DEFINING INTELLIGENT CONTROL

Report of the Task Force on Intelligent Control
IEEE Control Systems Society
Panos Antsaklis, Chair

December 1993

1 INTRODUCTION

In May 1993, a task force was created at the invitation of the Technical Committee on Intelligent Control of the IEEE Control Systems Society to look into the area of Intelligent Control and define what is meant by the term. Its findings are aimed mainly towards serving the needs of the Control Systems Society; hence the task force has not attempted to address the issue of intelligence in its generality, but instead has concentrated on deriving working characterizations of Intelligent Control. Many of the findings however may apply to other disciplines as well.

The charge to the task force was to characterize intelligent control systems, to be able to recognize them and distinguish them from conventional control systems; to clarify the role of control in intelligent systems; and to help identify problems where intelligent control methods appear to be the only viable avenues.

In accomplishing these goals, the emphasis was on working definitions and useful characterizations rather than aphorisms. It was accepted early on that more than one definition of intelligent systems may be necessary, depending on the view taken and the problems addressed.

In the remaining of this introduction, the different parts of this report are described and the process that led to this document is outlined. But first, a brief introduction to the types of control problems the area of intelligent control is addressing is given and the relation between conventional and intelligent control is clarified.

1.1 Conventional and Intelligent Control

The term "conventional (or traditional) control" is used here to refer to the theories and methods that were developed in the past decades to control dynamical systems, the behaviour of which is primarily described by differential and difference equations. Note that this mathematical framework may not be general enough in certain cases. In fact it is well known that there are control problems that cannot be adequately described in a differential/difference equations framework. Examples include discrete event manufacturing and communication systems, the study of which has led to the use of automata and queuing theories in the control of systems.

In the minds of many people, particularly outside the control area, the term "intelligent control" has come to mean some form of control using fuzzy and/or neural network methodologies. This perception has been reinforced by a number of articles and interviews mainly in the nonscientific literature. However intelligent control does not restrict itself only to those methodologies. In fact, according to some definitions of intelligent control (section 2) not all neural/fuzzy controllers would be considered intelligent. The fact is that there are problems of control which cannot be formulated and studied in the conventional differential/difference equation mathematical framework. To address these problems in a systematic way, a number of methods have been developed that are collectively known as intelligent control methodologies.

There are significant differences between conventional and intelligent control and some of them are described below. Certain of the issues brought forward in this introduction are discussed in more detail in section 3 of this report. 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. Intelligent control attempts to build upon and enhance the conventional control methodologies to solve new challenging control problems.

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, that study the control of continuous-state dynamic processes by discrete-state sequential machines. 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 nondynamic 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 and control theorists are invited to address these problems. 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. These issues are discussed in more detail in section 3.

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 recently 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 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.

Research areas relevant to intelligent control, in addition to conventional control include areas such as planning, learning, search algorithms, hybrid systems, fault diagnosis and reconfiguration, 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 now it is clear that computational complexity is a central issue, whenever one attempts to control complex systems.

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 report we do not even attempt to give general definitions of intelligence. Instead we introduce and discuss several characterizations of intelligent systems that appear to be useful when attempting to address some of the complex control problems mentioned above.

Some comments on the term "intelligent control" are now in order. Intelligent controllers are envisioned emulating 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. Certainly the term intelligent control has been abused and misused in recent years by some, and this is of course unfortunate. Note however that this is not the first time, nor the last that terminology is used to serve one's purpose. Intelligent control is certainly a catchy term and it is used (and misused) with the same or greater abundance by some, as for example the term optimal has been used (or misused) by others; of course some of the most serious offenses involve the word "democracy"! For better or worse, the term intelligent control is used by many. An alternative term 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; autonomous control is also discussed in sections 2 and 3. On the other hand, "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. This is the true challenge.

I would like to finish this brief outline with an optimistic note; and there are many reasons for being optimistic. This is an excellent time indeed to be in the control area. We are currently expanding our horizons, we are setting ambitious goals, opening new vistas, introducing new challenges and we are having a glimpse of the future that looks exciting and very promising.

1.2 Points of View

The list of the task force members can be found at the end of this report. This report represents a collective view of what intelligent control is and what are its main characteristics or dimensions. As usually happens, some of the members have had greater input to the process than others. Independently of the amount of individual contributions, however, it is fair to say that no member of the committee objects to the main points made in this report. In addition, in the second part of this report in section 3, task force members further explain and give reference to their own points of view and this gives an opportunity for further reading into the subject. Some additional references are also given.

1.3 The Process

Before I outline the different parts of this report, let me say a few words about the procedure that led to its final version. After the task force was formed in May, a position paper representing a particular point of view was aired to "get the ball rolling". It certainly achieved that! Views were exchanged over email and animated discussions were conducted off and on during the whole summer. A first outline of this report was sent to all members in late July. It tried to capture the main points of view and to establish a desirable format for the report. At the end of August a meeting took place at the 1993 International Symposium on Intelligent Control in Chicago, and several task force members and non- members exchanged views on the subject. It became apparent at that meeting that consensus was emerging. Participants of that meeting sent their comments in writing to all the task force members in September; a draft of this final report was put together in October, with the final version being prepared in November and December 1993.

1.4 Report Outline

This report consists of two main parts. The first part, in section 2, has the form of an executive summary and the second part in section 3 contains additional material and some references. Specifically, in section 2 definitions of intelligent systems and of degrees of intelligence are given, and the role of control in intelligent systems is explained. The different characteristics or dimensions of intelligent systems such as autonomy, learning and hierarchies are then discussed. Section 3 contains edited versions of some of the email exchanges and additional comments by the task force members, together with some references for further reading. They were included in an attempt to further clarify the issues brought forward in the first part of this report. They are meant to supplement the material in section 2 and to provide some guidance and references in exploring the area of Intelligent Control.

As the chair of the Task Force on Intelligent Control I had the role of coordinating the discussions and exchanges of the different points of view. I also drafted this report, which was then finalized with the help of the members of the task force, whom I would like to thank for their contributions and insights. I used my own judgement in selecting the format, the particular form of definitions, and in emphasizing particular aspects and characteristics of intelligent systems; and any errors are entirely mine. My aim was to extract the main points out of lengthy email exchanges and to write a report that represents the collective view of the Task Force on Intelligent Control. I hope that this report will be useful to the members of the Control Systems Society, that it will help identify and clarify the main issues in the area of Intelligent Control Systems, and will provide information and incentives for further study.

Submitted by
Panos J. Antsaklis
Chair, Task Force on Intelligent Control
IEEE Control Systems Society

2 INTELLIGENT CONTROL AND ITS DIMENSIONS

Intelligence and intelligent systems can be characterized in a number of ways and along a number of dimensions. There are certain attributes of intelligent systems, common in many definitions, that are of particular interest to the control community. These are emphasized in this report.

In the following, several alternative definitions and certain essential characteristics of intelligent systems are first discussed. A brief working definition of intelligent systems that captures their common characteristics is then presented. In more detail, we start with a rather general definition of intelligent systems, we discuss levels of intelligence, we explain the role of control in intelligent systems and outline several alternative definitions. We then discuss adaptation and learning, autonomy and the necessity for efficient computational structures in intelligent systems, to deal with complexity. We conclude with a brief working characterization of intelligent (control) systems, some examples and a list of important future research directions.

2.1 Intelligent Systems

We start 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 and in the following we mention several. They provide alternative, but related characterizations of intelligent systems with emphasis on systems with high degrees of intelligence.

The following definition emphasizes the fact that the system in question processes information, and it focuses on man-made systems and intelligent machines:

A. 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 leads to the principle of increasing precision with decreasing intelligence, which claims that: applying machine intelligence to a data base generates a flow of knowledge, lending an analytic form to facilitate modeling of the process.

Next, an intelligent system is characterized by its ability to dynamically assign subgoals and control actions in an internal or autonomous fashion:
B. 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:
C. 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.

2.2 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):
D. 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.

2.3 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 be "sparse" and it 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 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 Summary-A Working Definition

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 sparse types of functional architectures such as hierarchies.

2.4 Some Examples

Below, a list of man-made systems that solve complex problems and incorporate some of the above essential characteristics of intelligent control systems is given. The intention, in including such list, is to point out the fact that such systems do exist. Note that the list is far from complete, and it only contains the cases brought forward by task force members.

An example of a Hierarchically Intelligent Control System was designed and built at the NASA CIRSSE/RPI labs, to do truss construction remotely in deep space for the NASA Space Station "Freedom". The coordination and Execution levels were built using Petri nets, sensing (VSS) and motion control (CTOS) respectively. The innovation of the project was that a system (CTOS), 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. Thus the system was very efficient requiring a narrow bandwidth communication line. The system was tested by controlling a truss assembly at RPI, from NASA Johnson in Houston through a telephone line. The Organization level was replaced by a human manager; the design was completed using a Boltzmann machine Neural net, but was never built. An Intelligent controller for a mobile robot was also planned but never built at CIRSSE/RPI.

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; Robot Cleaning and Deburring Workstation; Robot Deburring and Chamfering Workstation; Multiple Autonomous Undersea Vehicles; NASA Space Station Telerobotics (NASREM); Army Field Material Handling Robot; DARPA Submarine Automation (SOAS); BOM Coal Mine Automation; Army Unmanned Land Vehicles: TEAM vehicle project, TMAP vehicle project. Robotics Testbed project, RT Demo I testbed; Air Force Next Generation Controller (NGC); NCMS Next Generation Inspection System (NGIS); DOT Intelligent Highway Vehicle Vision based road following; NIST RoboCrane; Navy/NIST/ARPA Enhanced Machine Controller.

Other examples include mobile robots that exhibit some autonomy at Oak Ridge National Lab, Robotic Division; an intelligent controller for OSPREY machine installed at navy research center developed at Drexel University; autonomous robots at Georgia Tech.

2.5 Future Research Directions

A list of important and promising research topics in intelligent control is given below. Although the list may not be complete, it includes some of the directions along which the field ought to be making progress in the next few years.

1. Mathematical modeling and analysis of intelligent control systems; in both discrete event and hybrid frameworks. Model identification; adaptive methods to derive higher level, more abstract models.

2. Fault detection and identification, control reconfiguration; also alarms and health monitoring.

3. Planning and learning control systems.
4. Efficient computational frameworks and algorithms to deal with complexity.
5. Emphasis on applications and on integrated intelligent control systems; important automotive, manufacturing and aerospace applications.

In section 3, the issues brought forward in this section are further discussed.

3 POINTS OF VIEW OF INTELLIGENT CONTROL

This section consists of additional material that helps clarify the issues addressed in the previous section and includes references for further reading. This material was contributed by the task force members, all recognized for their contributions in the area of intelligent control.

3.1 On Intelligence and its Dimensions

by J.S. Albus

A definition of intelligence is first given and then the dimensions of intelligence are discussed; see [1] for further discussion.

Definition of Intelligence

In order to be useful in the quest for a general theory, the definition of intelligence must not be limited to behavior that is not understood. A useful definition of intelligence should span a wide range of capabilities, from those which are well understood, to those which are beyond comprehension. It should include both biological and machine embodiments, and these should span an intellectual range from that of an insect to that of an Einstein, from that of a thermostat to that of the most sophisticated computer system that could ever be built. The definition of intelligence should, for example, include the ability of a robot to spotweld an automobile body, the ability of a bee to navigate in a field of wild flowers, a squirrel to jump from limb to limb, a duck to land in a high wind, and a swallow to work a field of insects. It should include what enables a pair of blue jays to battle in the branches for a nesting site, a pride of lions to pull down a wildebeest, a flock of geese to migrate south in the winter. It should include what enables a human to bake a cake, play the violin, read a book, write a poem, fight a war, or invent a computer.

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, prosper, and reproduce in a complex and often hostile environment.

From the viewpoint of control theory, intelligence might be defined as a knowledgeable "helmsman of behavior". Intelligence is the integration of knowledge and feedback into a sensory-interactive goal-directed control system that can make plans, and generate effective, purposeful action directed toward achieving them.

From the viewpoint of psychology, intelligence might be defined as a behavioral strategy that gives each individual a means for maximizing the likelihood of propagating its own genes. Intelligence is the integration of perception, reason, emotion, and behavior in a sensing, perceiving, knowing, caring, planning, acting system that can succeed in achieving its goals in the world.

For the purposes of this paper [1], intelligence will be defined as the ability of a system to act appropriately in an uncertain environment, where 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.

Both the criteria of success and the system's ultimate goal are defined external to the intelligent system. For an intelligent machine system, the goals and success criteria are typically defined by designers, programmers, and operators. For intelligent biological creatures, the ultimate goal is gene propagation, and success criteria are defined by the processes of natural selection.

There are degrees, or levels, of intelligence, and these are determined by:

1. the computational power of the system's brain (or computer),

2. the sophistication of algorithms the system uses for sensory processing, world modeling, behavior generating, value judgment, and global communication, and

3. the information and values the system has stored in its memory.

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. In artificial systems, growth in computational power and accumulation of knowledge derives mostly from human hardware engineers and software programmers. In natural systems, intelligence grows, over the lifetime of an individual, through maturation and learning; and over intervals spanning generations, through evolution.

Note that learning is not required in order to be intelligent, only to become more intelligent as a result of experience. Learning is defined as consolidating short-term memory into long-term memory, and exhibiting altered behavior because of what was remembered. In [1], learning is discussed as a mechanism for storing knowledge about the external world, and for acquiring skills and knowledge of how to act. It is, however, assumed that many creatures can exhibit intelligent behavior using instinct, without having learned anything.

Dimensions of Intelligence

The dimensions of intelligence may be thought of as elements in an intelligence-vector (or IQ vector) defined by parameters such as:

Computing power, number of processors, interprocess communications; Memory size, storage and retrieval functions; Knowledge representation mechanisms, including: Maps, Symbols, Attribute-value pairs, States and state- variables; Knowledge presentation systems such as: Query-reply, Question-answering, List searching; Functional capabilities such as: Motor skills, Perceptual skills, Reasoning and problem solving, Value judgment functions; Sensory resolution and range in terms of: Number and resolution of pixels (vision, touch, hearing), Spectral range and resolution, Temporal range and resolution (hearing, speech); Sensory processing: Signals to symbols, Detection and recognition, Recursive estimation, Haptic perception, Uncertainty and probability; Planning and predictive capabilities such as the ability to: Predict the results of actions, Predict actions of the world, Predict actions of other agents; Value judgment capabilities: Compute cost, risk, and benefits, Evaluate observed events, objects, and situations, Evaluate predicted outcomes, Generate rewards and punishments for learning, Assign priorities to behavioral tasks; Learning capabilities such as the abilities to: Remember objects, experiences, stories, symbols, Learn skills and tasks, Learn from experience, Learn from a teacher, Learn from symbolic text;

Along each of these dimensions, there are degrees or levels of capability. These dimensions define a space of intelligent systems, and the intellectual capabilities (or IQ) of any particular system at any particular time can be represented as a point (or vector) in that space. The origin of this space corresponds to the set of systems that have zero level of capability along all dimensions. Thus the origin of the space of intelligent systems consists of a point representing the set of non-intelligent systems.

The point in IQ space thus moves as the intelligence of the system grows or changes (possibly through learning or forgetting, or through acquiring new skills or losing skills).

[1]. Albus J.S., "Outline for a Theory of Intelligence", IEEE Transactions on Systems, Man and Cybernetics, Vol. 21, No.3, May/June 1991.
[2]. Albus J.S., "A Reference Model Architecture for Intelligent Systems Design", in Antsaklis P.J., Passino K.M., eds., An Introduction to Intelligent and Autonomous Control, Kluwer Academic Publishers, Norwell, MA, 1993.

3.2 On Autonomy and Intelligence in Control

by P.J. Antsaklis

In the design of controllers for complex dynamical systems there are needs today that cannot be successfully addressed by 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. More detailed treatment of the issues brought forward in the following can be found in [1], [2] and [3].

The need for quantitative methods to model and analyze the dynamical behavior of such autonomous systems presents significant challenges well beyond current capabilities. 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.

Conventional Control - Evolution

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 the 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. 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 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.

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 signalled the end of the era of intuitive invention. The flyball governor worked adequately for a long time meeting the needs of the period. As time progressed and more demands were put on the device there came a point when better and deeper understanding of the device was necessary, as it started exhibiting some undesirable and unexplained behavior, in particular unstable oscillations. This is quite typical of the situation in man made systems even today. Similarly to the flyball governor, one can rely on systems developed based mainly on intuitive invention so much. To be able to control highly complex and uncertain systems we need deeper understanding of the processes involved and systematic design methods, we need quantitative models and design techniques.

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 in the plant and the environment. 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.

At this point the seminal contributions of Norbert Wiener, the father of Cybernetics, to human-machine interaction should be mentioned. Note that many of the ideas in intelligent control have been influenced by past theories and methods. What is different now is that much faster, and better understood, machines are available today than ever before. So the dreams of yesterday may become reality in the not so distant future.

Intelligent Control for High Autonomy Systems

There are cases where 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. 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. It is useful to keep in mind an example which we may call the Houston control example . It is an example that sets goals for the future and it also teaches humility as it indicates how difficult, demanding and complex autonomous systems can be. Currently, if there is an unanticipated event on the space shuttle, such as a malfunction or a set of new tasks to be accomplished, the problem is addressed by the large number of engineers working in Houston Control, the ground station. After the problem is solved on the ground, the specific detailed instructions about how to deal with the problem are sent to the shuttle. Imagine the time when we will need all the tools and expertise of all Houston Control engineers, that are related to specific problems, aboard the space vehicle, or the space shuttle, for extended space travel. This is certainly not an easy problem! What is certainly possible in the near future is to incorporate some of this knowledge in the onboard computers to achieve higher degrees of autonomy in achieving and setting goals than it is the practice today, thus reducing the dependence on the ground stations and on communication links.

In view of the above it is quite clear that in the control of complex systems, there are requirements today that cannot be successfully addressed with the existing conventional control theory. They mainly pertain to the area of uncertainty, present because of poor models due to lack of knowledge, or due to high level models used to avoid excessive computational complexity. Normally the plant is 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 systems 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. Depending on the control performance specifications, 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 dynamics of the house or the office and consider only a temperature threshold model of the system to switch the furnace off or on. This naturally leads to the study of hybrid control systems, which are continuous-state systems controlled by sequential machines [3].

A number of research areas important to intelligent autonomous systems may be identified. They include the areas of: Hybrid Systems, Discrete Event Systems Theory and Simulation, Restructurable Control, Failure Detection and Identification (FDI), Intelligent Systems, Hierarchical Systems, Planning and Expert Systems, Machine Learning, Fuzzy Control, and Neural Networks.

Intelligent Autonomous 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 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 are not appropriate to address certain questions of importance to control systems such as stability. It is not that quantitative methods developed in other fields are inferior, it is the fact that these methods were developed to answer different questions. 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 that combine systems of continuous and discrete states [3]. 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 the control system ideas, such as feedback, and methods that are based on rigorous mathematical framework can provide the base for new theories and methods in those areas.

Hybrid System Modeling and Design

Being able to control a continuous-state system using a discrete-state supervisory controller is a central problem in the highly autonomous control of physical systems. The theory of hybrid system modeling and control addresses some of the important issues of extracting higher level abstract models from more detailed ones [3].

Concluding Remarks

Computational complexity is a major issue as the systems studied are typically very complex. Reduced computational complexity may mean that the controller can be implemented in real time. Without attempting to address the computational complexity issue it is impossible to achieve the levels of autonomy envisioned. Systematically deriving more abstract models so that only the necessary information is dealt with is essential; it is as essential or more than designing faster computers.

Incorporating dedicated sensors and actuators to identify changes and reconfigure the control laws may be necessary in high autonomy systems; this may be necessary for example in satellites to accomplish failure diagnosis. Technological breakthroughs are making large numbers of distributed sensors and actuators possible. This will certainly make reconfiguration and higher autonomy more common place. Areas such as sensor data fusion are becoming more important so to be able to deal with the mass of available data. And methods to extract only the necessary information from the data, which is related to the problem of extracting more abstract models, are becoming essential in the quest for higher autonomy.

Improving existing control systems by adding on new features is a plausible approach having high chances for success. This bottom-up approach builds upon experience and uses existing knowledge. It is also easier to justify in applications, where system failure is costly in human and material sense.

In summary, conventional control methods need to be enhanced, so that control systems can be designed that cope with significant changes in the plant, environment and objectives. Note that the goal is control systems with higher degree of autonomy in achieving and even setting control goals. It is stressed that autonomy is the design requirement and intelligent methods appear to offer some of the necessary tools to achieve higher degrees of autonomy. The research area of intelligent autonomous systems is a research area in its own right. It uses methods from a variety of areas but it modifies and extends them to address the particular problems of interest. There is need to answer questions and resolve novel problems in Planning and Expert Systems, in Learning and Neural Control, in Discrete Event Dynamical and Hybrid Systems, in Reconfigurable Control Systems and FDI Systems to mention but a few. There is great need for quantitative methods and mathematical rigor in the area; there is need for systematically generating less detailed, more abstract models. On going research in hybrid systems is attempting to address some of these problems.

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

[2]. Antsaklis P.J.and Passino K.M.,"Introduction to Intelligent Control Systems with High Degrees of Autonomy", in An Introduction to Intelligent and Autonomous Control, Antsaklis P.J., Passino K.M., eds., Kluwer Academic Publishers, Norwell, MA, 1993.

[3]. Antsaklis P. J., Stiver J. A. and Lemmon M. D., "Hybrid System Modeling and Autonomous Control Systems", Hybrid Systems, R L Grossman, A Nerode, A P Ravn, H Rischel Eds, pp 366-392, Lecture Notes in Computer Science, LNCS 736, Springer-Verlag, 1993.

3.3 On Intelligence and Learning

by M.D. Lemmon

A supervisory control system uses discrete event systems to control the plant. Such control systems can often be referred to as "intelligent" control systems because the actions of the controller attempt to mimic high level decision making processes of human operators. This notion of machine intelligence, however, is not entirely satisfying. At issue is the notion that mimicry of human decision making constitutes intelligence. The traditional formulation of such controllers involve the assignment of interpretations to logical symbols. Such interpretations allow us to "explain" what the controller is attempting to do. In a temperature control system, for instance, a certain range of temperatures might be designated as "TOO HOT", thereby necessitating a control action to cool the system. The "intelligence" of the system is buried in its interpretation of that symbol "TOO HOT". But where does this interpretation originate? In general, it is the designer who provides symbol interpretations. This means that it is not the system, but rather the system designer who is intelligent. Therefore if we are to have an "intelligent" control system, the system must have a capability for assigning symbol interpretations in an autonomous manner. This capability can be referred to as "symbol binding". The degree to which these associations can be done autonomously represents one way of quantifying the system's intelligence.

There are a number of consequences to this view of machine intelligence. 1). A desirable property of intelligent systems is that they are "adaptive". The ability to adaptively bind symbols with respect to an underlying organizational principle means that the system "understands" the meaning or significance of that organizational principle. 2). Intelligence is an internal property of the system. It is not a behavior. The immediate consequence of this observation is that a system's intelligence cannot always be determined by passive observation of behavior. Intelligence must be determined by actively testing to see whether or not the system is adaptively binding symbols with respect to an internally realized performance principle. 3). A pragmatic reason for focusing on "intelligent" control systems is that they endow the controlled system with enhanced autonomy. Examining the anticipated applications of intelligent control, it is apparent that they are meant for complex and unpredictable systems. This means that the system may change so that the original symbol bindings may no longer represent a valid interpretation of the system's symbolic behavior. If this occurs, then it is well within the realm of possibility for our controller to happily "chunk" away and produce a stream of nonsensical control directives. The reason this occurs is because the system has no "understanding" of the significance or meaning of the symbols it is manipulating. The result of this situation is a system whose autonomy is circumscribed by an a priori and possibly ad hoc set of symbol bindings. The pragmatic solution is to allow the system to adaptively fix bindings with respect to an internal organizational or performance principle. This is precisely what we should expect of an "intelligent" control system.

The preceding discussion has introduced a perspective on intelligent control which focuses on the way in which a system determines the interpretation of control directives or policies. It was argued that a desirable property of intelligent control systems is that they bind symbol interpretations using an internal representation of the plant's underlying control objective. In this regard, intelligent control can be viewed as the ability of a system to autonomously organize its controller to achieve a well-defined objective. Autonomy becomes an important attribute of intelligent control in which the degree of autonomy quantifies one aspect of the system's intelligence.

[1] Lemmon M.D. and Antsaklis P.J., "Towards a Working Characterization of Intelligent Supervisory Control", Technical Report of the ISIS Group (Interdisciplinary Studies in Intelligent Systems), University of Notre Dame, ISIS-93- 008, Notre Dame, IN, November, 1993.

[2]. Lemmon M. D., Stiver J. A. and Antsaklis P. J., "Event Identification and Intelligent Hybrid Control", Hybrid Systems, R L Grossman, A Nerode, A P Ravn, H Rischel Eds, pp 269-296, Lecture Notes in Computer Science, LNCS 736, Springer-Verlag, 1993.

[3]. Antsaklis P.J., Lemmon M. D. and Stiver J. A., "Learning to be Autonomous: Intelligent Supervisory Control", Technical Report of the ISIS Group (Interdisciplinary Studies in Intelligent Systems), No. ISIS-93-003, Univ of Notre Dame, April 1993. Also in Intelligent Control: Theory and Practice, Gupta M.M., Sinha N.K., eds., IEEE Press, Piscataway, NJ, 1994.

3.4 On Intelligent Control, Learning and Hierarchies

by A. Meystel

Why intelligent control? The need to deal with problems of uncertainty is not a new one. However, only during the last decades, the significant developments in the computer area enabled new approaches to these problems: approaches employing cognitive properties of intelligence including generalization, focusing attention, and combinatorial search and others considered to be properties of human intelligence. Note that all properties of intelligence including learning, recognition, existence of resolution levels, can be reduced to existence of generalization, focusing attention, and combinatorial search. Intelligent control focuses upon problems that otherwise cannot be solved, or can be solved in a unsatisfactory way.

1. Control: Control is to direct a system to a preassigned goal or to maximize a preassigned measure of utility under a set of specifications.

This directing can be done both in an open-loop as well as in a closed-loop fashion. Open loop control presumes existence of a model of the system. The open-loop control assignment invokes the process of "plan" generation ("planning") performed e.g. by searching. Since the model is usually incomplete and/or inadequate, the closed loop controller is required for error compensation which uses a feedback. Thus, Definition 1 presumes existence of a goal, a model, a plan, or a feedforward control law, and a feedback control law - all determined for a particular resolution of the control level.

2. Resolution (Scale, Granularity, Accuracy, Discretization): Resolution of the control level is the size of the indistinguishability zone (tile) for the representation of goal, model, plan and feedback law. Any control solution alludes to the idea of resolution (scale, granularity, accuracy, distinguishability zone, discrete) explicitly, or implicitly.

It turns out that resolution directly determines the complexity of computations. In complex systems and situations one level of resolution is not sufficient because the total space of interest is usually large, and the final accuracy is usually high enough. So, if the total space of interest is represented with the highest accuracy, the e (epsilon)-entropy of the system (the measure of its complexity) is very high. E (epsilon)- entropy=log(total volume of space/elementary discrete of space).

The total space of interest is to be considered initially with a much lower resolution. Only a subset of interest is considered at a higher resolution. The subset of this subset is considered with even higher resolution, and so on, until the highest resolution is achieved. This consecutive focusing of attention results in a multilevel task decomposition, and finding the intermediate (nested) plans at several resolution levels of the multiresolutional system.

We should start talking about complexity of the intelligent controller explicitly, remember that intelligence is a tool of fighting complexity, remember that this is why the level of resolution emerge.

3. Multiresolutional (Multiscale, Multigranular) representation system: Multiresolutional system is defined as a data (knowledge) structure for representing the model of our system at several resolution levels-scales. (A terminological comment: instead of the word "multiresolutional system", a word "heterarchy" can be used which is understood as follows: heterarchy - is a hierarchical organization of a het-

erogeneous information (knowledge). "Hierarchy" is a more general term, it can be related both to "homogeneous" and "heterogeneous" representations. We should not abstain from using these terms if they are clearly defined. Multiscale system seems to be a good term too).

In order to construct a multiresolutional (multiscale) system of representation,the process of generalization is consecutively applied to the representation of the higher levels of resolution. Generalization usually presumes clustering of the subsets and substitution of them by entities of the higher level of abstraction. This is why instead of the term "resolution levels" we use sometimes an expression "abstraction levels" (which is the same as "generalization levels", "granularity levels", etc.).

4. Intelligent Control: Intelligent Control is a computationally efficient procedure of directing to a goal of a complex system with incomplete and inadequate representation and under incomplete specifications of how to do this (acting appropriately in an uncertain environment).((We can talk about the degree of completeness of the system representation and the specifications formulation: then all levels of intelligence will be presented)) Sometimes, Intelligent Control presumes working under not completely specified problem which requires further clarification during the functioning of the system. Intelligent control as a rule combines planning with on-line error compensation; it requires learning of both the system and the environment to be a part of the control process. Most importantly: intelligent control usually employs generalization (G), focusing attention (FA), and combinatorial search (CS) as its primary operators (GFACS) which leads to multiscale structures.

In all intelligent controllers, one can easily demonstrate presence of the GFACS operators. It is also possible to demonstrate that using the set of GFACS operators is not typical for conventional controllers, although the elements of GFACS are often utilized. The following attributes of Intelligent Control are presumed: multiresolutional (multiscale) system of goals, multiresolutional (multiscale) system of model representation,multiresolutional (multiscale) system of plans, and multiresolutional (multiscale) system of feedback control laws.

(A terminological comment: fuzzy logic controllers are tools of generalization and focusing attention; neural networks are tools of generalization , focusing attention, and combinatorial search; combinatorial search has many particular instantiations: A-star, exhaustive search, complete, or approximate dynamic programming, etc.)

5. Intelligence

Intelligence is a control tool.(for the system at hand) that has emerged as a result of evolution. Intelligence is oriented toward complexity reduction. Intelligence allows for an increase in functionality with a reduction of computational complexity.

Intelligence grows through generation of multiresolutional (multigranular, multiscale) system of knowledge processing. Multiresolutional system of knowledge processing is not hardwired. These multi-level systems are not 'hardwired' hierarchies (although they can be in some cases); they are rather virtual hierarchies of perception, representation of the World Model, i.e. knowledge representation, and representations of decisions. As the new concept emerges - a new 'node'is created. Intelligence can be evaluated by a"degree of intelligence". The definition of intelligent control should be based on the properties of intelligence as we understand them rather than the virtue of using some particular hardware components. The following properties can be used for evaluating the degree of intelligence a) Error compensation - low level of intelligence (Level 1), b) Planning+error compensation within the vocabulary of the designer - medium level of intelligence (Level 2), c) Planning+error compensation with creation of new alternatives not introduced previously by the designer - high level of intelligence (Lever 3), d) Reformulation of assignment in changing situation-very high level of intelligence (Level 4).

6. Learning.

Maintenance of the multiresolutional system of representation is done by learning which employs the same set of GFACS operators. Levels of resolution are selected to minimize the complexity of computations. Planning and determining of the feedback control laws is also done by joint using of generalization, focusing of attention, and combinatorial search (GFACS). When generalization is continuously done in the course of time-varying of variables, it becomes a key tool of learning.

The process of generalization upon the time-varying functions is called learning of a control system. It results in constant updating of the multiresolutional system of representation, and thus, in improvement of plans and feedback control laws. Learning is a component of this multiresolutional knowledge processing.

The operation of learning was associated with the layers: each layer is learning separately: all learning processes (at particular levels) are connected via their results. Learning experiences can be organized ONLY by using a multiresolutional structure! (This is how it is done in the neural nets too.) Levels are not hardwired,they are constructed from the information at hand.

7.Nesting

Nesting is a property of recursively applying the same procedures of multiresolutional knowledge processing within the operator of processing at a level. Levels of the multiresolutional intelligent controller are nested one within another. The levels function as separate controllers.

a. This separation in levels is a result of a need to reduce the complexity of computations. Thus, instead of solving in one shot the whole problem with the maximum volume of the state space and with the amount of high resolution details one chose to solve several substantially simpler problems nested one within another. b. Assigning of resolution levels is probably the most urgent problem of the area of intelligent control: it should be done so to minimize the complexity of the controller. However, each level of control has its perception, its world model, and its decision making. Perceptions of all levels are nested one within another, world models are nested one within another, decisions are nested one within another too. These system would be impossible without generalization, focusing attention and combinatorial search. c. All these level controllers have limited resolution and they can be defined as fuzzy controllers. All these controllers are part of the overall learning process and cannot function unless an NN-like structure, recognizes of motion and primitives of the world that can be correlated with each other.

8. Additional relevant issues:

8.1 Why has the property of intelligence emerged in the living creatures? In the evolving Nature, the evolution of intelligence can be demonstrated as a tool of survival that emerged in order to control the the systems in better correspondence with the ever changing environment conditions and with the evolving needs and 'desires'. As the complexity of needs is growing, fighting this complexity is becoming a major role of intelligence.

8.2. Increasing functionality with reducing computational complexity-a fundamental result of the evolution of intelligence. This evolution of intelligence can be unequivocally interpreted as development of a system (together with its controller) which allows for increasing the functionality of the system while fighting with computational complexity. This is why the ability to 'generalize' emerges ('lumping' for better storage and quicker computation). Generalization is a tool of creating a new, abridged representation. The new representation 'in generalities', forms the level of lower resolution . Since at lower resolution we can afford the larger scope of attention - we can solve the problem of a larger picture. So, the decision of the required resolution can be preceded by the decision at lower resolution, and so on.

8.3. Why hierarchies? We can call it "a hierarchy", or we can use another term, but we cannot avoid labeling the structure of intelligent controller that by and large boils down to a hierarchy. About 'hierarchy': GFACS recursively constructs levels of representation, and levels of decision making, obviously supplemented by levels of perception. So, the system which employs GFACS as an elementary computational package builds itself as a system of multiple levels of representation. As a result of this consecutive levels construction it arrives with a low resolution level which contains maximum of what the system knows - in a compact, generalized, aggregated form. The next level is dealing with a subset of this low resolution picture - but with more details, i.e. at a higher resolution (and so on recursively). For the whole process of decision making several resolution levels are required. Each of these levels executes the same chain: perception-knowledge processing- decision making. Should we call this multilayer system a hierarchy? I think, the term to be chosen is a secondary issue. At least for the knowledge representation system it can be considered as such. Object-orientedness emerges when we are dealing with entities. Multiresolutional 'nestedness' is obvious when no entities are listed and we describe textures. Is it a single-principle hierarchy? Not at all. It can be - for a simple case. In general case it is a mixture of hierarchies based on many principles - a heterarchy.

But no matter what term we will agree upon: hierarchy, heterarchy, hierarchical network, multiresolutional hierarchy, multigranular network - it is a layered system of multiresolutional representation with decision making processes performed at each level. Representations are nested and decomposable. Decisions are nested and decomposable too. Processes of the higher resolution can be guided by processes of the lower resolution.

[1]. Meystel A., "Nested Hierarchical Control ", in Antsaklis P.J., Passino K.M., eds., An Introduction to Intelligent and Autonomous Control, Kluwer Academic Publishers, Norwell, MA, 1993.
[2]. Meystel A., Autonomous Mobile Robots, World Scientific, 1991.
[3]. Meystel A., "Intelligent Control ", in Encyclopedia of Physics and Technology, Academic Press, 1993.

3.5 On the Relevance of Control Engineering

by K.M. Passino

In this section we explain the control engineer's perspective on intelligent control systems. Let us begin by defining a "control methodology" to be the set of techniques and procedures used to construct and/or implement a controller for a dynamical system. For many intelligent control systems (e.g., fuzzy/neural controllers, expert controllers, learning controllers, hierarchical intelligent controllers) the controller construction methodology is largely heuristic and based on certain principles from Artificial Intelligence or Operations Research. The intelligent controllers are constructed to emulate, e.g., certain human cognitive functions to control complex dynamical processes. In the end implementation, however, nothing magical is created. The resulting intelligent controller is just a heuristically constructed nonlinear, perhaps adaptive system which is therefore amenable to control theoretic approaches to analysis. For instance, the simple direct fuzzy controller is a static nonlinear map, the expert controller may model certain "IF- THEN" statements in a control implementation (a type of nonlinearity) to ensure reliable operation, and many (numerical) learning controllers are types of nonlinear adaptive systems. More complex, multi-layer intelligent controllers are very complex adaptive decision making systems, but nevertheless they are nonlinear controllers (to convince yourself of this think of the implementation or simulation of the intelligent control system - if you can simulate it, you can write down equations representing it as a nonlinear dynamical system).

Hence, from a control engineer's perspective the focus should be on whether intelligent controllers are able to achieve higher performance with a greater degree of autonomy than their conventional predecessors. To develop this focus further, consider a general control system where P is a model of the plant, C represents the controller, and T represents specifications on how we would like the closed loop system to behave. For some classical control problems the scope is limited so that C and P are linear and T simply represents, for example, stability, rise time, and overshoot specifications. In this case intelligent control techniques may not be needed. As engineers, the simplest solution that works is the best one. We tend to need more complex controllers for more complex plants (where, for example, there is a significant amount of uncertainty) and more demanding closed loop specifications T. Consider the case where: (i) P is so complex that it is most convenient to represent it with ordinary differential and discrete event system models (or some other hybrid mix of models) and for some parts of the plant the model is not known (or it is too expensive to find), and (ii) T is used to characterize the desire to make the system perform well and act with high degrees of autonomy (i.e., so that the system performs well under significant uncertainties in the system and its environment for extended periods of time, and compensates for significant system failures without external intervention).

The general control problem is how to construct C, given P, so that T holds. From a control engineer's perspective, researchers in the field of intelligent control are trying to use intelligent (and conventional) control methodologies to solve this general control problem. It is important to note that researchers in intelligent control have been naturally led to focus on the very demanding general control problem described above (i) in order to address pressing needs for practical applications, and

(ii) since often there is a need to focus on representing more aspects of the plant so that they can be used to reduce the uncertainty in making high level decisions about how to perform control functions that are normally performed by humans.

Viewed as a control problem, the following research areas become very important for the field of intelligent control: - mathematical models for intelligent control systems (differential equations, discrete event systems, hybrid systems) - systematic (or perhaps automatable) design procedures for intelligent controllers - techniques for nonlinear analysis to study stability, boundedness, convergence issues, limit cycles, controllability, observability, robustness, etc. - performance analysis - simulation techniques for intelligent systems (particularly, hybrid systems) - implementation issues Hence, although intelligent controllers may operate in much more complex fashion than many conventional controllers to solve control problems that are beyond the focus of conventional control, we can find much in common with the field of conventional control in the areas of methodology and research issues.

Generally speaking the field of intelligent control is helping to expand the horizons of the field of control theory. Much of the drive to expand the focus of conventional control, through the field of intelligent control comes from the ever expanding frontiers of technology. Clearly, computer science, engineering, and technology drive the development of control theory, control engineering, and control technology by providing alternative strategies for the functionality and implementation of controllers for dynamical systems. For instance, the introduction of the microprocessor had significant impacts on: (i) the implementation and wide spread use of controllers, (ii) the expansion of the role of control systems over the times when they were implemented solely in an analog fashion, and (iii) the development of extensive theoretical results in control theory. While a portion of control theory naturally developed driven by technology, certain theoretical results allowed the technology to expand its role due to the fact that they provided methods to "guarantee" that the technology would work in critical environments (e.g., the use of stability theory for ensuring the safe operation of controllers for nuclear reactors and aircraft).

Analogous statements can be made relative to more recent developments in computer science and technology. For instance: What will the impact of highly parallel processing (e.g., via neural networks), fuzzy processors, or techniques from AI have on control engineering and the implementation of controllers? Is there a role for theoretical and experimental engineering analysis in expanding the use of intelligent control? From a control engineer's perspective, the field of intelligent control is trying to answer important questions such as these. Overall, we have computers with enhanced capabilities and we are trying to figure out what we can do with this added capability in the solution of control problems.

[1] Passino K.M., "Bridging the Gap Between Conventional and Intelligent Control", Special Issue on Intelligent Control, IEEE Control Systems Magazine, Vol. 13, pp. 12-18, 1993; See an expanded version of this paper: "Towards Bridging the Perceived Gap Between Conventional and Intelligent Control", to appear in Gupta M.M., Sinha N.K., eds., Intelligent Control: Theory and Practice, IEEE Press, Piscataway, NJ, 1994.
[2] Antsaklis P.J., Passino K.M., eds., An Introduction to Intelligent and Autonomous Control, Kluwer Academic Publishers, Norwell, MA, 1993.

3.6 On the Analytic Formulation of Intelligent Controls

by G.N. Saridis

The theory of Intelligent Control systems, developed by Saridis combines the powerful high-level decision making of the digital computer with advanced mathematical modeling and synthesis techniques of system theory and with linguistic methods of dealing with imprecise or incomplete information [1]. This produces a unified approach for the design of Intelligent Machines. The theory may be thought of as the result of the intersection of the three major disciplines of Artificial Intelligence, Operations Research, and Control Theory. The theory is aimed at establishing Intelligent Controls as an engineering discipline, with the purpose of designing Intelligent Autonomous Systems of the future.

Intelligent Control provides the fusion between the mathematical and linguistic methods and algorithms applied to systems and processes. It combines effectively the results of cognitive systems research, with various mathematical programming control techniques.

The control intelligence is hierarchically distributed according to the Principle of Precision with Decreasing Intelligence (IPDI), evident in all hierarchical management systems [2]. The analytic functions of an Intelligent Machine are implemented by Intelligent Controls, using Entropy as a measure. The resulting structure is composed of three basic levels of controls, each level of which may contain more than one layer of tree-structured functions:

The organization level; is modeled after a Boltzmann machine for abstract reasoning, task planning and decision making; The coordination level; is composed of a number of Petri Net Transducers supervised, for command exchange, by a dispatcher, which also serves as an interface to the organization level; The execution level; includes the sensory, planning for navigation and control hardware which interacts one-to-one with the appropriate coordinators, while a VME bus provides a channel for database exchange among the several devices.

The functions involved in the upper levels of an intelligent machine are imitating functions of human behavior and may be treated as elements of knowledge-based systems. Actually, the activities of planning, decision making, learning, data storage and retrieval, task coordination, etc., may be thought of as knowledge handling and management [3].

[1] Saridis G.N. and Valavanis K.P., RAnalytical Design of Intelligent MachinesS, Automatica the IFAC Journal, 24, No. 2, pp. 123-133, March 1988.
[2] Saridis G.N., RAnalytic Formulation of the IPDI for Intelligent MachinesS, Automatica the IFAC Journal, 25, No. 3, pp. 461-467, 1989.
[3] Valavanis K.P., Saridis G.N., Intelligent Robotic System Theory: Design and Applications, Kluwer Academic Publishers, Boston, MA., 1992.

3.7 On Intelligence and Intelligent Control

by P. Werbos

A naive way of responding to the issue of "What is intelligent control" is to regard this as an essentially empty issue of semantics. But there is more than semantics at stake here, and semantics do have some real significance. The word "intelligence" has a long and important history, and it would be a great shame if we in control decided to throw out its historical meaning. If we are attracted to concepts like autonomy, perhaps we should call them "autonomous" control. Instead of arguing over which concept gets into the definition, why not have separate words for separate concepts?

The word "intelligent" has two kinds of meanings, historically. Above all, it refers to the kinds of capabilities that the human brain possesses, in toto. Secondarily, it refers to ideas from AI intended to replicate some vision of the key components – planning over time, reasoning, etc.. If we have lots of nice little incremental improvements to make in control, I really wish we could agree to call them "smart control" or "brilliant control," and leave the word " intelligent" alone. I'm afraid I tend to view the usual supervisory control as one of those incremental improvements.

The formal statement I would personally propose to make about intelligent control is simply the one made in the two- page foreword in [1]: True intelligent control – control which replicates the most critical aggregate capabilities of human intelligence– does not exist in any artificial system today. A true intelligent controller would, above all, have to be capable of maximizing some notion of goal-achievement or utility over time in an uncertain, nonlinear environment, through a learning process which can be implemented efficiently on distributed hardware analogous to networks of neurons in the brain. It would also have to be capable of true real-time performance and learning. The learning and planning capabilities should be enough to allow the ability to learn higher-order symbolic reasoning, in principle, if enough hardware were available, to the extent that humans are capable of learning symbolic reasoning. Even though no one has built such a system yet, there has been substantial progress in understanding the key prerequisites to building such a system. The field of intelligent control may be defined as that community of researchers who believe that they have a clear plan or vision of how such a controller might be built, through a strategic vision of current research opportunities. The visions of how to reach this end point are in fact very varied. Some of us hope that small incremental improvements of existing controllers may do the job. Others believe that that is like trying to build an airplane by successive improvements in an auto engine used on the ground. It is very clear, however, that the end point here cannot be achieved without a greater synthesis of new concepts from neural network theory, from adaptive and optimal control theory, and from various strands of AI, such as machine learning or fuzzy logic. It is also clear that this will require true intellectual synthesis, and not just cutting and pasting; for example, there are concepts involving approximate dynamic programming which may provide a basis for unifying and implementing concepts from all these different fields in a unified learning control system. The development and understanding of true intelligent control will require a great deal of boldness, but the potential benefits are also enormous. In addition to the obvious technological benefits, it may well be that a mathematical, engineering-grounded understanding of intelligent control will be an absolute prerequisite to a true understanding of intelligence as it exists in the human brain and the human mind. From a strict scientific point of view, such an understanding would be comparable in importance to Newton's discovery of gravity (for which calculus was a prerequisite).

[1]. White D.A., Sofge D.A., eds., Handbook of Intelligent Control Neural, Fuzzy, and Adaptive Approaches, Van Nostrand 1992.
[2]. Werbos P., "Elastic fuzzy logic: a better fit to neurocontrol and true intelligence", Journal of Intelligent and Fuzzy Systems, Vol. 1, No. 4, 1993.
[3]. Werbos P., Roots of Backpropagation: From Ordered Derivatives to Neural Networks and Political Forecasting, Wiley, 1993.

3.8 Additional References

The following books are good sources of additional references on Intelligent Control Systems:

[1]. An Introduction to Intelligent and Autonomous Control, Antsaklis P.J., Passino K.M., eds., Kluwer Academic Publishers, Norwell, MA, 1993.
[2]. Handbook of Intelligent Control Neural, Fuzzy, and Adaptive Approaches, White D.A., Sofge D.A., eds., Van Nostrand 1992.
[3]. Intelligent Control: Theory and Practice, Gupta M.M., Sinha N.K., eds., IEEE Press, Piscataway, NJ, 1994.
[4]. Meystel A., Autonomous Mobile Robots, World Scientific, 1991.
[5]. Valavanis K.P., Saridis G.N., Intelligent Robotic System Theory: Design and Applications, Kluwer Academic Publishers, Boston, MA., 1992.

4 LIST OF TASK FORCE MEMBERS

James S. Albus National Institute of Standards and Technology
Panos J. Antsaklis University of Notre Dame
Kishan Baheti National Science Foundation
J. Douglas Birdwell University of Tennessee
Michael D. Lemmon University of Notre Dame
Maja J. Mataric MIT, AI Lab
Alex Meystel Drexel University
Kumpati S. Narendra Yale University
Kevin M. Passino The Ohio State University
Herbert E. Rauch. Lockheed, Palo Alto, CA
George N. Saridis Rensselaer Polytechnic Institute
Harry Stephanou Rensselaer Polytechnic Institute
Paul Werbos National Science Foundation

Also thanks to M. Kokar of Northeastern University for his contribution to section 2.