Proc. of IEEE Intelligent Robots & Systems '91 -- Intelligence for Mechanical Systems; November 1991; Osaka, Japan

# A Theory of Intelligent Machine Systems

**James S. Albus**
Robot Systems Division
National Institute of Standards and Technology
Gaithersburg, MD 20899

### *Abstract*

*A theoretical model is proposed consisting of six basic elements: actuators, sensors, sensory processing, world modeling, behavior generation, and value judgment. These elements are integrated into a hierarchical system architecture wherein: a) control bandwidth decreases about an order of magnitude at each higher level, b) goals expand in scope and planning horizons expand in space and time about an order-of-magnitude at each higher level, and c) perception of spatial and temporal patterns, models of the world, and memories of events decrease in resolution and expand in spatial and temporal range by about an order-of-magnitude at each higher level.*

*At each level, functional modules perform behavior generation, world modeling, sensory processing, and value judgment. Behavior generation modules perform task planning and execution. The world model is the intelligent system's best estimate of the external world. Value judgments provide an evaluation of hypothesized plans, and perceived objects, events, and situations. Sensory processing modules transform sensory maps into world model maps and extract entity attributes and states.*

*Sensory feedback control loops are closed at every level.*

## Introduction

Much is unknown about intelligence, and much will elude human understanding for a very long time. Yet much is known, both about the mechanisms and function of intelligence. It is not too soon to propose at least the beginnings of a theory of intelligent systems.

The study of intelligent machines and the neurosciences are extremely active fields. Many millions of dollars per year are now being spent in Europe, Japan, and the United States on computer integrated manufacturing, robotics, and intelligent machines for a wide variety of military and commercial applications. Around the world, researchers in the neurosciences are searching for the anatomical, physiological, and chemical basis of behavior. Research in learning automata, neural nets, and brain modeling has given insight into learning and the similarities and differences between neuronal and electronic computing processes. Computer science and artificial intelligence is probing the nature of language and image understanding, and has made significant progress in rule based reasoning, planning, and problem solving. Game theory and operations research have developed methods for decision making in the face of uncertainty. Robotics and autonomous vehicle research has produced advances in real-time sensory processing, world modeling, navigation, trajectory generation, and obstacle avoidance. Research in automated manufacturing and process control has produced intelligent hierarchical controls, distributed databases, representations of object geometry and material properties, data driven task sequencing, network communications, and multiprocessor operating systems. Modern control theory has developed precise understanding of stability, adaptability, and controllability under various conditions of feedback and noise. Research in sonar, radar, and optical signal processing has developed methods for fusing sensory input from multiple sources, and assessing the believability of noisy data.

Progress is rapid, and there exists an enormous and rapidly growing literature in each of the areas mentioned above. What is lacking is a general theoretical model of intelligent systems which ties all these separate fields of knowledge into a unified framework. It is time to begin the construction of such a framework.

To have a Theory of Intelligent Machines, one must first have a theory of intelligence. A theory implies that one can define what intelligence is, explain where it came from, describe how it works, and tell how to build it.

## Definition

In order to be useful, the definition of intelligence must not be limited to behavior that is beyond our understanding. 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.

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.

For the purposes of this paper, 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.

The measure of intelligence is not a scalar, but a vector or tensor value 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, 3) the information and values the system has stored in its memory, and 4) the elegance of the system architecture in which the computational modules are organized.

The magnitude of 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.

Learning is not required in order to <u>be</u> intelligent, only to <u>become</u> more intelligent as a result of experience. Learning can be defined as consolidating short-term memory into long-term memory, and exhibiting altered behavior because of what was remembered. Learning is a mechanism for storing knowledge about the external world, and for acquiring skills and knowledge of how to act. However, many creatures exhibit intelligent behavior using instinct, without having learned anything.

## Origin

Natural intelligence, like the brain in which it appears, is a result of the process of natural selection. The brain is first and foremost a control system. Its primary function is to produce successful goal-seeking behavior in finding food, avoiding danger, competing for territory, attracting sexual partners, and caring for offspring. All brains that ever existed, even those of the tiniest insects, generate and control behavior. Some brains produce only simple forms of behavior, while others produce very complex behaviors. Only the most recent and highly developed brains show any evidence of abstract thought. For each individual, intelligence provides a mechanism for generating biologically advantageous behavior.

Intelligence improves an individual's ability to act effectively and choose wisely between alternative behaviors. All else being equal, a more intelligent individual has many advantages over less intellgient rivals in acquiring choice territory, gaining access to food, and attracting more desirable breeding partners. The intelligent use of aggression helps to improve an individual's position in the social dominance hierarchy. Intelligent predation improves success in capturing prey. Intelligent exploration improves success in hunting and establishing territory. Intelligent use of stealth gives a predator the advantage of surprise. Intelligent use of deception improves the prey's chances of escaping from danger.

Higher levels of intelligence produce capabilities in the individual for thinking ahead, planning before acting, and reasoning about the probable results of alternative actions. These abilities give to the more intelligent individual a competitive advantage over the less intelligent in the competition for survival and gene propagation. Intellectual capacities and behavioral skills that produce successful hunting and gathering of food, acquisition and defense of territory, avoidance and escape from danger, and bearing and raising offspring tend to be passed on to succeeding generations. Intellectual capabilities that produce less successful behaviors reduce the survival probability of the brains that generate them. Competition between individuals and groups thus drives the evolution of intelligence within a species and between species.

## The Elements of Intelligence

The elements of intelligence and their relationship to each other are illustrated in Figure 1.

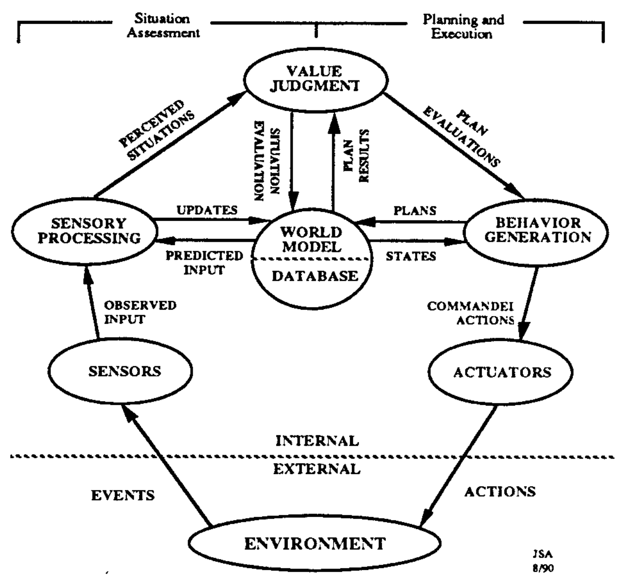


**Figure 1.** The elements of intelligence and the functional relationships between them.

a. ACTUATORS -- Output from an intelligent system derives from actuators which move, exert forces, and position arms, legs, hands, and eyes. Actuators generate forces to point sensors, excite transducers, move manipulators, handle tools, steer and propel locomotion. An intelligent system may have tens, hundreds, or even thousands of actuators, all of which must be coordinated in order to perform tasks and accomplish goals. Natural actuators are muscles and glands. Machine actuators are motors, pistons, valves, solenoids, and transducers.

b. SENSORS -- Input to an intelligent system derives from sensors. These may include visual brightness and color sensors; tactile, force, torque, position detectors; velocity, vibration, acoustic, range, smell, taste, pressure, and temperature measuring devices. Sensors may be used to monitor both the state of the external world and the internal state of the intelligent system itself. Sensors provide input to a sensory processing system.

c. SENSORY PROCESSING -- Perception takes place in a sensory processing system that compares observations with expectations generated by an internal world model. Sensory processing algorithms integrate similarities and differences between observations and expectations, over time and space, so as to detect events, and recognize features, objects, and relationships in the world. Sensory input data from a wide variety of sensors over extended periods of time may be fused into a consistent unified perception of the state of the world. Sensory processing algorithms compute distance, shape, orientation, surface characteristics, physical and dynamical attributes of objects and regions of space. Sensory processing may include recognition of acoustic signatures, speech, and interpretation of language.

d. WORLD MODEL -- The world model is the intelligent system's best estimate of the state of the world. The world model includes a database of knowledge about the world, plus a database management system that stores and retrieves information. The world model also contains a simulation capability which generates expectations and predictions. The world model thus can provide answers to requests for information about the present, past, and probable future states of the world. The world model provides this information service to the behavior generation system, so that it can make intelligent plans and behavioral choices, and to the sensory processing system, in order for it to perform correlation, model matching, and model based recognition of states, objects, and events. The world model is kept up-to-date by the sensory processing system.

e. VALUES -- The value system makes value judgments as to what is good and bad, rewarding and punishing, important and trivial. The value system evaluates both the observed state of the world and the predicted results of hypothesized plans. It computes costs, risks, and benefits both of observed situations and of planned activities. The value system thus provides the basis for choosing one action as opposed to another, or for acting on one object as opposed to another. The value system also computes the probability of correctness and assigns believability and uncertainty parameters to world model state estimations.

f. BEHAVIOR GENERATION -- Behavior is generated in a behavior generating system that plans and executes tasks by decomposing them into subtasks, and by sequencing these subtasks so as to achieve goals. Goals are selected and plans generated by a looping interaction between behavior generation, world modeling, and value judgment functions. The behavior generation system hypothesizes plans, the world model predicts the results of those plans, and the value judgment system evaluates those results. The behavior generation system then selects the plans with the best evaluations for execution. Behavior generation monitors the execution of task plans, and modifies existing plans whenever the situation requires.

In many cases, intelligent behavior generation requires the ability to reason about space and time, geometry and dynamics, and to formulate or select plans based on values such as cost, risk, utility, and goal priorities. Task planning and execution often must be done in the presence of uncertain, incomplete, and sometimes incorrect information.

In order for behavior generation to succeed in a dynamic and unpredictable world, it must be accomplished in real-time. In order to achieve real-time behavior generation, it is necessary to partition the planning problem into a hierarchy of levels with different temporal planning horizons and different degrees of detail at each hierarchical level. Once this is done, it is possible to employ a multiplicity of planners to simultaneously generate and coordinate plans for many different subsystems at many different levels of resolution.

## The System Architecture of Intelligence

Each of the elements of intelligent systems are reasonably well understood. The phenomena of intelligence, however, requires more than a set of disconnected elements. Intelligence requires an interconnecting system architecture that enables the various system components to interact and communicate with each other in intimate and sophisticated ways.

A system architecture is what partitions the elements of intelligence into computational modules, and interconnects the modules in networks and hierarchies. It is what enables the behavior generation system to direct sensors, and to focus sensory processing algorithms on objects and events worthy of attention, ignoring things that are not important to current goals and task priorities. It is what enables the world model to answer queries from behavior generation modules, and make predictions and receive updates from sensory processing modules. It is what communicates the value state-variables that characterize the success of behavior and the desirability of states of the world.

A number of intelligent system architectures have been proposed, and a few have been implemented. [1] The model of intelligence that will be discussed here is largely based on the Real-time Control System (RCS) that has been implemented in a number of versions over the past 13 years at the National Institute for Standards and Technology (NIST formerly NBS) [2,3].

The RCS system architecture organizes the elements of intelligence so as to create the functional relationships and information flow shown in Figure 1. In all intelligent systems, a sensory processing system processes sensory information to acquire and maintain an internal model of the external world. In all systems, a behavior generating system controls actuators so as to pursue behavioral goals in the context of the perceived world model. In systems of higher intelligence, the behavior generating system element may interact with the world model and value judgment system to reason about space and time, geometry and dynamics, and to formulate or select plans based on values such as cost, risk, utility, and goal priorities. The sensory processing system element may interact with the world model and value judgment system to assign values to perceived entities, events, and situations.

The proposed system architecture replicates and distributes the relationships shown in Figure 1 over a hierarchical computing structure with the logical and temporal properties illustrated in Figure 2. On the left is an organizational hierarchy wherein computational nodes are arranged in layers like command posts in a military organization. Each node in the organizational hierarchy contains four types of computing modules: behavior generating (BG), world modeling (WM), sensory processing (SP), and value judgment (VJ) modules. Each chain of command in the organizational hierarchy, from each actuator and each sensor to the highest level of control, can be represented by a computational hierarchy, such as is shown in the center of Figure 2.

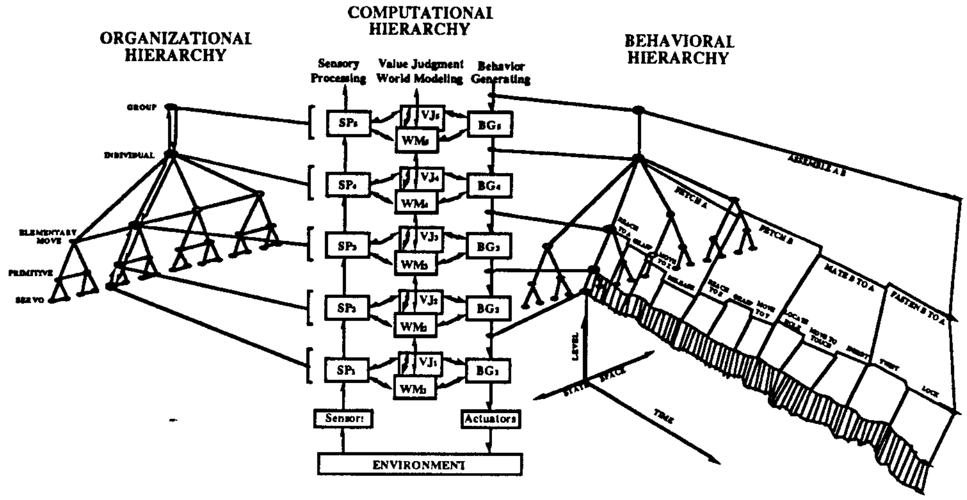


**Figure 2.** Relationships in hierarchical control systems. On the left, is an organizational hierarchy consisting of a tree of command centers, each of which possesses one supervisor and one or more subordinates. In the center, is a computational hierarchy consisting of BG, WM, SP, and VJ modules. Each actuator and each sensor is serviced by a computational hierarchy. On the right, is a behavioral hierarchy consisting of trajectories through state-time-space. Commands at each level can be represented by vectors, or points in state-space. Sequences of commands can be represented as trajectories through state-time-space.

At each level, the nodes, and computing modules within the nodes, are richly interconnected to each other by a communications system. Within each computational node, the communication system provides intermodule communications of the type shown in Figure 1. Queries and task status are communicated from BG modules to WM modules. Retrievals of information are communicated from WM modules back to the BG modules making the queries. Predicted sensory data is communicated from WM modules to SP modules. Updates to the world model are communicated from SP to WM modules. Observed entities, events, and situations are communicated from SP to VJ modules. Values assigned to the world model representations of these entities, events, and situations are communicated from VJ to WM modules. Hypothesized plans are communicated from BG to WM modules. Results are communicated from WM to VJ modules. Evaluations are communicated from VJ modules back to the BG modules that hypothesized the plans.

The communications system also communicates between nodes at different levels. Commands are communicated downward from supervisor BG modules in one level to subordinate BG modules in the level below. Status reports are communicated back upward through the world model from lower level subordinate BG modules to the upper level supervisor BG modules from which commands were received. Observed entities, events, and situations detected by SP modules at one level are communicated upward to SP modules at a higher level. Predicted attributes of entities, events, and situations stored in the WM modules at a higher level are communicated downward to lower level WM modules. Output from the bottom level BG modules is communicated to actuator drive mechanisms. Input to the bottom level SP modules is communicated from sensors.

The communications system can be implemented in a variety of ways. In a biological brain, communication is mostly via neuronal axon pathways, although some messages are communicated by hormones carried in the bloodstream. In artificial systems, the physical implementation of communications functions may be a computer bus, a local area network, a common memory, a message passing system, or some combination thereof. In either biological or artificial systems, the communications system may include the functionality of a communications processor, a file server, a database management system, a question answering system, or an indirect addressing or list processing engine. In the system architecture proposed here, the input/output relationships of the communications system produce the effect of a virtual global memory, or blackboard system.

The string of input commands to each of the BG modules at each level generates a trajectory through state-space as a function of time. The set of command strings to all BG modules creates a behavioral hierarchy, as shown on the right of Figure 2. Actuator output trajectories (not shown in Figure 2) correspond to observable output behavior. All the other trajectories in the behavioral hierarchy constitute the deep structure of behavior.

## Hierarchical vs. Horizontal

Figures 3 shows the organizational hierarchy in more detail, and illustrates both the hierarchical and horizontal relationships involved in the proposed architecture. The architecture is hierarchical in that commands and status feedback flow hierarchically up and down a behavior generating chain of command. The architecture is also hierarchical in that sensory processing and world modeling functions have hierarchical levels of temporal and spatial aggregation. The command hierarchy is a tree, in that at any instant of time any BG module has only one supervisor.

The SP modules are also organized hierarchically, but as a layered graph, not a tree. At each higher level, sensory information is processed into increasingly higher levels of abstraction, but the sensory processing pathways may branch and merge in many different ways.

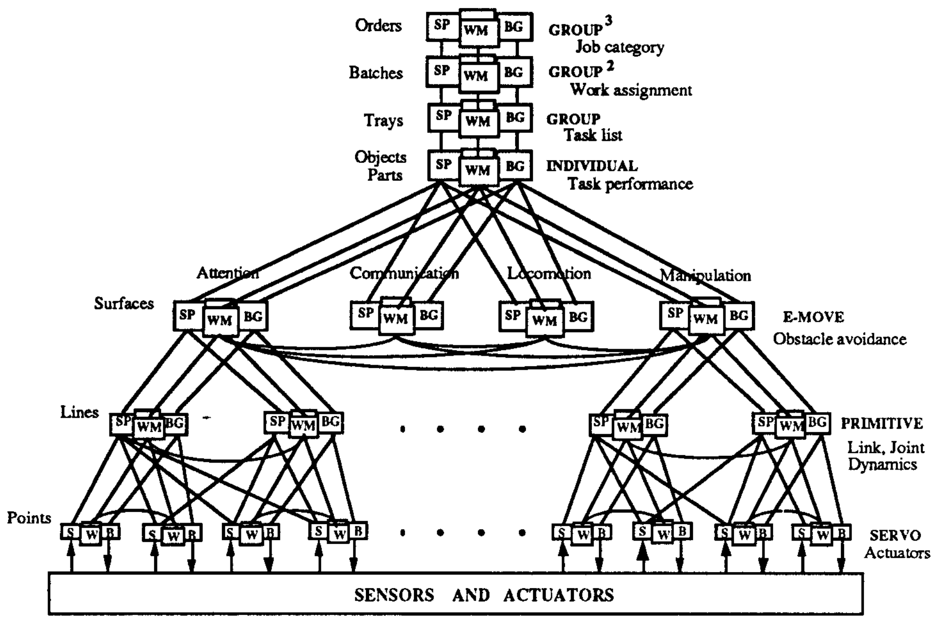


**Figure 3.** An organization of processing nodes such that the BG modules form a command tree. On the right are examples of the functional characteristics of the BG modules at each level. On the left are examples of the type of visual and acoustical entities recognized by the SP modules at each level. In the center of level 3 are the type of subsystems represented by processing nodes at level 3.

The architecture is horizontal in that data is shared horizontally between heterogeneous modules at the same level. At each hierarchical level, the architecture is horizontally interconnected by wide-bandwidth communication pathways between BG, WM, SP, and VJ modules in the same node, and between nodes at the same level, especially within the same command subtree. The horizontal flow of information is most voluminous within a single node, less so between related nodes in the same command subtree, and relatively low bandwidth between computing modules in separate command subtrees.

The volume of information flowing horizontally within a subtree may be orders of magnitude larger than the amount flowing vertically in the command chain. The volume of information flowing vertically in the sensory processing system can also be very high, especially in a vision system.

The specific configuration of the command tree is task dependent, and therefore not necessarily stationary in time. Figure 3 illustrates only one possible configuration that may exist at a single point in time. During operation, relationships between modules within and between layers of the hierarchy may be reconfigured in order to accomplish different goals, priorities, and task requirements. This means that any particular computational node, with its BG, WM, SP, and VJ modules, may belong to one subsystem at one time and a different subsystem a very short time later. For example, a mobile robot may belong to one workstation at one instant in time, and a different workstation a very short time later. Similarly, a robot may employ one gripper for one task, and a different gripper with different capabilities for another task.

In the biological brain, command tree reconfiguration can be implemented through multiple axon pathways that exist, but are not always activated, between BG modules at different hierarchical levels. In intelligent machine systems, multiple communications pathways may enable each BG module to receive input messages and parameters from several different sources. During operation, goal driven switching mechanisms in the BG modules assess priorities, negotiate for resources, and coordinate task activities so as to select among the possible communication paths. As a result, each BG module will accept task commands from only one supervisor at a time, and hence the BG modules form a command tree at every instant in time.

## Hierarchical Levels

Levels in the behavior generating hierarchy are defined by temporal and spatial decomposition of goals and tasks into levels of resolution. Temporal resolution is manifested in terms of loop bandwidth, sampling rate, and state-change intervals. Temporal span is measured by the length of historical traces and planning horizons. Spatial resolution is manifested in the branching of the command tree and the resolution of maps. Spatial span is measured by the span of control and the range of maps.

Levels in the sensory processing hierarchy are defined by temporal and spatial integration of sensory data into levels of aggregation. Spatial aggregation is best illustrated by visual images. Temporal aggregation is best illustrated by acoustic parameters such as phase, pitch, phonemes, words, sentences, rhythm, beat, and melody.

Levels in the world model hierarchy are defined by temporal resolution of events, spatial resolution of maps, and by "parent-child" relationships between entities in symbolic data structures. These are defined by the needs of both SP and BG modules at the various levels.

The interconnections between modules in nodes, and between nodes within and between levels result in a hierarchically structured goal-driven, sensory-interactive, intelligent control system architecture wherein:

a) control bandwidth decreases about an order of magnitude at each higher level,

b) goals expand in scope and planning horizons expand in space and time about an order-of-magnitude at each higher level, and

c) perception of spatial and temporal patterns, models of the world, and memories of events decrease in resolution and expand in spatial and temporal range by about an order-of-magnitude at each higher level.

## Timing

The timing diagram in Figure 4 illustrates how the range of the time scale increases, and resolution decreases, exponentially by about an order of magnitude at each higher level. The seven hierarchical levels in Figure 4 span a range of time intervals from three milliseconds to one day. Three milliseconds was arbitrarily chosen as the shortest servo update rate because that is adequate to reproduce the highest bandwidth reflex arc in the human body. One day was arbitrarily chosen as the longest historical-memory/planning-horizon to be considered. Shorter time intervals could be handled by adding another layer at the bottom. Longer time intervals could be treated by adding layers at the top, or by increasing the difference in loop bandwidths and sensory chunking intervals between levels.

The origin of the time axis in Figure 4 is the present, i.e. t=0. Future plans lie to the right of t=0, past history to the left. The open triangles in the right half-plane represent task goals in a future plan. The filled triangles in the left half-plane represent recognized task-completion events in a past history. At each level there is a planning horizon and a historical event summary interval. The planning horizon and event summary interval increases, and the loop bandwidth and frequency of subgoal events decreases, exponentially at each higher level.

The heavy cross-hatching on the right shows the planning horizon for the current task. The light shading on the right indicates the planning horizon for the anticipated next task. The heavy cross-hatching on the left shows the event summary interval for the current task. The light shading on the left shows the event summary interval for the immediately previous task.

Figure 4 suggests a duality between the behavior generation and the sensory processing hierarchies. At each hierarchical level, planner modules decompose task commands into strings of planned subtasks for execution. At each level, strings of sensed events are summarized, integrated, and "chunked" into single events at the next higher level.

## Planning

Planning implies an ability to predict future states of the world. Prediction algorithms based on Baysian statistics, Fourier transforms, or Kalman filters typically use recent historical data to compute parameters for extrapolating into the future. Predictions made by such methods are typically not reliable for periods longer than the historical interval over which the parameters were computed. Thus at each level, planning horizons extend into the future only about as far, and with about the same level of detail, as historical traces reach into the past.

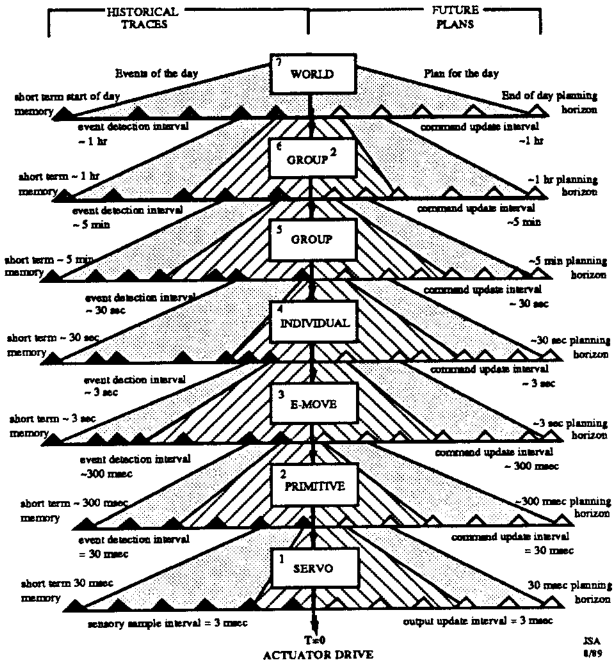


**Figure 4.** A timing diagram illustrating the temporal flow of activity in the task decomposition and sensory processing systems. At the world level, high level sensory events and circadian rhythms react with habits and daily routines to generate a plan for the day. Each element of that plan is decomposed through the remaining six levels of task decomposition into action.

Predicting the future state of the world often depends on assumptions as to what actions are going to be taken and what reactions are to be expected from the environment, including what actions may be taken by other intelligent agents. Planning implies a search over the space of possible future actions and probable reactions. Planning search takes place via a looping interaction between the BG, WM, and VJ modules whereby various possible futures are simulated and evaluated.

Planning complexity grows linearly with the number of possible actions at each step in the plan, and exponentially with the number of steps (i.e. the number of layers in the search graph). If real-time planning is to succeed, any given planner must operate in a limited search space. If there are too much resolution in the time line, or in the space of possible actions, the size of the search graph can easily become too large for real-time response. One method of resolving this problem is to use a multiplicity of planners in hierarchical layers so that at each layer no planner needs to search more than a given number (for example ten) steps deep in a game graph, and at each level there are no more than (ten) subsystem plans that need to be simultaneously generated and coordinated.

These criteria give rise to hierarchical levels with exponentially expanding spatial and temporal planning horizons, and characteristic degrees of detail for each level. The result of hierarchical spatio-temporal planning is illustrated in Figure 5. At each level, plans consist of at least one, and on average 10, subtasks. The planners have a planning horizon that extends about one and a half average input command intervals into the future.

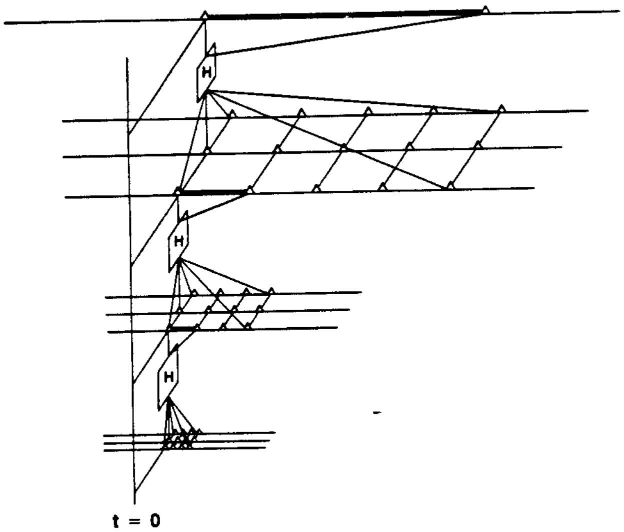


**Figure 5.** Three levels of real-time planning illustration the shrinking planning horizon and greater details at successively lower levels of the hierarchy. At the top level a single task is decomposed into a set of four planned subtasks for each of three subsystems. At each of the next two levels, the first task in the plan of the first subsystem is further decomposed into four subtasks for three subsystems at the next lower level.

In a real-time system, plans must be regenerated periodically to cope with changing and unforeseen conditions in the world. Cyclic replanning may occur at periodic intervals. Emergency replanning begins immediately upon the detection of an emergency condition. Under full alert status, the cyclic replanning interval should be about an order of magnitude less than the planning horizon (or about equal to the expected output subtask time duration). This requires that real-time planners be able to search to the planning horizon about an order of magnitude faster than real time. This is possible only if the depth and resolution of search is limited through hierarchical planning.

Plan executors within each BG have responsibility for reacting to feedback every control cycle interval. Control cycle intervals are inversely proportional to the control loop bandwidth. Typically the control cycle interval is an order of magnitude less than the expected output subtask duration. If the feedback indicates the failure of a planned subtask, the executor branches immediately (i.e. in one control cycle interval) to a preplanned emergency subtask. The planner simultaneously selects or generates an error recovery sequence which is substituted for the former plan which failed.

When a task goal is achieved at time t=0, it becomes a task completion event in the historical trace. To the extent that a historical trace is an exact duplicate of a former plan, there were no surprises; i.e. the plan was followed, and every task was accomplished as planned. To the extent that a historical trace is different from the former plan, there were surprises. The average size and frequency of surprises (i.e. differences between plans and results) is a measure of effectiveness of a planner (indeed of an intelligent control system).

At each level in the control hierarchy, the difference vector between planned (i.e. predicted) and observed events is an error signal, that can be used by executor submodules for servo feedback control (i.e. error correction), and by VJ modules for evaluating success and failure.

## Conclusion

The RCS system architecture outlined above has been elaborated many times for many different applications. RCS was first implemented by Barbera for laboratory robotics in the mid 1970's. It was adapted by Albus, Barbera, and others for manufacturing control in the NIST Automated Manufacturing Research Facility (AMRF) during the early 1980's [3]. Since 1986, RCS has been implemented for a number of additional applications, including the NBS/DARPA Multiple Autonomous Undersea Vehicle (MAUV) project and the Army TMAP and TEAM semi-autonomous land vehicle projects. RCS also forms the basis of the NASA/NBS Standard Reference Model Telerobot Control System Architecture (NASREM) being used on the space station Flight Telerobotic Servicer [4]. NASREM is also being used as a development guideline for the European Space Agency Telerobotics program, the U.S. Bureau of Mines Coal Mine Automation program, the DARPA Advanced Submarine Technology program, the Army Robotics Testbed program, and the Air Force Next Generation Controller program.

Current work at NIST and many other laboratories is directed at developing intelligent control systems based on RCS. Efforts are being directed toward developing a Standard Architecture for Real-Time Intelligent Control Systems along with a Methodology for Real-Time Intelligent Control System Development.

## References


[1] J.S. Albus, "Outline for a Theory of Intelligence", IEEE Transactions on Systems, Man, and Cybernetics (In press)

[2] J.S. Albus, *Brains, Behavior, and Robotics*, BYTE/McGraw-Hill, Peterbourough, N.H.,1981 (Also available in Japanese version)

[3] A.J. Barbera, J.S. Albus, M.L. Fitzgerald, and L.S. Haynes, "RCS: The NBS Real-time Control System", *Proceedings Robots 8 Conference and Exposition*, Detroit, MI, June 1984

[4] J.S. Albus, H.G. McCain, and R. Lumia, "NASA/NBS Standard Reference Model for Telerobot Control System Architecture (NASREM)" National Institute of Standards and Technology, Technical Report 1235, Gaithersburg, MD 1989.

A more complete bibliography appears in [1].