

UNCLASSIFIED

Defense Technical Information Center Compilation Part Notice

ADP010329

TITLE: Distributed Intelligence, Sensing, and
Control for Fully Autonomous Agents

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Advances in Vehicle Systems Concepts and
Integration. [les Avancees en concepts systemes
pour vehicules et en integration]

To order the complete compilation report, use: ADA381871

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, ect. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP010300 thru ADP010339

UNCLASSIFIED

Distributed Intelligence, Sensing, and Control for Fully Autonomous Agents

(January 1999)

Dr. Allen Moshfegh and Mr. David S. Siegel
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217-5660
United States of America

1. Introduction

Future naval missions at sea or shore will require effective and intelligent utilization of real-time information and sensory data to assess unpredictable situations, identify and track hostile targets, make rapid decisions, and robustly influence, control, and monitor various aspects of the theater of operation. Littoral missions and operations are expected to be highly dynamic and extremely uncertain. Communication interruption and delay are likely, and active deception and jamming are anticipated.

There is an evolving need for a new generation of unmanned aerial vehicles (UAVs) to perform the tasks traditionally attributed to manned aircraft. For example, UAVs such as Global Hawk are rapidly becoming integral part of military surveillance and reconnaissance operations. UAVs are economical, capable of carrying powerful sensors, and complement manned aircraft missions. Other inherent advantages are (a) removal of personnel from hazardous environments; (b) elimination of error-prone repetitive tasks; (c) reduction of cost associated with operational safety and training; (d) expansion of operational envelope; and (e) performing long endurance mission.

Recent advances in high speed computing, information processing, sensors, wireless communications, Internet technologies, and mobile telecommunications have led to emergence of *network-centric systems*. The technology focus is shifting from individual platforms with limited number of agents to multiple platforms with transparent agents. The software and hardware agents are becoming smarter and capable of continuously adapting to changes in the operational environment. The agents can strategize and make decisions to achieve the desired objectives of mission.

At the Office of Naval Research (ONR) we envision airborne intelligent autonomous agents will have the ability to collect, process, fuse, and disseminate real-time information while exploiting and/or denying an enemy similar opportunities. These airborne intelligent autonomous agents are referred to as unmanned combat

air vehicles (UCAV). This new capability will enhance the notion of network-centric warfare. It is well understood that network-centric operations can deliver to the US military a distinct edge over the enemy. At the strategic level it provides, not simply raw data but a detailed understanding and situational awareness of the appropriate competitive space. At the tactical level, network-centric warfare allows forces to develop rapid response capability and the ability to command and control the littoral environment in real-time settings.

ONR's approach to the development of the unmanned combat air vehicle systems is based on the premise of decentralized intelligence and cooperative behavior in a distributed fashion. The UCAV's decentralized intelligence resides in its organization of its multiple hosts with wide variety of sensing capabilities and functionality that will enable it to protect mission integrity in hostile, uncertain, and spatially extended environment with no single point failure. This organization will be able to accomplish missions that individual agents cannot. This UCAV system of systems organization is composed of: information systems; sensing systems; control and actuation systems; knowledge discovery, learning, and inference systems; planning and decision-making systems; and communications and networking systems.

To date, autonomous agents have extremely limited intelligence and responsiveness (agility and maneuverability) and lack flexibility. Time latency is a major hindrance in the following areas: adaptation to new operational conditions or component failure, learning new tasks, decision-making, and performing cooperative maneuvers.

This paper outlines ONR's conception of cooperative intelligent autonomous airborne agents with application toward intelligent unmanned combat air vehicles. We will describe how our programs are addressing the architectural issues and design techniques needed for the development of the information, connectivity, dynamic networking, communications, intelligent autonomy, and hybrid and intelligent control elements of the vehicle that comprise the envisioned capabilities.

2. Concept of Operation

Figure 1, illustrates a battlefield scenario in which there are several agent teams. There is a ground vehicle




Figure 1. Cooperative UCAV Concept of Operation

team and an air vehicle or UCAV team located in different sectors of the battlefield theater. Two members of the UCAV team are engaged in a strike mission, others in surveillance, and the ground vehicle team is waiting for the opportunity to seize ground control. Though these agent teams may appear to be localized in different sectors with different tasks, they are actually interlocking components commanded by mission control located offshore on a manned control platform. The organization of agents into teams, and the coordination of teams by mission control, transforms a set of agents with localized sensing and actuation capabilities into an organic system that operates over a wide area. Figures 2 and 3, show the hierarchical structure of this organization. Data is shared across layers of the hierarchy and in between peer entities at each layer of the hierarchy.




Figure 2: Multi-Agent Organization

The anticipated UCAV missions are close air support, surveillance, reconnaissance, and strike, see figure1.

The principal objective of the vehicles is to enable time-critical over-the-horizon target detection, identification, tracking, and precision engagement where targets could be stationary or mobile and often in clutter environment. To support these missions, the system of UCAVs will be composed of a set of independent and highly maneuverable platforms that individually will support specialized sensors and some will have weapon deployment capability, but in aggregate provide a robust, survivable, and flexible combat capability. A key feature of the UCAVs is their ability to perform autonomous operation for prolonged periods of time, with multiple options for connectivity to higher authority as required for command, control, and mission retasking.

It is highly likely that the UCAVs will be operating in an actively jammed littoral environment where the lines of communication with human command and control centers are cut off and GPS signal nonexistent. Connectivity outages or lack of GPS signal may last for protracted periods of time, from several minutes to a few hours, nevertheless the UCAVs are expected to




Figure 3. UCAVs Decentralized Hierarchical Architecture

continue their missions safely and reliably until the communication links and/or GPS signals are reestablished, see figure 4. Therefore, the system of UCAVs must be able to self-organize and adjust to unpredictable events while operating in such harsh environments. Consequently, the vehicles must adhere to the most stringent operational requirements for safety and reliability. Following is a partial list of the expected UCAV operational constraints:

- Operate in jammed environment with limited bandwidth;
- Function with incomplete information;
- Navigate without GPS signal;
- Handle unanticipated events;
- Operate in a fault-tolerant and survivable manner;
- Perform new tasks based on real-time information autonomously;
- Operate beyond line of sight;

- Carry lethal payload;
- Engage in lethal operations such as air-to-air strike and air-to-ground strike for suppression of enemy air defenses;
- Maintain connectivity with remote human-decision centers (naval vessels, aircraft, land-based facilities) from which the decision-maker can interact, intervene, and ultimately override various phases of a mission.




Figure 4. Cooperative UCAVs Connectivity in Jammed Environment


3. Information, Connectivity, Dynamic Networking, and Communications

In the ONR conception, as illustrated in figure 5, connectivity and dynamic networking of UCAVs are based on a decentralized hierarchical organization, where the vehicles have varying domains of responsibility at different levels of the hierarchy. Clusters of UCAVs will operate at low altitude (2K-15K feet) to perform combat missions with a focus on target identification, combat support, and close-in weapons deployment. Mid-altitude clusters (15-50K feet) will execute knowledge acquisition, for example, surveillance and reconnaissance missions such as detecting objects of interest, performing sensor fusion/integration, coordinating low-altitude vehicle deployments, and medium-range weapons support. The high altitude cluster(s) (50K-100K feet) provides the connectivity. At this layer, the cluster(s) has a wide view of the theater and would be positioned to provide maximum communications coverage and will support high-bandwidth robust connectivity to manned command and control elements located over-the-horizon from the littoral/targeted areas.

This hierarchical agent organization has architectural features useful for the design of the dynamic network architecture. Higher levels of the hierarchy mostly operate over a greater spatial extent but at slower time-scales. The reason is that the transfer of data over larger spaces usually requires more time, because data transfer requires multiple hops, and in a wireless

environment the reliability of a link can degrade rapidly with increasing range. The bandwidth requirements could be derived from the space-time locus of data. Following are some of the key communication requirements for UCAVs:

- **Secure communication** to deny information to hostile forces. This is particularly challenging because the envisioned strength of the UCAVs stem from their ability to share information and perform distributed information processing and fusion;
- **Low-Probability-of-Detection (LPD)/Low-Probability-of-Interception (LPI)/Anti-jamming (AJ)** capability to penetrate deep into hostile territory. Once UCAVs are detected, hostile forces will attempt to disrupt the UCAV's communication system with jamming techniques ranging from broadband noise to optimum fraction-of-the-band jammers;
- **Dynamic resource allocation:** data quality, high throughput, and high performance, for example, low bit error rate, frame error rate, lost data, and delay;
- **Channel and network capacity:** reliability, redundancy, availability, interoperability of communication links to insure a high degree of connectivity, e.g., alternate transmission routes and multihop communications, in hostile environments.



Functional flexibility and interoperability of the UCAVs are essential to the overall mission effectiveness, that is, loss of individual UCAV or malfunction should only result in marginal degradation of the mission. This self-healing/self-preservation characteristic relies on the autonomy which includes redundant functionality, adaptation, and self-

reconfiguration, as well as robust connectivity of the aggregate system through:

- Distribution and reallocation of essential functions amongst the vehicles in a given cluster;
- Transfer of UCAVs from one cluster to another.

These capabilities can only be realized through adaptable dynamical communication networks allowing reliable, secure, high throughput connectivity [13,14,15]. These networks can be grouped as:

- Intra-network for secure communications among the vehicles within the local network/line-of-sight;
- Inter-network for secure communications between the vehicles in adjacent networks.

Other significant and challenging issues that our program is addressing are as follows:

- (a) Network capacity and resource allocation to perform a specific task or mission. This will depend on the category of the information flow, e.g., command and control, navigation, sensor aggregation, target designation, and network management. The portion of total capacity allocated to each function will vary with mission profile and assigned degree of autonomy.
- (b) Adaptive Communications, UCAV's mission diversity and cooperative networking configurations coupled with the vehicle's dynamics and mobility will demand communication infrastructure that is adaptive and dynamic. Therefore, the architecture must accommodate adjustments to changing channels, network configurations, data requirements, and security. Our focus is on developing adaptive connectivity techniques at various levels of the hierarchy, including the physical layer, network layer, data/information layer, and security layer. In contrast to non-adaptive schemes that are designed relative to the worst-case channel conditions, adaptive techniques, take advantage of the time-varying nature of wireless channels. That is, in adaptive techniques the goal is to vary the transmitted power level, symbol rate, coding rate/scheme, configuration size, or any combination of these parameters in order to improve the link performance which includes data rate, latency, and bit error rates (BER), while meeting the system performance specifications. Adaptive modulation has been shown to increase the data rates on flat-fading channels by a factor of five or more. Additional coding can be used to obtain a reduction in transmit power or BER or resistance to jamming. Moreover, the BER in

adaptive modulation remains constant independent of channel variations, which greatly improves reliability of the wireless link.

- (c) Adaptive Quality-of-Service (QoS), UCAVs will require unique protocols for the QoS. The QoS stands for end-to-end performance metrics for communications link such as bandwidth, latency, and packet dropping probability. Depending on the application, performance metrics defines the minimum requirements needed for good performance. However, for UCAVs, networks are based on dynamic nodes with a dynamic backbone structure. Moreover, network characteristics and applications are mission driven. To secure an acceptable end-to-end performance, the QoS must be adaptive to the network's mission. This adaptation may take the form of variable-rate or multiresolution compression, variable-rate error correction coding, and message prioritization relative to delay constraints, etc.

4. Intelligent Autonomy

Complexity, massive uncertainty, and real-time demands can characterize the operational environment of UCAVs. Crucial elements of intelligence are reasoning, situational awareness, adaptability, learning, decision-making, and contingency planning. Current systems typically lack the ability to learn or to handle unexpected events, either failing, aborting, or referring all such events back to a central human controller. Therefore, UCAVs require a combination of new technologies for sensing, control, learning, communications, and high-level decision making.

Hierarchical structuring is key to the overall design of autonomous intelligent agents. The replication of human optimal decision making process for systems in such UCAV environment, is intractable by the complexity of the task environment. In general, the only way to manage intractability has been to provide a hierarchical organization for complex activities. Although it can yield suboptimal policies, top-down hierarchical control often reduces the complexity of the decision making from exponential to linear in the size of the problem. For example, hierarchical task network planners can generate solutions containing tens of thousands of steps, whereas "flat" planners can manage only tens of steps. The goal is to achieve similar improvements in the ability of the systems to construct complex plans including contingency planning and handling unpredictable events in environments, such as UCAV environment, that are characterized by massive uncertainty.

In both Control Theory and Artificial Intelligence (AI), there is now a consensus that probabilistic and decision-theoretic methods provide a rigorous foundation for optimal decision making in

environments with partial and uncertain sensory data and uncertain dynamics. For example, stochastic optimal control relates directly to AI work on rational agent design. In control theory, online and offline design of control laws is used to address continuous-domain events in environments; in AI, online decision making is used to handle environments with large numbers of discrete variables. ONR's approach is to merge AI and Control-theoretic approaches by developing technology tools for handling the representational and inferential complexity inherent in large, hybrid environments such as UCAV's.

In rational agent design and stochastic control theory, a key concept is known as the belief state: the current joint probability distribution over states of the environment, conditioned on all prior observations [1,2]. With incomplete and noisy sensors, optimal decisions must be computed from the current belief states. In the case of hybrid domains with both discrete and continuous variables, the belief state if explicitly represented grows exponentially with the number of variables. Avoiding this exponential growth is essential. Probabilistic networks (PN) also known as Bayesian networks produce structured representations for complex environments and they are now in wide use for static tasks such as diagnosis, help functions in software products, and situation assessment. Dynamic probabilistic networks (DPNs) extend PNs by including multiple connected copies (called time slices) of a static PN, thereby enabling the modeling of stochastic temporal processes [3]. DPNs serve a number of purposes, see figure 6:

- **Monitoring:** This requires computing the belief state incrementally as new sensor data arrives over time and it easily handles multiple noisy sensors, sensor failure, etc.
- **Prediction:** This requires computing a probability distribution over possible future evolutions of the observed system, and is done by adding slices into the future (this is called filtering).
- **Hindsight:** This requires computing the posterior distribution at any past time given all evidence up to the present time.
- **Decision Making:** By combining prediction with decision nodes representing possible actions by the system itself, one can achieve approximately optimal decision-making with a limited horizon.

The DPNs are expected to model processes that operate at a wide range of time scales. For example, the UCAV must be able to reason about the weather @ 0.01Hz and the behavior of other UCAVs @ 100Hz or manned aircraft @ 10Hz. There is a close relationship between DPNs and Nonlinear Filtering Theory

including Kalman Filtering concepts which are widely used in modern Guidance, Navigation, and Control of dynamic vehicles.

On the intelligent agent architecture the issues of prime importance are: real-time decision-making, adaptation, learning, and hierarchical decomposition. Real-time control is handled by metareasoning and by the integration of multiple execution architectures ranging from compiled control laws to online planning. Adaptation and learning must take place at all levels of the hierarchy, since one can not assume that the environment and correct system structure are known at the outset. This includes learning the environment from sensory inputs, direct learning of control laws in supervised and unsupervised setting, and verification for learning systems – that is proving the resulting systems configuration and the strategy will be effective. Figure 7 illustrates the architecture of an intelligent autonomous agent. Such agents are designed to recognize the inadequacy of their information in an unfamiliar situation and respond by mining available data sources to create new information.




Figure 6. DPN Structure

For example, if a UCAV flying over a sector detects and geolocates a static target that it cannot recognize, it may attempt to augment its information by searching for pre-existing maps that show an object in the same location or ATR logs of other UCAVs' that have passed over the sector, see figure 8. This implies that the information library (see figure 7) should be a shared resource, populated on the basis of the collective experience of the agents, and accessible to all. This is analogous to the way libraries are managed in large institutions. This concept of a heterogeneous knowledge base is a key feature of cooperative agents such as UCAVs. The space-time locus of the knowledge base should track the space time locus of the agents and their data needs.

Real-time decision-making is a crucial capability for UCAVs. However, it is essential to make the following distinctions among decision situations:

- **Low-level open and/or closed-loop control:** Control over actuators during maneuvers, e.g., landing on a moving deck, requires very fast execution. At this level of control, rapid execution is possible because only a few aspects of the environment are relevant and uncertainty is constrained.





Figure 7. Intelligent Agent Architecture

- **High-level precomputed decision strategies:** Some high-level decisions must be made very rapidly, e.g., what maneuvers to execute when faced with multiple incoming threats. These decisions, which deal with a large number of variables and considerable uncertainty, are intrinsically complex and must therefore be precomputed offline. Reinforcement learning, dynamic programming, and genetic/evolutionary learning methods can do this [9,11].
- **High-level decisions in combinations of different circumstances:** It is inevitable that a UCAV will

face some combinations of circumstances that have not been anticipated during earlier offline learning and precomputation. For example, a new mission that requires a new route may require some deliberation before the UCAV can decide which route to start flying. Such decisions need not be made instantaneously; on the other hand, it is simply unacceptable for a UCAV to deliberate for ten minutes. The amount of deliberation must be appropriate to the urgency of the situation and to the value of deliberation for further improvements in decision quality. This can be handled using *rational metareasoning* and composition of *anytime algorithms*.

Rational metareasoning means deciding optimally or nearly optimally which computations to carry out. This can be done by comparing the estimated benefit in terms of improved decision quality with the estimated cost in terms of time (and the implied deterioration of the situation). For example, if the UCAV has decided to return to base because of a serious fuel leak, it is pointless to deliberate further about the location of possibly interesting naval operations in the battle arena. On the other hand, at the beginning of the mission it might be worth spending a minute or two plotting an efficient, safe route and gathering additional intelligence.

Rational metareasoning, along with various other iterative algorithms for generating successive approximations, results in *anytime algorithms* whose decision quality increases monotonically with the amount of computation allocated. UCAVs are expected to contain many such algorithms, e.g., for visual scene interpretation, course computation, weather prediction, cooperative planning with other UCAVs. Thus, it is crucial to be able to allocate computational resources optimally among a large collection of anytime processes.

The principal representation tools for environment models are PNs and DPNs. The PN learning is a local update process using information obtained directly from the inference algorithm. Thus, a simple local update process allows the PN to adapt itself optimally to the environment. This form of learning can be performed offline or online. The DPN learning is similar to PN but it is a dynamic learning process, e.g., the sensor and state evolution models are replicated across time steps.

Reinforcement learning (RL) is the process of learning based on rewards, e.g., short-term payoff information from the environment (useful in UCAV tactics maneuvers). Partially observable environments, which constitute the vast majority of UCAV's missions, require optimal decision-making on the basis of the current belief state. Solving partially observable decision problems is NP-hard, RL can help to reduce

the complexity, e.g., RL focuses on the states arising in the UCAV's actual flight experience. The approach is based on hierarchical reinforcement learning known as hierarchical abstract machine (HAMs). A HAM is a partial specification of behavior that can range from very general, e.g., "fly around over a region of interest, identify moving objects, then come back and land on one of the following ships", to very specific, e.g., "execute the following flight path and maneuvers". Thus, HAMs can be used to place constraints on the behavior of UCAVs, such that the UCAV will execute the optimal strategy that is consistent with the planned mission specifications.

5. Hybrid and Intelligent Control Architectures

Intelligent autonomous systems such as UCAVs are viewed as hybrid multi-agents systems that sense and manipulate their environment by gathering multi-modal sensor data, and compressing and representing it in symbolic form at various level of granularity [6]. The representations are then used by the vehicles to reason and learn about how to optimally interact with the environment. In real world, environments are complex, spatially extended, dynamic, stochastic, and largely unknown. Intelligent systems must also accommodate massive sensory and motor uncertainty and must act in real-time. The hybrid dynamics arise from the interactions between continuous and discrete events and coordination protocols [5,7,8,10,12]. At the continuous level, each agent chooses its own optimal strategy, while discrete coordination is used to resolve conflicts.

The new paradigm that ONR is pursuing for the UCAVs is known as the hybrid distributed hierarchical perception and control. This paradigm is composed of the following key elements:

- Intelligent hierarchical control architectures for autonomous agents that share a single environment;
- Decentralized information and control to maximize a successful and fault-tolerant mission through rapid and dynamic reconfiguration of the inter-agent coordination protocols;
- Perception Systems: (a) hierarchical aggregation of decision and control; (b) wide area situational awareness; and (c) low-level perception.

For safety purposes, the UCAVs are expected to have multiple levels of autonomy and controllability, ranging from teleoperation, to interactive, to fully autonomous meaning that autonomy with intelligence to enable the vehicle to respond rapidly to dynamically changing environments.

It is expected that in a large spatially distributed theater of operations, the sensory systems of individual agents are able to obtain localized and noisy/incomplete information, though mission objectives demand that the agents act quickly, decisively, and cooperatively to optimize mission objectives. One approach is to decompose the process that maps sensory information to control actions in two steps:

- First step is mapping of information about the unpredictable, partially modeled, internal and external environment of the agent into a top-level control decision, which is accomplished through soft computing techniques. These techniques are characterized as goal oriented planning, perceptual reasoning, optimal decision making in stochastic control, and pattern recognition in neural networks;
- Second step is the process that maps the top level control decision to the sequence of control and coordination actions that cascade through the multi-agent system, and ultimately result in the activation of various agent effectors.

For the development of the intelligent control architecture, there is a continuum of design choices for systems decomposition, ranging from strict hierarchical control to a fully distributed multi-agent system. We envision an architecture that allows different choices that are appropriate at different levels of abstraction:

- Continuous domain, for low-level control systems, is concerned with safety and smooth execution;
- Discrete domain, for symbolic and discrete strategic levels, is concerned with optimization and planning for high-level goals;
- Interface and organization of hybrid systems to attain emergent behavior of the collective system of agents for the usage of scarce resources by many agents operating with varying degree of autonomy.

The conceptual underpinning for intelligent, multi-agent systems is the ability to verify that the sensory-motor hierarchies perform as expected. The UCAV will need to have multiple modes of operation, including takeoff, land, track, etc. It is important for the vehicle to verify the modes, e.g., the vehicle should self check the control algorithms that switch between the modes based on high level commands and vision data to prevent the vehicle to enter unstable or unsafe states. In the event of failure or damage, the UCAV must maintain the integrity of the vehicle and safety with possible gradual degradation in the performance of the system.

On the multi-layered hierarchical architecture the higher layers are typically modeled by discrete-event systems, which plan and reason under uncertainty, and assume strategic decisions in coordination with other agents. The lower-layers involve continuous dynamics and perform path planning and regulation tasks. Figure 9 illustrates such a multi-layer hierarchical organization of diagnostics and control layers required for fault management of autonomous vehicles. The hierarchy consists of multiple levels where each level is functionally autonomous. The information flows both ways between the layers while the control commands are passed one way from higher layers to the lower layers. The lower levels of the hierarchy exercise localized control and operate at higher speeds. As one moves up the hierarchy, the domain of influence becomes more global and the decision time cycles grow longer. At each level of the hierarchy, an appropriate world-view can be developed and converted into a model for inference and decision, for example:

- **Vehicle Layer:** Represents UCAV airframe, engines, actuators, control effectors, vision and other sensors, etc. This level provides accurate measurements and assures fast and reliable response of the UCAV to the commands generated by other levels.
- **Regulation Layer:** (a) Adaptive Reconfigurable Flight Control Sub-layer: performs on-line failure detection and identification, control reconfiguration, and signal processing; (b) Autonomous Intelligent Flight Control System Sub-layer: provides trajectory optimization and tracking, and set-point control.
- **Tactical Layer:** This layer executes the plan generated by the Strategic Planning layer. Speed is critical at this level. The main objective of the level is to coordinate the activities of various UCAV missions and dynamically execute tasks such as target assignment, flight mode switching, and trajectory planning.
- **Strategic Layer:** This layer performs autonomous decision making, learning, and verification. It performs threat detection and assessment, and fault management. At this level the supervisor essentially generates long term plans that will result in a successful mission and performs some level of inter-agent coordination. Tasks performed at this layer are computationally intensive.
- **Mission Layer:** This is a mission supervisory layer (such as reconnaissance and surveillance, strike, resupply) and provides human-machine interaction. The supervisor at this level coordinates its mission with other agents in the

network, allocates resources, performs tasks at a discrete level such as route planning, and resource allocation.




Figure 9. Multi-Layer Hierarchical Control Architecture

Intelligent control can be broadly defined as a set of strategies combined in a suitable manner to achieve the desired control objectives in the presence of large uncertainties, fast variations in the system dynamics, and constraints. The emphasis is on large uncertainties and fast variations, which is the main feature of intelligent control systems in comparison with, for instance, robust or adaptive controllers. Intelligent controllers can be designed using Multiple Models, Switching, & Tuning (MMST) technique [4]. The MMST framework is closely related to the intelligent decision-making framework encountered in biological systems. A biological system continuously learns by building different models of its environment and storing this information in the memory. In any new situation, it compares the current information with that stored in memory, based on the model that is closest in some sense to that of the current environment, takes appropriate actions. The MMST concept, shown in figure 10, has been developed using similar ideas. In the figure, models 1 thru n are different operational/event models (observers), while controllers 1 thru n are the corresponding decision-making mechanisms. In the context of intelligent reconfigurable control, the observers are built using linear or nonlinear models associated with different modes of operation (normal mode and failure modes of the system and its components). For each of these models there is a corresponding adaptive reconfigurable controller. The mixing and switching mechanism compares the information obtained from the observers with available measurements and, based on the model that is closest in some sense to the current operating regime of the plant, chooses the corresponding controller.

One of the main features of the MMST methodology is its flexibility. Both observers and controllers can be either fixed or adaptive, and linear or nonlinear. The above structure can be used for identification, estimation, and prediction. It can also be used to estimate the current environment and available resources, and make appropriate decisions. The same structure can be used at higher hierarchical levels where, based on inference methodologies and decision making algorithms, appropriate decisions related to the overall system can be made. Besides its flexibility, the other important feature of the MMST methodology is that in many cases the stability, robustness, and performance of the systems containing MMST observers and controllers can be explicitly evaluated.

6. Summary

We described the envisioned architectures and technology capabilities needed to develop a robust, adaptive, and dynamic communications, information, decision, and control infrastructure supporting a class of cooperative and distributive intelligent autonomous air vehicles or UCAVs. This infrastructure will be used to organize the mobile agents into dynamic teams supporting complex missions such as surveillance, target detection and tracking, and coordinated attack. The dynamic network provides connectivity for real-time situational awareness in the littoral environment and exchange of real-time information between agents. The connectivity requirements of the underlying mission include real and non-real time data transfer as well as data retrieval from fixed and mobile databases. This communications infrastructure will provide a unique and powerful mechanism for the Navy to carry out complex missions with minimal risk and vulnerability.




Figure 10. Multiple Model Switching and Tuning

7. References

1. R. Parr and S. Russell, "Approximating Optimal Policies for Partially Observable Stochastic Domains" Proc. Fourteenth Intl. Joint Conf. On AI, Montreal, Canada, 1995.
2. S. Russell, "Rationality and Intelligence", Artificial Intelligence, 94, 1997.
3. J. Binder, K. Murphy, S. Russell, "Space-Efficient Inference in Dynamic Probabilistic Networks", Proc. Fifteenth Intl. Joint Conf. On AI, Nagoya, Japan, 1997.
4. K. S. Narendra and J. Balakrishnan, "Adaptive Control Using Multiple Models", IEEE Trans on Automatic Control, Vol. 42, No. 2, Feb. 1997, pp. 171-187.
5. J. Lygeros, D.N. Godbole, S. Sastry, "Hybrid Systems III, chapter A Game Theoretic Approach to Hybrid System Design", Springer Verlag, New York, 1996.
6. J. Lygeros, D.N. Godbole, S. Sastry, "Optimal Control Approach to Multiagent, Hierarchical System Verification", In Proc. of the IFAC World Congress, 1996.
7. J. Lygeros, D.N. Godbole, S. Sastry, "A verified hybrid controller for automated vehicles", In Proc. of Control and decision Conference, 1996.
8. D.N. Godbole, J. Lygeros, S. Sastry, "Hybrid Systems II, Chapter Hierarchical Hybrid Control: A Case Study", Springer Verlag, LNCS, New York, 1995.
9. D.E. Goldberg, "Genetic algorithms in search, optimization, and machine learning", Addison-Wesley, Reading, MA, 1989.
10. P.J. Antsaklis, J.A. Stiver, M. Lemmon, "Hybrid System, chapter Hybrid System Modeling and Autonomous Control Systems", Springer Verlag, Series LNCS, No. 736, New York, 1993.
11. Davis, "Handbook of genetic algorithms", Van Nostrand Reinhold, New York, 1991.
12. Deshpande, "Control of Hybrid Systems", Ph.D. thesis, University of California, Berkeley, 1994.
13. T. S. Rappaport, Wireless Communications Principles and Practice, IEEE Press/Prentice-Hall, 1996.
14. J. D. Parsons, The Mobile Radio Propagation Channel, Wiley, 1992.
15. J. B. Andersen, T. S. Rappaport, and S. Yoshida, "Propagation Measurements and Models for Wireless Communications Channels," IEEE Comm. Mag., Jan. 1995, pp. 42-49.