providing the right kind or control policies to the plant; or alternatively, a controller will be called an *autonomous controller* if it causes the system to become autonomous). The controller may achieve only partial autonomy with or without the human in the loop (meaning that it may need extra help from humans or other systems to attain full autonomy), in which case we will call such controller, *controller with high or low degree of autonomy* (or a *sub- or a semi-autonomous controller*). An example of an autonomous controller in an automobile is the system consisting of the human driver and all the control systems in the car with the plant being the vehicle and the goals of the autonomous controller being to provide the right steering and gas pedal commands so the vehicle maintains its course within a lane and at certain (approximately) constant speed. If one considers in this case the controller to consist of just the control systems of the car without the driver then the controller is not autonomous but semi-autonomous.

This article presents a view of autonomous systems and autonomous controllers and it is based on earlier work [1-3], which was in turn based on the more detailed papers [4-6] and the references therein. It is an updated version and includes extensions based on lessons learned. One of the main issues in the above quoted literature [1-6] was the connection to and the meaning of the term “Intelligent Control.” In [7,8] for example several definitions of Intelligent Control were presented. The difficulty was, and still is, the fact that what constitutes intelligence is not universally agreed upon (today the IQ tests are still controversial and are not widely used around the world) and the issue is still debated, as there are many different strong views. However, throughout though the above quoted literature the main point that was consistently made was that *autonomy should be the property of main interest and if for high degrees of autonomy methods that are considered intelligent are used then the name Intelligent Control may be justified* [5, p.345].

Note that the themes of “Quest for Autonomy” and “Autonomy is the goal and intelligent methods are ways to achieve it”, together with discussions of characteristics and properties of intelligent autonomous controllers were repeated in plenary lectures and

talks and guest editorials [9-17], in books [18-21] and in publications on hybrid and discrete event systems, see for example [22-24]. The expression “*Quest for Autonomy*” makes this exact point that autonomy is the property of interest, while the term Intelligent describes in a eye catching way the methodologies used, not unlike today’s term “Smart” which is used in many applications, such as Smart Grid, Smart Phones, Smart Buildings etc. So autonomy is the goal! Appendix B includes materials on Intelligent Control, and its relation to Autonomous Control.

In the following, Autonomous Systems and Autonomous Controllers are discussed in Section II. Section III describes the functions of an autonomous controller and Section IV provides brief historical comments on controller design. Section V describes a hierarchical functional architecture to implement the functions of autonomous controllers. The characteristics of autonomous control systems are discussed in Section VI and mathematical models are discussed in Section VII. Concluding remarks may be found in Section VIII and references in section IX. Appendix A contains comparisons of autonomous and conventional control and in Appendix B intelligent control is discussed, including characteristics or dimensions of intelligence and the relation of intelligent to autonomous control.

II. Autonomous Systems and Autonomous Controllers

It is important to stress again that in our discussion of autonomy, the system under consideration always has a set of goals to be achieved and a control mechanism, a controller, to achieve them. This implies that *every autonomous system is a control system*.

Autonomous means having the ability and authority for self-government. ***A system is autonomous regarding a set of goals, with respect to a set of measures of intervention (by humans or other systems)***. A regular feedback control system for example is autonomous regarding stability goals with respect to (certain level or degree of external and internal) disturbances. This is because stability is maintained even when there are internal system parameter variations and external disturbances. This robustness is due to feedback closed-loop mechanism that compensates for uncertainties; on the other hand, an open-loop system with feed-forward control has none of these robustness properties and no autonomy regarding stability with respect to parameter variations and disturbances.

Alternatively, a perhaps more useful working definition of an autonomous system is that a system has *high or low degree or level of autonomy regarding a goal*. By high degree/level of autonomy it is meant that the degree/level of human intervention (or perhaps intervention by other engineered systems) is low, while by low degree/level of autonomy, a high degree/level of human intervention is implied.

Human in the Loop and Adaptive Autonomy: Humans or other systems may insert themselves at certain levels of the functional hierarchy (that correspond to levels of

autonomy), as described later in the present paper, and take over control functions. For example, humans may insert themselves to take over planning, FDI, learning functions (see descriptions of functions later in the paper). Or they may insert themselves to take over lower control functions e.g. a driver may want to take over from the ABS system and perform the braking pumping action herself. As mentioned above, this reference to levels connects with the hierarchical functional architecture for the autonomous controller discussed later in this paper.

Autonomous controllers have the ability and authority for self-governance in the performance of control functions. They are composed of a collection of hardware and software, which can perform the necessary control functions, without external intervention, over extended time periods. *Note that a controller will be called autonomous controller when its functions make the system of interest an autonomous system.* Alternatively, an autonomous controller can be seen itself as an autonomous system with goals to apply appropriate controls that make the system of interest autonomous.

There are several ***degrees or levels of autonomy***. A fully autonomous controller should perhaps have the ability to perform even hardware repair, if one of its components fails. Note that conventional fixed controllers can be considered to have a *low degree* of autonomy since they can only tolerate a restricted class of plant parameter variations and disturbances. To achieve a *high degree* of autonomy, the controller must be able to perform a number of functions in addition to the conventional control functions such as tracking and regulation. These additional functions, which include the ability to accommodate for drastic system failures, are discussed in this article. Autonomous controllers can of course be used in a variety of systems from manufacturing to unmanned space, atmospheric, ground, and underwater exploratory vehicles. This introduction to autonomous control will be developed around a space vehicle application so that a) concrete examples for the various control functions, and fundamental characteristics of autonomous control can be given; and b) so that the development addresses relatively well defined control needs rather than abstract requirements. Furthermore, the autonomous control of space vehicles is highly demanding; consequently the developed architecture is general enough to encompass all related autonomy issues. It should be stressed that all the results presented here apply to any autonomous control system. In other classes of applications, the architecture, or parts of it, can be used directly and the same fundamental concepts and characteristics identified here are valid.

We begin by describing a conceptual functional architecture of the autonomous controller necessary for the operation of future advanced space vehicles that was developed in [4-6]. This hierarchical architecture is certainly one of many possible control architectures. The choice is dependant on the particular problem addressed. We refer to it as a hierarchical functional architecture-hierarchies make it possible for us to handle complexity better (see Appendix and the discussion on characteristics of intelligent control)-but the architecture in fact is a *heterarchy* as it also allows direct communication among elements on the same level.

The concepts and methods needed to design successfully such an autonomous controller are introduced and discussed. A hierarchical functional autonomous controller architecture for a future spacecraft is described; it is designed to ensure the autonomous operation of the control system and it allows interaction with the pilot/ground station and the systems on board the autonomous vehicle. A command by the pilot or the ground station is executed by dividing it into appropriate subtasks, which are then performed by the controller. The controller can deal with unexpected situations, new control tasks, and failures within limits. To achieve this, high-level decision making techniques for reasoning under uncertainty and taking actions must be utilized. These techniques, if used by humans, are attributed to *intelligent* behavior. Hence, one way to achieve autonomy, in some applications, is to utilize high-level decision making techniques, “intelligent” methods, in the autonomous controller. Remember that *autonomy is the objective, and “intelligent” or “smart” controllers are one way to achieve it.*

III. Autonomous Controller Functions

Autonomous control systems must perform well under significant uncertainties in the plant and the environment for extended periods of time and they must be able to compensate for system failures without external intervention. Such autonomous behavior is a very desirable characteristic of advanced systems. An autonomous controller provides high level *adaptation* to changes in the plant and environment. To achieve autonomy the methods used for control system design should utilize both

- a) algorithmic-numeric methods, based on the state-of-the-art conventional control, identification, estimation, and communication theory, and
- b) decision making-symbolic methods, such as the ones developed in computer science (e.g., automata theory), and specifically in the field of AI.

In addition to supervising and tuning the control algorithms, the autonomous controller must also provide a high degree of tolerance to failures. To ensure system reliability, failures must first be detected, isolated, and identified (and if possible contained), and subsequently a new control law must be designed if it is deemed necessary.

The autonomous controller must be capable of planning the necessary sequence of control actions to be taken to accomplish a complicated task.

It must be able to interface to other systems as well as with the operator, and it may need learning capabilities to enhance its performance while in operation. It is for these reasons that advanced planning, learning, and expert systems, among others, must work together with conventional control systems in order to achieve autonomy.

The need for quantitative methods to model and analyze the dynamical behavior of such autonomous systems presents significant challenges. The development of autonomous controllers requires significant interdisciplinary research effort as it integrates concepts and methods from areas such as control, identification, estimation, and communication theory, computer science, artificial intelligence, and operations research.

Autonomous controllers evolve from existing controllers in a natural way fueled by actual needs, as is now discussed.

IV. Design Methodology – History

Conventional control systems are designed 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 software 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 or a region. The first mathematical model to describe plant behavior for control purposes is attributed to J.C. Maxwell 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 (the first feedback control mechanism in the historical record is the water clock of Ktesibios, 3rd century BC).




Fig 1 Ktesibios Water Clock




Fig. 2 Watt's Flyball Governor

Control theory made significant strides in the past 140 years, with the use of frequency domain methods and Laplace transforms in the 1930s and 1940s and the introduction of the state space analysis in the 1960s. Optimal control in the 1950s and 1960s, stochastic, robust and adaptive 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. The control methods and the underlying mathematical theory were

developed to meet the ever increasing control needs of our technology. The evolution in the control area was fueled by three major needs:

- a) The need to deal with increasingly complex dynamical systems.
- b) The need to accomplish increasingly demanding design requirements.
- c) The need to attain these design requirements with less precise advanced knowledge of the plant and its environment, that is, the need to control under increased uncertainty.




Fig 3 Conventional Fixed Controller for Robust Control

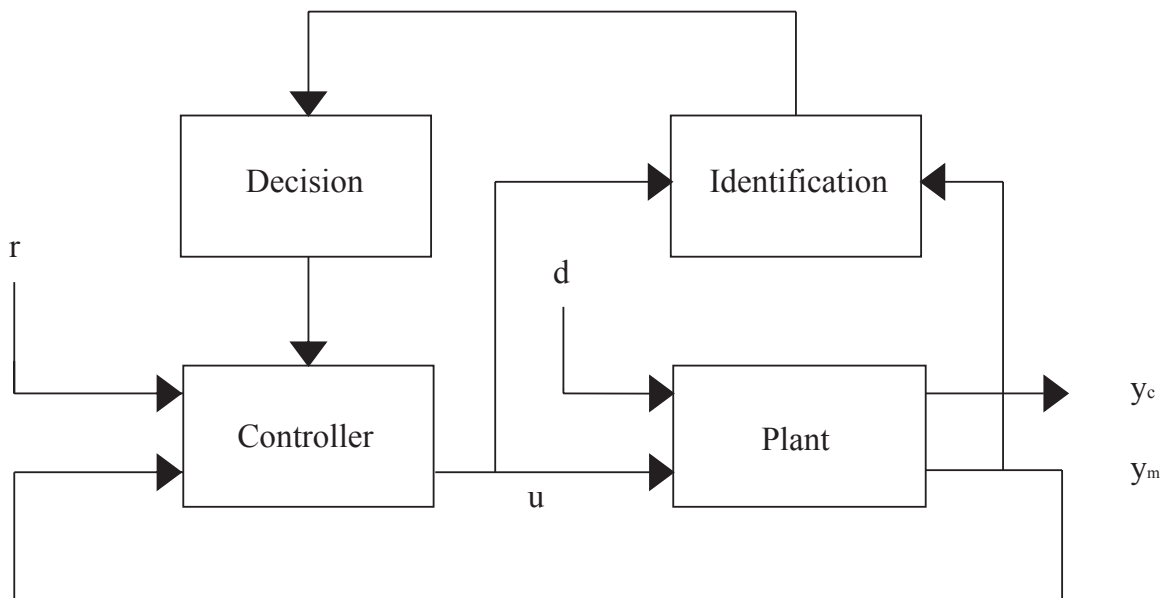


Fig 4 Conventional Indirect Adaptive Controller

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. 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. *In autonomous control systems we need to significantly increase the operating range;* 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. This will involve the use of intelligent decision making processes to generate control actions so that a performance level is maintained even though there are drastic changes in the operating conditions.




Fig 5 Highly Adaptive Controller for Autonomous Control

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 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. Many of these functions have been performed in the past by human operators. *To increase the speed of response, to relieve the pilot from mundane tasks, to protect operators from hazards, autonomy is desired.* It should be pointed out that several functions proposed in later sections, to be part of the autonomous controller, 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 reconfiguration systems in aircrafts, planning the sequence of order execution in steel mills and setting control set-points.

In the next section the functions, characteristics, and benefits of autonomous control are outlined. Next it is explained that plant complexity and design requirements dictate how sophisticated a controller must be. From this it can be seen that often it is appropriate to

use methods from operations research and computer science to achieve autonomy. An autonomous control functional architecture for future space vehicles is then presented, which incorporates the concepts and characteristics described earlier. *The controller is hierarchical, with three levels, the execution level (lowest level), the coordination level (middle level), and the management and organization level (highest level).* The general characteristics of the overall architecture, including those of the three levels are explained, and an example to illustrate their functions is given. In the following section the fundamental issues and attributes of intelligent autonomous systems are described. Then we discuss mathematical models for autonomous systems including “logical” discrete event system models. A “hybrid” approach that includes both conventional analysis techniques based on difference and differential equations, together with new techniques for the analysis of systems described with a symbolic formalism such as finite automata appears to offer advantages.

V. Functional Architecture of an Autonomous Controller

Intelligent Autonomous Control Motivation: Sophistication and Complexity in Control.

The complexity of a dynamical system model and the increasingly demanding closed loop system performance requirements, necessitate the use of more complex and sophisticated controllers. For example, highly nonlinear systems normally require the use of more complex controllers than low order linear ones when goals beyond stability are to be met. The increase in uncertainty, which corresponds to the decrease in how well the problem is structured or how well the control problem is formulated, and the necessity to allow human intervention in control, also necessitate the use of increasingly sophisticated controllers. *Controller complexity and sophistication is then directly proportional to both the complexities of the plant model and of the control design requirements.*

Based on these ideas, Saridis in [25] suggested a hierarchical ranking of increasing controller sophistication on the path to *intelligent* controls. At the lowest level, deterministic feedback control based on conventional control theory is utilized for simple linear plants. As plant complexity increases, such controllers will need for instance, state estimators. When process noise is significant, Kalman or other filters may be needed. Also, if it is required to complete a control task in minimum time or with minimum energy, optimal control techniques are utilized. When there are many quantifiable, stochastic characteristics in the plant, stochastic control theory is used. If there are significant variations of plant parameters, to the extent that linear robust control theory is inappropriate, adaptive control techniques are employed. For still more complex plants, self-organizing or learning control may be necessary. At the highest level in their hierarchical ranking, plant complexity is so high, and performance specifications so demanding, that intelligent control techniques are used. In the hierarchical ranking of increasingly sophisticated controllers described above, the decision to choose more sophisticated control techniques is made by studying the control problem using a controller of a certain complexity belonging to a certain class. When it is determined that the class of controllers being studied (e.g., adaptive controllers) is inadequate to meet the required objectives, a more sophisticated class of controllers (e.g., intelligent controllers) is chosen. That is, if it is found that certain higher level decision making processes are

needed for the adaptive controller to meet the performance requirements, then these processes can be incorporated. These intelligent autonomous controllers are the next level up in sophistication. They are *enhanced adaptive controllers*, in the sense that they can adapt to more significant global changes in the plant and its environment than conventional adaptive controllers, while meeting more stringent performance requirements. *One turns to more sophisticated controllers only if simpler ones cannot meet the required objectives. The need to use intelligent autonomous control stems from the need for an increased level of autonomous decision making abilities in achieving complex control tasks.*

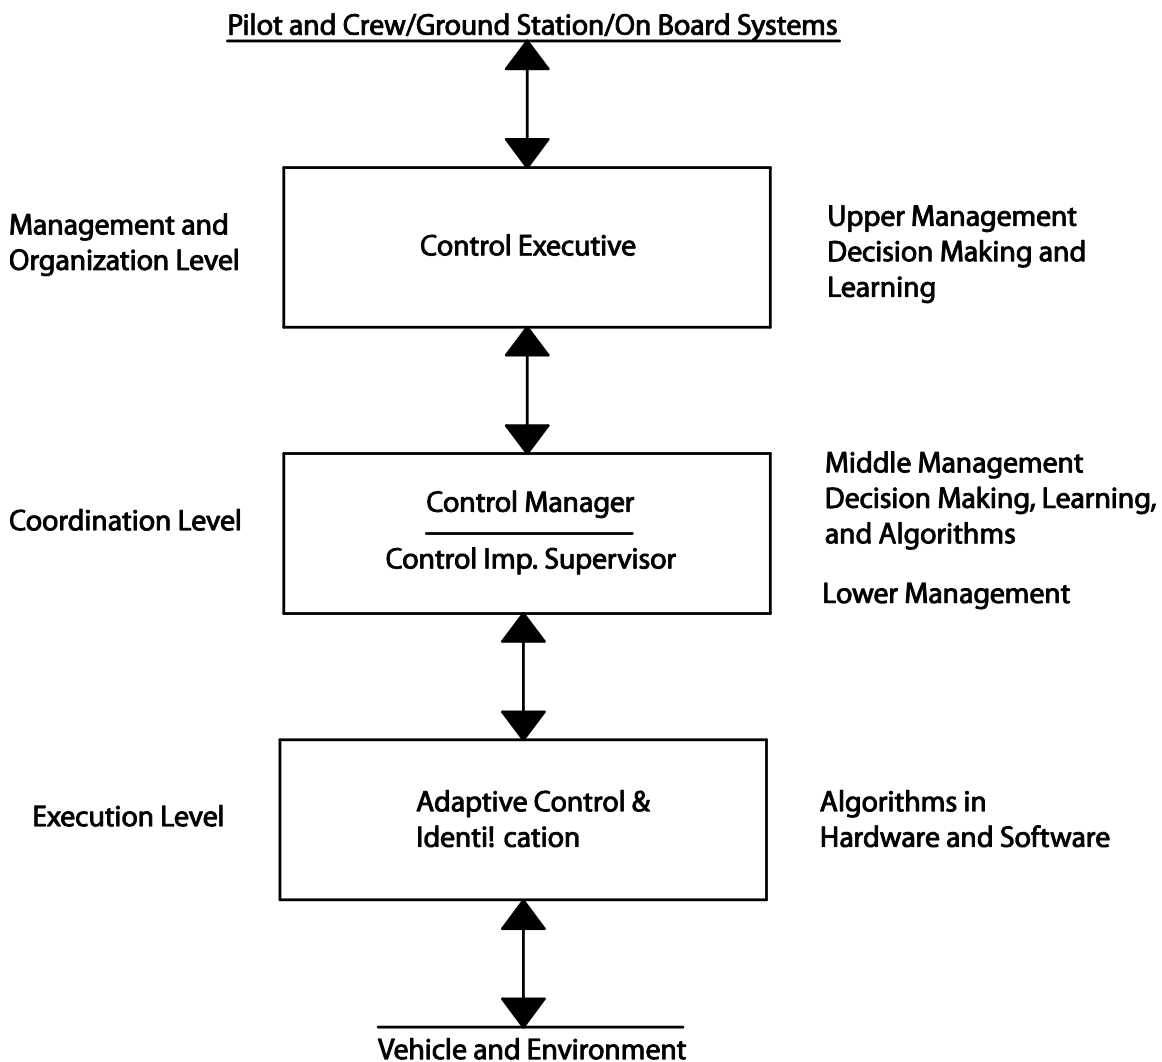


Fig 6 Autonomous Controller Functional Architecture – Spacecraft JPL

A brief literature overview of the early literature on autonomous intelligent control may be found in [1,5]. The architecture in Fig. 6 has three levels. At the lowest level, the execution level, there is the interface to the vehicle and its environment via the sensors and actuators. At the highest level, the management and 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. *It must be stressed that the system may have more or fewer than three levels. Some characteristics of the system, which dictate the 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, the hierarchical Characteristics of the plant. Note however that the three levels shown here in Fig. 6 are applicable to most architectures of autonomous controllers, by grouping together sublevels of the architecture if necessary; the levels are the lower execution level, the higher management level with everything else in between being included in the mid coordination level.* Notice that as it is indicated in the figure, the lowest, execution level involves conventional control algorithms, while the highest, management and organization level involves only higher level decision making methods. The middle, 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. They are the connection between the physical system and the controller. 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. For example, there may be one control manager, which directs a number of different adaptive control algorithms to control the flexible modes of the vehicle via appropriate sensors and actuators. Another control manager is responsible for the control functions of a robot arm for satellite repair. The control executive issues commands to the managers and coordinates their actions. Note that the autonomous controller is only one of the autonomous systems on the vehicle. It is responsible for all the functions related to the control of the physical system and allows for continuous online 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: 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.

The specific functions at each level are described in detail in [5,6]. Here we present 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 control executive, the task is refined to “repair satellite using robot A”. This is arrived at using the capability assessing, performance monitoring, and planning functions of the control executive. The control executive decides if the repair is possible under the current performance level of the system, and in view of near term planned functions. The control executive, using its planning capabilities, sends a sequence of subtasks sufficient to achieve the repair to the control manager. This sequence could be to order robot A to: “go to satellite at coordinates xyz ”, “open repair hatch”, “repair”. The control manager, 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 the control implementation supervisor, 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 m, to the right 15”, etc. These control algorithms are executed in the controller at the execution level utilizing sensor information; the control actions are implemented via the actuators.

Some Design Guidelines for Autonomous Controllers: There are certain functions, characteristics, and behaviors that autonomous systems should possess. These are outlined below. Some of the important characteristics of autonomous controllers are that they relieve humans from time consuming mundane tasks thus increasing efficiency, enhance reliability since they monitor health of the system, enhance performance, protect the system from internally induced faults, and they have consistent performance in accomplishing complex tasks. There are autonomy guidelines and goals that should be followed and sought after in the development of an autonomous system. Autonomy should reduce the work-load requirements of the operator or, in the space vehicle case discussed here, of the pilot/crew/ground station, for the performance of routine functions, since the gains due to autonomy would be superficial if the maintenance and operation of the autonomous controller taxed the operators. Autonomy should enhance the functional capability of the system. Since the autonomous controller will be performing the simpler routine tasks, persons will be able to dedicate themselves to even more complex tasks. There are certain autonomous system architectural characteristics that should be sought after in the design process. The autonomous control architecture should be amenable to evolving future needs and updates in the state of the art. The autonomous control architecture should be functionally hierarchical; for lower level subsystems to take some actions, they have to clear it with a higher level authority. The system must, however, be able to have lower level subsystems, that are monitoring and reconfiguring for failures, and act autonomously to certain extent to enhance system safety. There are also certain operational characteristics of autonomous controllers. Human operators should have ultimate supervisory override control of autonomy functions. Autonomous activities should be highly visible, “transparent”, to the operator at he maximum extent possible. Finally, there must be certain features inherent in the autonomous system design.

Autonomous design features should prevent failures that would jeopardize the overall system mission goals or safety. These features should enhance safety, and avoid false alarms and unnecessary hardware reconfiguration. This implies that the controller should have self-test capability. Autonomous design features should also be tolerant to transient errors, they should not degrade the reliability or operational lifetime of functional elements, they should include adjustable fault detection thresholds, avoid irreversible state changes, and provide protection from erroneous or invalid external commands.

VI. Characteristics of Autonomous Control Systems

Based on the architecture described above we identify the important fundamental concepts and characteristics that are needed for an autonomous control theory. Note that several of these have been discussed in the literature as outlined above. Here, these characteristics are brought together for completeness. Furthermore, the fundamental issues which must be addressed for a quantitative theory of intelligent autonomous control are introduced and 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, 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" described in [25,26]. 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*. 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. The execution level of a highly dexterous controller is very sophisticated and it can accomplish complex control tasks. The control implementation supervisor can issue high level commands to a dexterous controller, or it can completely dictate each command in a less dexterous one. The simplicity, and level of abstractness of macro commands in an autonomous controller depends on its dexterity, which really corresponds to its level of autonomy. The more able 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

and this is now discussed. It was pointed out above that complex control problems required a controller sophistication that involved the use of AI methodologies. 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”.

VII. Mathematical Models for Autonomous Systems

For autonomous control problems, normally the plant is so complex that it is either impossible or inappropriate to describe it with conventional 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 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.* For example, a simple temperature controller could ignore almost all heat related dynamics of the house or the office and consider only a temperature threshold model of the system to switch the furnace off or on. Logical discrete event system (DES) models and Petri nets are quite useful for modeling the higher level decision making processes in the autonomous controller together with logics, semantic networks, rule based descriptions etc. Queuing network models, Markov chains, etc. will be useful in the study. 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 autonomous controllers do not exist to a similar degree. This is of course because of their novelty, but for the most part, it is due to the “*hybrid*” structure (nonuniform, nonhomogeneous nature) of the dynamical systems under consideration. The systems are hybrid since in order to examine autonomy issues, a more global,

macroscopic view of a dynamical system must be taken than in conventional control theory. Modeling techniques for autonomous systems must be able to support this macroscopic view of the dynamical system, hence it is necessary to represent both numeric and symbolic information. We need modeling methods that can gather all information necessary for analysis and design. For example, we need to model the dynamical system to be controlled (e.g., a space platform), we need models of the failures that might occur in the system, of the conventional adaptive controller, and of the high level decision making processes at the management and organization level of the intelligent autonomous controller (e.g., an AI planning system performing actions that were once the responsibility of the ground station). The heterogeneous components of the autonomous 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. It is our viewpoint that research should begin by using different models for different components of the autonomous controller. Full hybrid models that can represent large portions or even the whole autonomous system should be examined but 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, research in the area of systems that are modeled with a logical DES model at the higher levels and a difference equation at the lower level should be examined. In any case, our modeling philosophy requires the examination of *hierarchical* models. Much work needs to be done on hierarchical DES modeling, analysis, and design, let alone the full study of hybrid hierarchical dynamical systems. Abstractions are of course at the center of any such study.

VIII. Concluding Remarks

The fundamental issues in autonomous control system modeling and analysis were identified and briefly discussed, thus providing an introduction to the research problems in the area. A hierarchical functional autonomous controller architecture was also presented. It was proposed to utilize a hybrid approach to modeling and analysis of autonomous systems. This will incorporate conventional control methods based on differential equations and new techniques for the analysis of systems described with a symbolic formalism. In this way, the well developed theory of conventional control can be fully utilized. It should be stressed that autonomy is the design requirement and intelligent control methods appear, at present, to offer some of the necessary tools to achieve autonomy.

IX. References

- [1] P.J.Antsaklis, K.M.Passino and S.J.Wang, "An Introduction to Autonomous Control Systems," IEEE Control Systems Magazine, Vol 11, No 4, pp 5-13, June 1991. Reprinted in Neuro-Control Systems: Theory and Applications, M.M. Gupta and D.H. Rao Eds., Chapter 4, Part 1, pp 81-89, IEEE Press 1994.
- [2] P. J. Antsaklis and J. C. Kantor, "Intelligent Control for High Autonomy Process Control Systems," Intelligent Systems in Process Engineering, pp. 37-46, AIChE Symposium Series, J.F. Davis, G. Stephanopoulos and V. Venkatasubramanian, Eds., No.

312, Vol. 92, 1996. Proceedings of the International Conference on 'Intelligent Systems in Process Engineering', Snowmass Village, Colorado, July 9-14, 1995. This is a brief version of the Technical Report of the ISIS Group, ISIS-95-004, Dept. of Electrical Engr, Univ of Notre Dame, July 1995, which was presented and distributed at the conference.

[3] P.J. Antsaklis, "Intelligent Control," Encyclopedia of Electrical and Electronics Engineering, Vol. 10, pp. 493-503, John Wiley & Sons, Inc., 1999.

[4] P. J. Antsaklis, K. M. Passino and S. J. Wang, "Autonomous Control Systems: Architecture and Fundamental Issues" (Invited), Proc. of the 1988 American Control Conference, pp. 602-607, Atlanta, Georgia, June 15-17, 1988. See also, P.J. Antsaklis and K.M. Passino, "Autonomous control systems: Architecture and concepts for future space vehicles," Final Rep., Jet Propulsion Laboratory Contract 957856, Oct. 1987.

[5] P. J. Antsaklis, K. M. Passino and S. J. Wang, "Towards Intelligent Autonomous Control Systems: Architecture and Fundamental Issues," Journal of Intelligent and Robotic Systems, Vol. 1, pp. 315-342, 1989.

[6] P.J.Antsaklis, K.M.Passino and S.J.Wang, "An Introduction to Autonomous Control Systems," Proc of the 5th IEEE International Symposium on Intelligent Control, pp.21-26, Philadelphia, Pennsylvania, Sept. 5-7, 1990.

[7] P. J. Antsaklis, "On Intelligent Control: Report of the IEEE CSS Task Force on Intelligent Control," Technical Report of the ISIS (Interdisciplinary Studies of Intelligent Systems) Group, No. ISIS-94-001, Univ of Notre Dame, January 1994.

[8] P. J. Antsaklis, "Defining Intelligent Control," Report of the Task Force on Intelligent Control, P.J Antsaklis, Chair. In IEEE Control Systems Magazine, pp. 4-5 & 58-66, June 1994. Also in Proceedings of the 1994 International Symposium on Intelligent Control, pp. (i)-(xvii), Columbus, OH, August 16-18, 1994.

[9] P.J.Antsaklis, "Towards Autonomous Intelligent Control Systems," Keynote Address. Summary in Proc of the 1991 AI Simulation and Planning in High Autonomy Systems, pp 2-3, Cocoa Beach, FL, April 1-2, 1991.

[10] P. J. Antsaklis, "On Modeling, Analysis and Design of High Autonomy Control Systems," Proc of the The 1992 Workshop on Intelligent Autonomous Control Systems, pp. 128-147, Leshem, Israel, November 2-3, 1992.

[11] P. J. Antsaklis, "Towards Autonomous Control Systems," Proc IEEE Mediterranean Symposium on New Directions in Control Theory and Applications, pp. 274-277, Chania, Crete, Greece, June 21-23, 1993.

[12] Panos Antsaklis and Kevin Passino, "Introduction to Intelligent Control Systems with High Degrees of Autonomy," in Greek, Technical Review Monthly, Athens, Greece, pp 24-35, No 26, March 1994.

[13] P. J. Antsaklis, "Intelligent Control for High Autonomy in Unmanned Underwater Vehicles," Proceedings of the NSF/ISR Workshop on 'Undersea Robotics and Intelligent Control', pp. 25-32, Lisboa, Portugal, March 2-3, 1995.

[14] P. J. Antsaklis, "Intelligent Learning Control," Guest Editor's Introduction, IEEE Control Systems, Vol.15, No.3, pp. 5-7, June 1995; Special Issue on 'Intelligence and Learning' of the IEEE Control Systems Magazine, Vol.15, No.3, pp. 5-80, June 1995. .

[15] P.J. Antsaklis, "Setting the Stage: Some Autonomous Thoughts on Autonomy," Introduction to the Panel Discussion: "Autonomy in Engineering Systems: What is it and Why is it Important?," Proceedings of the ISIC/CIRA/ISAS'98 A Joint Conference on the Science and Technology of Intelligent Systems, pp. 520-521, National Institute of Standards and Technology (NIST), Gaithersburg, MD, September 14-17, 1998.

[16] P.J. Antsaklis, "Emerging Control Technologies: What does the Future Hold?," Keynote Address Summary, Proceedings of the IFAC Workshop on Control Applications and Ergonomics in Agriculture, pp. 1-3, Athens, Greece, June 14-17, 1998.

[17] P.J. Antsaklis, "Towards High Autonomy Complex Engineering Systems," in the Panel Discussion "The Control Challenge in the 21st Century," P.P. Groumpos Organizer, Proc of the IFAC 14th World Congress, Plenary and Index Volume, pp. 83-88, Beijing, P.R. China, July 5-9, 1999.

[18] An Introduction to Intelligent and Autonomous Control, P.J.Antsaklis and K.M.Passino, Eds, 448 pages, Kluwer Academic Publishers, 1993.

[19] P. J. Antsaklis and K. M. Passino, "Introduction to Intelligent Control Systems with High Degree of Autonomy," Introduction to Intelligent and Autonomous Control, P.J.Antsaklis and K.M.Passino, Eds., Chapter 1, pp. 1-26, Kluwer,1992.

[20] K. M. Passino and P. J. Antsaklis, "Modeling and Analysis of Artificially Intelligent Planning Systems," Introduction to Intelligent and Autonomous Control, P.J.Antsaklis and K.M.Passino, Eds., Chapter 8, pp. 191-214, Kluwer, 1992.

[21] P. J. Antsaklis, M. D. Lemmon and J. A. Stiver, "Learning to be Autonomous: Intelligent Supervisory Control," Technical Report of the ISIS (Interdisciplinary Studies of Intelligent Systems) Group, No. ISIS-93-003, Univ of Notre Dame, April 1993. Also in Intelligent Control: Theory and Practice, Chapter 2, pp. 28-62, Gupta M.M., Sinha N.K., eds., IEEE Press, Piscataway, NJ, 1995.

[22] P.J. Antsaklis, "Hybrid and Supervisory Control Systems in Autonomous Underwater Vehicles," First International Workshop on Autonomous Underwater Vehicles for Shallow Waters and Coastal Environments, pp. 135-139, Lafayette, Louisiana, February 17-19, 1998.

[23] X.D. Koutsoukos and P.J. Antsaklis, "Computational Issues in Intelligent Control: Discrete-event and Hybrid Systems," in Soft Computing and Intelligent Systems: Theory and Practice, N.K. Sinha and M.M. Gupta, Eds., pp. 39-69, Academic Press, 2000. Also Technical Report isis-99-005, Dept. of Electrical Engr., Univ. of Notre Dame, May 1999.

[24] P.J. Antsaklis, "Hybrid Supervisory Control in Intelligent Autonomous Systems," *Proceedings of the 1st IFAC/IEEE Symposium on Systems Structure and Control*, Prague, Czech Republic, 29-31 August 2001.

[25] G.N. Saridis, "Toward the realization of intelligent controls," *Proc. IEEE*, Vol. 67, pp. 1115-1133, Aug. 1979.

[26] G.N. Saridis, "Analytic formulation of the principle of increasing precision with decreasing intelligence for intelligent machines," *Automatica*, V01.25, pp. 461-467, 1989.

APPENDIX A

Conventional and Autonomous Control

Comparing Conventional Controllers to Controllers with High Degree of Autonomy

The word control in "autonomous 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 autonomous 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 autonomous controllers. It is also clear that task planning is an important area in autonomous system design. So *the control problem in autonomous 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 control for high degree of autonomy is in fact interdisciplinary, and it attempts to combine and extend theories and methods from areas such as control, computer science, communications and operations research to attain demanding control goals in complex systems.

Note that the theories and methodologies from the areas of computer science communications and operations research 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; this issue is being addressed in the literature. Concepts such as reachability and deadlock developed in operations research and computer science are useful in autonomous 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 autonomous 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 autonomous 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 autonomous 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. Autonomous 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 autonomous 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 autonomous 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 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 autonomous 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 autonomous control, in addition to conventional control include hybrid systems, planning and knowledge based systems, communication protocols, security, 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.

APPENDIX B

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 Appendix we 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 was discussed 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 (*this was the comment made over 20 years ago and has proven to be correct*). 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 (*which is exactly what has happened*). 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, which are of particular interest in the control of systems [7,8]. 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 sub-goals 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 sub-goals 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.