PROCEEDINGS

of the

1994 IEEE International Symposium on INTELLIGENT CONTROL

16 - 18 August 1994
Holiday Inn Crowne Plaza
Columbus, Ohio
USA

IEEE

IEEE Control Systems Society
94CH3453-8

CSS

PROCEEDINGS

of the

1994 IEEE International Symposium on INTELLIGENT CONTROL

16 - 18 August 1994
Holiday Inn Crowne Plaza
Columbus, Ohio
USA

IEEE

IEEE Control Systems Society
94CH3453-8

CSS

1994 IEEE International Symposium on Intelligent Control

IEEE Catalog Number 94CH3453 - 8

ISBN 0 - 7803 - 1990 - 7 (Softbound)
0 - 7803 - 1991 - 5 (Casebound)
0 - 7803 - 1992 - 3 (Microfiche)

Library of Congress Number 90 - 655042

Sponsor: IEEE Control Systems Society

Operating Committee:

General Chair:	Ümit Özgüner *The Ohio State University*
Program Chairs:	Levent Acar *National Institute of Standards and Technology* Michael B. Leahy *USAF: SA-ALC/TIEST*
Local Arrangements Chair:	James F. Davis *The Ohio State University*
Publications Chair:	John S. Bay *Virginia Polytechnic Institute and State University*
Finance Chair:	Kevin M. Passino *The Ohio State University*
Registration Chair:	Jeffrey D. Donne *AEG, Automation Systems Group*

1994 IEEE International Symposium on Intelligent Control

International Program Committee

J. Albus, *National Institute of Standards and Technology*
B. Balaram, *California Institute of Technology JPL*
J. Bay, *Virginia Polytechnic Institute & State University*
H. Berenji, *NASA Ames Research Center*
M. Buehler, *Center of Intelligent Machines*
J. Farrell, *Charles Stark Draper Laboratory*
J. Feddema, *Sandia National Laboratory*
T. Fukuda, *Nagoya University, Japan*
O. Kaynak, *Bogazici University, Turkey*
M. Kokar, *Northeastern University*
M. Lemmon, *University of Notre Dame*
A. Meystel, *Drexel University*
K. Passino, *Ohio State University*
V. Rao, *University of Missouri-Rolla*
H. Rauch, *Palo Alto Research Laboratory*
T. Sobh, *University of Utah*
K. Valavanis, *University of Southwestern Louisiana*
Y. Zheng, *Ohio State University*

Forward

INTELLIGENT CONTROL:
PROBLEMS AND METHODOLOGY

The ever increasing technological demands of today call for very complex systems, which in turn require highly sophisticated controllers to ensure that high performance can be achieved and maintained under adverse conditions. There are needs in the control of these complex systems which cannot be met by conventional approaches to control. For instance, there is a significant interest in enhancing current avionic systems so that they can reconfigure the aircraft controls to maintain adequate levels of performance even if there are complete failures in one or more of the actuators or sensors. In a similar manner, there is a significant need to achieve higher degrees of autonomous operation for robotic systems, spacecraft, manufacturing systems, automotive systems, underwater and land vehicles, and others. To achieve such highly autonomous behavior for complex systems one can enhance today's control methods using intelligent control systems and techniques.

The area of Intelligent Control is a fusion of a number of research areas in Systems and Control, Computer Science, and Operations Research among others, coming together, merging and expanding in new directions and opening new horizons to address the problems of this challenging and promising area. Intelligent control systems are typically able to perform one or more of the following functions to achieve autonomous behavior: planning actions at different levels of detail, emulation of human expert behavior, learning from past experiences, integrating sensor information, identifying changes that threaten the system behavior, such as failures, and reacting appropriately. This identifies the areas of Planning and Expert Systems, Fuzzy Systems, Neural Networks, Machine Learning, Multi-sensor Integration, Failure Diagnosis, and Reconfigurable Control, to mention but a few, as existing research areas that are related and important to Intelligent Control. While these techniques provide several key approaches to Intelligent Control, for complex systems they are often interconnected to operate within an architecture which is hierarchical and often distributed. It is for this reason that the areas of hierarchical intelligent control, distributed intelligent control, and architectures for intelligent systems are of significant importance in the design and construction of the overall intelligent controller for complex dynamical systems.

Finally, it is of fundamental importance to recognize that (i) intelligent controllers are nonlinear (possibly hierarchical and distributed) controllers that are constructed in nonconventional ways, and (ii) intelligent controllers are often designed to operate in "critical environments" where, for example, the safety of a crew (e.g., in an aircraft/spacecraft), or environmental issues are of concern (e.g., from nuclear power plants or process control). Hence, it is both possible, and of significant importance to introduce mathematical modeling and analysis techniques to be used in the verification and certification of the behavior of intelligent control systems.

Topics of interest include:

- Architectures for intelligent control
- Hierarchical intelligent control
- Distributed intelligent systems
- Modeling intelligent systems
- Mathematical analysis of intelligent systems
- Discrete event systems
- Hybrid systems
- Design techniques for intelligent controllers
- Knowledge-based and expert systems
- Planning systems
- Fuzzy systems / fuzzy control
- Machine learning / adaptive systems
- Genetic algorithms
- Neural networks / neural control
- Reconfigurable control
- Multisensor fusion / integration
- Pattern recognition
- Failure detection and identification
- Applications / Implementations
  - Robotics
  - Manufacturing systems
  - Automotive/transportation systems
  - Intelligent vehicle highway systems
  - Aircraft / spacecraft / satellites
  - Underwater / land vehicles
  - Process operations and control
  - Consumer products

Ümit Özgüner
General Chair

ISIC Past And Present

The 1994 ISIC is the 9th International Symposium on Intelligent Control. From the beginning, the symposium filled a real need and attracted researchers who felt that a forum for discussion and exchange of ideas in this area was necessary. Today the term "Intelligent Control" is part of everyday vocabulary. Its use, and sometimes its misuse, signifies the importance this area has attained, and this symposium was where the early steps were taken.

The following summary of locations, chairs and dates outlines the history of the symposium. Note that "GC" and "PC" denote General Chair and Program Chair respectively.

1985 IEEE International Symposium on Intelligent Control
GC: G.N. Saridis, PC: A. Meystel
Rensselaer Polytechnic Institute, Troy, New York
August 26, 1985
Sponsored by the National Science Foundation

1987 IEEE International Symposium on Intelligent Control
GC: A. Meystel, PC: J.Y.S. Luh
Penn Tower Hotel, Philadelphia, Pennsylvania
January 19-20, 1987
Sponsored by the IEEE Computer Society and the Philadelphia section

1988 IEEE International Symposium on Intelligent Control
GC: H.E. Stephanou, PC: A. Meystel, J.Y.S. Luh
Key Bridge Marriott Hotel, Arlington, Virginia
August 24-26, 1988
Sponsored by the IEEE Control Systems Society; in collaboration with the IEEE Systems, Man, and Cybernetics Society and the IEEE Industrial Electronics Society

1989 IEEE International Symposium on Intelligent Control
GC: A.C. Sanderson, PC: A.A. Desrochers, K. Valavanis
Desmond Americana, Albany, New York
September 25-26, 1989
Sponsored by the IEEE Control Systems Society

1990 IEEE International Symposium on Intelligent Control
GC: A. Meystel, PC: H. Kwatny, S. Navathe, H. Wechsler
Penn Tower Hotel, Philadelphia, Pennsylvania
September 5-7, 1990
Co-Sponsored by the IEEE Control Systems Society and the IEEE Society of Robotics and Automation; in cooperation with the IEEE Computer Society

1994 IEEE International Symposium on Intelligent Control

1991 IEEE International Symposium on Intelligent Control
GC: H.E. Stephanou, PC: A.H. Levis
Key Bridge Marriott Hotel, Arlington, Virginia
August 13-15, 1991
Sponsored by the IEEE Control Systems Society. Concurrently with the 1991 IFAC International Symposium on Distributed Intelligence Systems

1992 IEEE International Symposium on Intelligent Control
GC: E. Grant, PC: T.C. Henderson
The Albany Hotel, Glasgow, Scotland
August 11-13, 1992
Sponsored by the IEEE Control Systems Society

1993 IEEE International Symposium on Intelligent Control
GC: P. Antsaklis, PC: K. Passino, U. Özgüner
The Knickerbocker Hotel, Chicago, Illinois
August 25-27, 1993
Sponsored by the IEEE Control Systems Society

DEFINING INTELLIGENT CONTROL

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

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