

Agent-based Simulation and Support of Command, Control, and Communication Decisionmaking Utilizing a Distributed Research Network

Christopher Barnes

Warfighter Training Research Division
Air Force Research Laboratory
2504 Gillingham Drive
Brooks AFB, TX 78235-5100
(210) 536-2177
christopher.barnes@brooks.af.mil

Dr. Plamen V. Petrov

21st Century Systems Inc.
12612 Decatur Street
Omaha, NE 68154
402-384-9093
Plamen@21csi.com

Dr. Linda R. Elliott

Veridian Engineering
2504 Gillingham Drive
Brooks AFB, TX 78235-5100
210-534-1134 x 19
linda.Elliott@brooks.af.mil

Abstract: This paper describes approach, methodology, and potential application areas for agent technologies in C2 performance and primarily focuses on the development of agent-based constructed forces. The platform involved is the AWACS-AEDGE™. This is a distributed, real-time team decision support environment comprised of simulators, entity framework, intelligent agents and user interfaces. The AEDGE is constructed as a federation of intelligent agent-based functions that enable user-friendly scenario construction, emulation of friendly and hostile entities, and dynamic scenario control. Its architecture and decision making algorithms are examined, as well as agent technology and utilization in the realms of constructive forces, synthetic team members, and decision support.

Introduction

This paper describes the approach, methodology, and potential application areas for agent-based synthetic task platforms (STE) to enhance C2 performance, as exemplified in the AWACS-AEDGE™. We focus particularly on the development of agent-based constructed forces and decision support to enhance C2 training and performance (Barnes et al, 2002). The AWACS-AEDGE was developed to represent core characteristics of the Airborne Warning and Control System (AWACS) Weapons Director (WD) team.

The AWACS WD team serves as a vital airborne Command and Control (C2) node, providing airborne surveillance, control, and communications functions for tactical and air defense forces. It is clear that AWACS WD duties exemplify core characteristics of a C2 team. They perform in highly interdependent roles, tracking and coordinating some type of tactical action, in a manner consistent with overall strategic goals and procedures, for a defined sector of air and/or land space over a sustained period of time (Elliott et al., 2001). To accomplish this, they exchange, interpret and effectively weight information and optimize resource allocation decisions across team members. As a prime example of a command, control, and communications paradigm, the AWACS team is an area where C2 research has often centered.

The Air Force has recently focused on investigations and enhancement of operational expert training through an internet-based research infrastructure to enhance Distributed Mission Training (DMT). The USAF DMT program is national in scope, with the goal of enhancing operational training through the use of high-fidelity military simulation systems that are networked using secure, classified systems. The DMT-Research Network (DMT-RNet) project, the subject of this paper, is a local project that will support DMT through basic research accomplished using PC-based systems networked through the Internet. The project will establish an I-2 based infrastructure for collaborative research and training, along with identification of specific research issues related to enhanced skill acquisition and operational performance. This research will guide improvements in the operational USAF DMT training environment.

The USAF DMT project relies on a network of highly realistic battle simulators that allow expert operators to train in a virtual battlespace across a highly secure and classified communication network. In contrast, the DMT-RNet project is developing less expensive PC-based systems that can run in unclassified mode on I-2. This allows the PCs to be distributed and deployed as training systems in almost any setting. It also allows research to be conducted in nonclassified environments, using these simulation systems that reflect essential components of operator expertise. For example, multiple universities and research companies can and have been networked to enable real-time multidisciplinary collaborative research.

As a pioneering technology, DMT-RNet shows how other collaborations in research will be possible. Other universities, agencies, and companies will be able to link up to similar sorts of networks and pool their talents and resources to produce high level research. The applications of these concepts and technologies to other realms are limited only by imagination, initiative, and support (Barnes, Elliott, & Entin, 2001).

The initial phase of the DMT-RNet project utilized the dynamic distributed decision making (DDD) team-in-the-loop simulation environment (Hess, MacMillan, Elliott, & Schiflett, 1999; Kleinman and Serfaty 1989). An internet-based version of the DDD was developed, the DDD Network (DDDnet), which allows players in distributed locations to connect and perform a distributed mission in real time. The DDDnet is an internet-ready version of a Linux-based collaborative gaming space that connects players to each other and to others, such as observers, confederates, trainers, or researchers. In the DDDnet observers at any location in the network are able to observe the scenario play in real time. They can view the screen display and electronic communications of any player, and communicate to one another via email or voice. In addition, the DDDnet can connect players to one another for interactive mission planning, debriefings and after-action reviews.

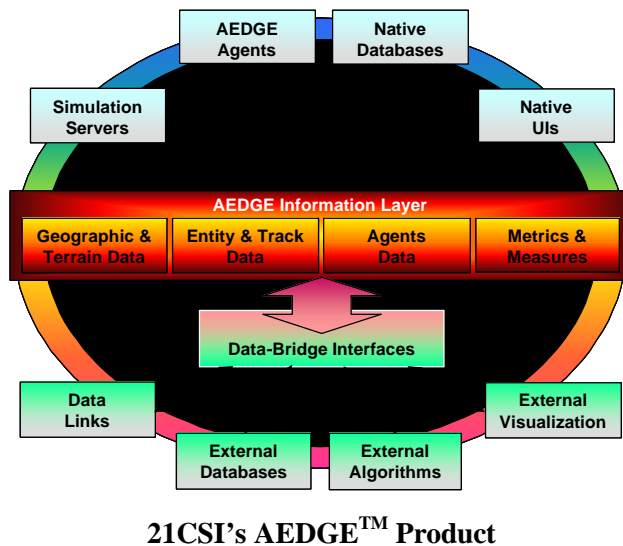
The second platform developed is the AWACS-AEDGE™ (Agent Enabled Decision Guide Environment). The AEDGE is constructed as a federation of intelligent agent-based functions that enable user-friendly scenario construction, emulation of friendly and hostile entities, and dynamic

scenario control. The remainder of this paper goes into the detail of this platform and the agent technology that it employs.

It has been stated that advanced technology in itself is not a training system and should not be seen as a training solution (Salas, Cannon-Bowers, & Kozlowski, 1997). We are in complete agreement. The purpose of this effort was to build a platform with advanced-technology features that will augment, enable, and/or investigate new approaches, methods and measures in training and performance research.

AEDGE™¹

The Airborne Warning and Control System (AWACS) Weapons Director (WD) team serves as a vital airborne Command and Control (C2) node, providing airborne surveillance, control, and communications functions for tactical and air defense forces. In this paper, we describe development of an agent-based C2 team decisionmaking platform for research and training, the AWACS-AEDGE™ (Agent Enabled Decision Guide Environment).



The AWACS-AEDGE, built using 21st Century Systems Inc.'s AEDGE™ infrastructure, is a distributed, real-time team decision support environment comprised of simulators, entity framework, intelligent agents and user interfaces. The environment supports a wide variety of air, sea (surface and sub-surface), and ground assets in a combat environment (Chiara & Stoyen, 1997), primarily based on the roles and responsibilities of AWACS WD team members. The environment has been tested with an excess of two hundred physical entities (planes, ships, SAM sites, etc.) operating with realistic yet non-classified performance characteristics in an interactive

environment in which real-time decision support is available to each WD.

The behavior and decisionmaking of all hostile and friendly entities not controlled by humans is directed by agent-based technology. If a human decides to “log in” as a particular entity, he/she may choose to view recommendations generated by the agent for that entity. Even if the human operator chooses not to view recommendations, the agent recommendations are still logged by the computer. This enables direct comparison of human to agent decisionmaking. We expect these capabilities to facilitate skill acquisition, decisionmaking, skills assessment, and human/team performance modeling.

AEDGE agent capabilities enable more detailed and innovative approaches to measurement and modeling of individual and team workload, communication and decisionmaking. Tracking the number and type of recommendations generated by the agent at any given time contributes toward new ways of conceptualizing and representing cognitive workload of individuals and teams. Agent-based recommendations may also serve as a standardized benchmark by which human tactics and

¹ AEDGE, AWACS-AEDGE, and 21CSI are registered trademarks of 21st Century Systems Inc.

decisions can be compared. In addition, the AEDGE platform can operate through speech – operators can speak to the system using predefined jargon, request tasks be performed or information provided/transferred, and the agents will respond verbally to the speech-driven requests, using voice generation technology. All agent communications to each other, as well as to humans, are transcribed, logged to data output files, and are available online.

The AWACS-AEDGE was conceived through cognitive and functional analysis of team member roles, responsibilities, and decisionmaking (Dalrymple, 1991), to optimize generalizability of results to operational settings. Systematic descriptions of AWACS roles, responsibilities, requirements, interdependencies, tactics, strategies, and task demands were collected from subject matter experts, cognitive task analyses (Fahey et al., 1998; MacMillan et al., 1998) and focal-group interviews (Elliott et al., 1999; Elliott et al., 2001). These data were examined to identify decision events, which were generic to performance, regardless of mission scenario, and likely to bottleneck under high tempo situations.

AEDGE architecture

The AEDGE product is based on an extensible distributed component-based architecture, which defines entities, agents, players/users, and their interactions. The interaction and communications among AEDGE elements are based on the Service Provider/Service Requester protocol (SPSRP), using flexible services and messages to exchange information among any two components.

In SPSRP, Service Providers implement a number of services and register service-templates with a Component Registry, which maintains the location of all components and the services provided or required by each. The registry is used by Service Requesters to locate components that provide the services required by the requester. After the requester is matched with one or more Service Providers, a direct connection is established between Service Providers and Service Requesters. This prevents the Registry from being a service dispatcher and a potential bottleneck. Let us consider the interaction between users (via user interface components) and agent components. In most cases users will interface with specialized Agents, called Agent Managers. The managers are designed to coordinate, synchronize and manage the work of multiple “worker” agents. Without an Agent Manager, the user will need to interface with each “worker” agent individually, while using the Manager, the user is able to issue higher-level requests (e.g. “Send me your current recommendations”) by letting the Agent Manager (who knows the capabilities of its workers) to distribute and correlate individual agent tasks.

The Service Requester then sends Service objects directly to one or more Service Providers, who respond with ServiceResult objects. A Requester may wish to subscribe for service updates, in which case, the Service Provider will send a Message object to the requester every time it needs to advertise an update; it is up to the Requester to respond to that message by requesting the actual update (i.e. data is advertised, not pushed, to avoid client-side congestion).

On figure 3, we see that the user may trigger a request for recommendations, which is sent to the Agent Manager via a Service request object. After receiving the request, the Agent Manager finds the best-suited collection of agents to perform the job (it may take one or more worker agents) and forwards specialized Service requests to each of them. The worker agents formulate their responses (usually sets of recommendations and rationale) and send them back to the Agent Manager via ServiceResult objects. The Agent Manager may correlate (and even remove redundant or inconsistent recommendations) all service results and then sends the combined set of

recommendations to the user, again via a ServiceResult object. The user interface component then knows how to extract the recommendations and present them to the human user for evaluation.

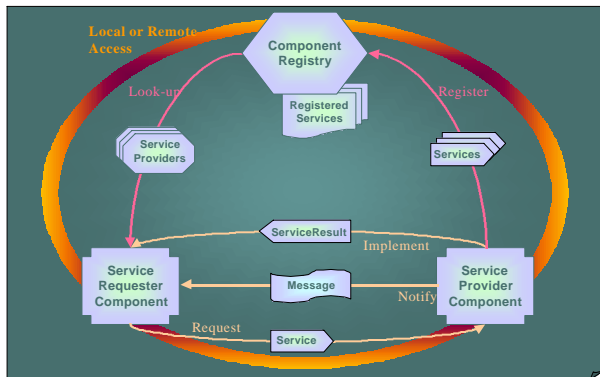


Figure 2. SPSRP protocol interactions

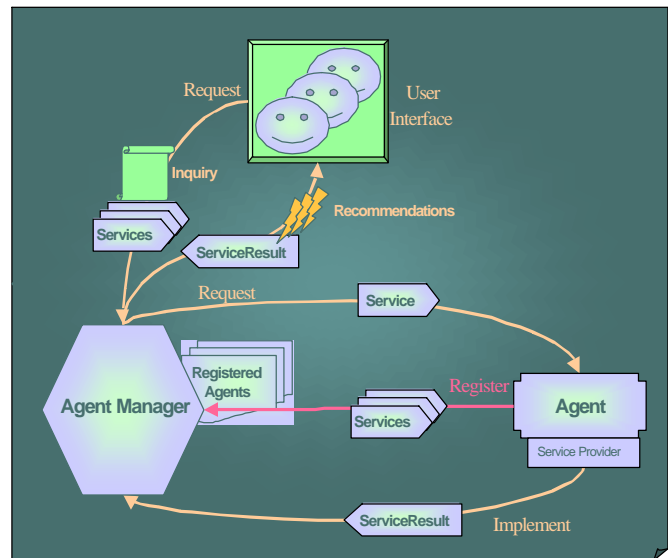


Figure 3. User Request

AEDGE decision making algorithms

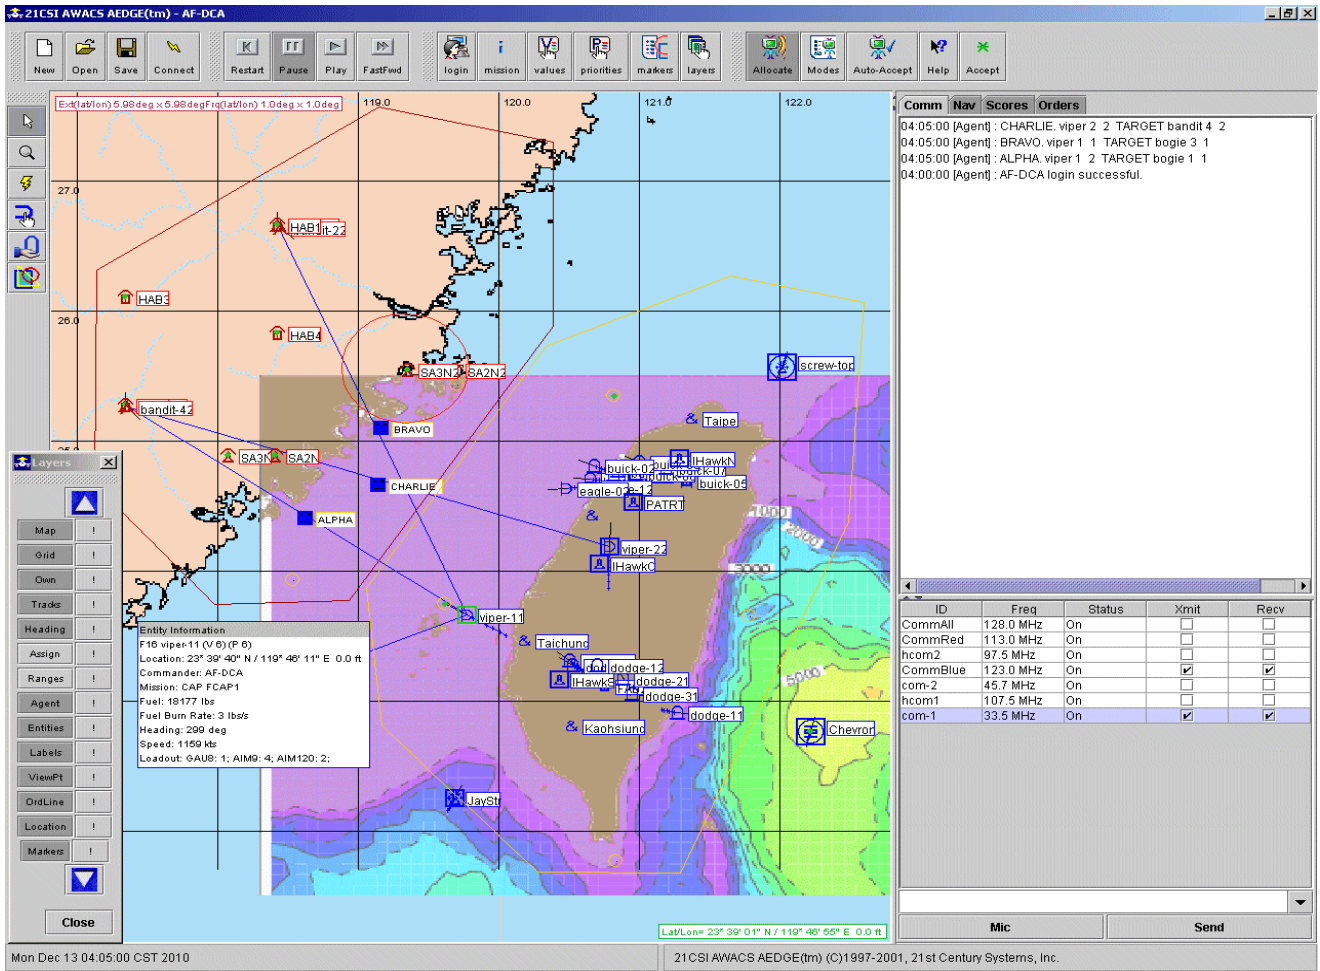
The AEDGE architecture provides multiple levels of agent-based algorithms. Generic resource allocation, search and optimization algorithms are a core part of the AEDGE product. Each AEDGE application can use and further extend these fundamental agent algorithms by either providing parameters and applications-specific values, functions and rules, or by combining, modifying or supplying new algorithms. All new and modified algorithms must comply with a well-defined agent interaction interfaces, similarly to the generic algorithms.

The AWACS-AEDGE extends resource allocation, optimization and other algorithms with AWACS/WD-specific objective functions and constraints. For example, the AWACS weapon-target allocation algorithm, based on a generic resource-allocation with heuristic function evaluation, defines extended constraints such as Table 1.

Similarly, the AWACS weapon-target allocation algorithms define objective and cost functions for any potential allocation and let the generic allocator agent arrive at a (near-)optimal set of weapon-target pairings. The objective functions are based on the individual target values (as well as other factors, such as target priorities, probabilities for success and so forth) and cost functions are based on the risk for the team if the allocation is to be committed.

Further, the AWACS-AEDGE agents use the extended algorithms as a model of the desired WD performance. Thus, the agents are able to generate a set of recommendations pertinent to a particular tactical position and the events that lead to it – agents keep event history. Such recommendations can either be presented to the user (who may choose to accept or ignore them) or be used for internal evaluation of the user’s performance as a function of the similarity of recommended-action versus actually-executed-action. A new application of the agent

recommendation analysis involves the measurement of AWACS WD cognitive workload based on the volume and complexity of agent recommendations at any given time (Chaiken et al, 2001).



```

IN_RANGE(TargetCandidate,
WeaponCandidate ) OR
INTERCEPT_TIME( WeaponPlatform,
TargetCandidate ) < MAX_TIME
AND
Pk( TargetCandidate, WeaponCandidate )
> DesiredPk( TargetCandidate )
AND
FUEL_TO_INTERCEPT(
WeaponPlatform, TargetCandidate ) +
FUEL_TO_BASE( InterceptPoint,
WeaponPlatform ) < WeaponPlatform.currentFuel

```

Table 1: Decision Making Algorithms

To enhance the utility of AWACS-AEDGE as a decision support tool, recommendations must not only be presented to the user in a unambiguous and intuitive manner, but in some cases they may need to be pre-processed to ensure that the human user can maintain strong situational awareness and be alert. For example, in periods of exceptionally high-activity, unsupervised agents may tend to generate large number of recommendations that will be confusing and even detrimental to the human performance. The Agent Manager must thus not only coordinate and synchronize recommendations, but also prioritize and reduce the number of presented recommendations to only the top-most critical ones (Chaiken et al, 2001). Conversely, in periods of lull, artificially increasing the number of recommendations may help keep the human alert and situationally aware.

Agent Technology

Intelligent agent technology is rapidly demonstrating its value to operational simulation and training. Within the AEDGE, "agent" is a broadly defined term with three dominant functions. The first involves the simulation of scenario entities, referred to as constructed forces. Related to that is the use of agents to substitute for other human roles, to create "synthetic" team members. The third function is that of decision support.

Constructive Forces. Agent technology defines the operating characteristics and behavior of hostile and friendly entities (e.g. speed, radar range, and weapons range of aircraft). Agent specification is complete and detailed to the point where the entire scenario can be played out through a federation of numerous "agents" (i.e., a simulation with no live players).

Synthetic team members. Agent technology was also applied to simulate other WD team members within the AEDGE. The distinction between "agent as constructive forces" and "agent as player" is largely one of degree. However, the latter sort of agent is typically far more complex, as these agents are designed to simulate another simulation player, not just a battle entity. Such agents give users the option to play with other "live" participants or participate alone, with the simulation acting out other roles in a realistic fashion. This sort of agent also defines the pedagogical goal of the simulation, in the sense that these agents can be used to implement (e.g. set policy for) optimal performance. They can also be used to demonstrate the results of flawed performance. "Player" agents are an extremely useful and yet rare capability for team task simulations. However, their development is expected to increase given the great utility of allowing individual training within a team-like context. Both entity and player agent technology are equally important to our effort. The former sense is what gives the AWACS-AEDGE its fidelity to the real task; the latter sense provides both the model of normative behavior the user should strive for and the means (algorithms/knowledge) to effectively accomplish that behavior.

Decision Support. The third manner in which agent technology was utilized is to provide decision support. This type of agent is not so much a simulation of a player (or simulation of an entity) but a simulation of a "coach" or "adviser". This variant of agent can be very broad, and the distinction between such agents and operational interfaces can be blurry because such agents manifest themselves through the interface. For instance, such agents can be imbued with the ability to seek out information over distributed networks; search through information databases; manipulate information through filtering, transforming, aggregation, and fusing of multiple, independent information streams; and to report information to the human requester. There may also be multiple agents working on several tasks at any point of time, e.g., several agents monitoring and filtering information from disparate channels, agents to aggregate and fuse relevant information, agents to select an appropriate visualization of the data to report, and so on. Some agents may be

imbued with a high level of autonomy, allowing them to make critical recommendations based on information found without human solicitation or guidance.

Utilization

In the AEDGE, the experimenter can control the autonomy and configuration of the decision aide agent. If the agent is allowed to make all decisions, the scenario is effectively being run independently of any human intervention. This allows (a) assessment of reliability of recommendations, (b) assessment of effects of uncertainty in a dynamic environment, and (c) investigation of “what-if” scenarios, where algorithms underlying recommendations are manipulated.

The aide is, and should be, configurable. First, it needs to be configurable in order to maintain effectiveness. The manner in which decisions are made can and will change according to the particular mission scenario and rules of engagement. The software allows variations in its rule structure. It also allows changes (in real-time) to the perceived value and priority of various assets and targets. This enables fine-tuned research in decision process, as the agent can be tailored to be more or less risk-taking (when information is uncertain), have directional bias (more or less “aggressive”, “passive” in threat assessment or rules of engagement), or bias in central tendency (decisions are always “moderate”). In addition the probabilistic nature of the environment can be manipulated by specifying the probability that the decision made will be successfully executed. When that probability is very high, the environment is deterministic and very reliable. When probabilities are lowered, scenario events will unfold in different ways, each time the scenario is run.

Configurable decision algorithms enable in-depth descriptive and prescriptive investigations of decision process. Particular heuristics, biases, and models of decision choice can be predicted and compared to algorithm function. Results from descriptive investigations can inform refinement of the decision tool. In turn the algorithm can be modified to reflect a particular decision model, and compared to other models with regard to the degree to which either model accounts for performance data. For example, threat assessment decision events have been shown to be sensitive to order effects, in that information presented first or last (depending on tempo) is given more weight, even when other information is more important. The algorithm can be adjusted to reflect this tendency, and results compared to actual data. Other facets relevant to the decision process can also be investigated with this approach, such as risk-taking, aggressiveness, and information uncertainty.

Discussion

Despite agent limitations, we expect the AWACS AEDGETM to enhance research, training, and performance in complex high-tempo scenarios. The aide as decision support is particularly useful where multiple decisions must be made within a short time frame or where an “out-of-the-ordinary” event must be spotted among numerous seemingly normal processes. The aide has the advantage of complete reliability regardless of stress, sustained operations, or consequences of failure. It will never forget to refuel a plane because of a tense situation occurring at a different part of the scope.

While the usefulness of the DSS function is apparent, the potential utilization of this platform for training and performance research is its greatest asset. The benefits of this general

approach to STE-based research is detailed elsewhere (Schiflett & Elliott, 2000). To summarize briefly, the AWACS AEDGE™ was developed to primarily to support trainers and researchers. In fact, every characteristic and feature within this platform was developed in order to empower trainers and researchers with regard to methods, measures, manipulations, and transfer of training.

First, internal validity is enhanced by providing researchers with more detailed performance measures, increased scenario realism, ease in generating scenario events, agent-based performance models, and comprehensive data output files. Cognitive task analyses augment operational relevance of assets, events, and decisions. Advanced scenario generation features enable rapid configuration of scenario assets, events, and role assignments. Communication and information flow networks can also be configured toward research goals. Characteristics of information (e.g. accuracy, ambiguity, tempo) and decision rules (e.g., priorities, risk, rules of engagement) can be configured to further define behavior of agent-based forces and investigate impact on human decision making processes.

The creation and refinement of synthetic team members allow investigations of individual performance within more highly controlled team contexts. Synthetic team members can be programmed to different decision rules, thus enabling the creation of team members with more or less expertise, decision bias, and/or risk aversion. One line of research will investigate the impact of fatigue on individual and team performance, such that synthetic forces can represent performance at various levels of fatigue.

The configurability of the decision support system is consistent with that of the synthetic forces. The DSS features can be programmed to provide recommendations toward specific decision events, as determined by the experimenter or trainer. This enables investigations of DSS attributes and operator utilization and trust in the system.

AEDGE also includes provision of online scenario revision capabilities and configuration of visual online performance feedback for operators. Data logging is extensive, providing several different files targeted toward mission outcomes, team communications, and agent recommendations. Recommendations are logged regardless of DSS use, such that human performance can be compared to agent performance, for each significant decision.

While use of this system is no guarantee of good training or research per se, we hope it will accomplish its purpose—to provide tools that empower experts to more easily accomplish research and/or training goals.

References

- 21st Century Systems Inc (2000). AWACS Weapons Director Trainer Tool: Phase II Final Report. Brooks AFB: Air Force Research Laboratory.
- Barnes, C., Elliott, L. R., & Entin, E. (2001). Employing Internet2 Technology to Enable Collaborative Research and Distributed Training in Complex Multi-Operator Settings. WebNet Journal, 3 (4), 24-31.
- Barnes, C. M., Petrov, P. V., Elliott, L. R., & Stoyen, A. (2002). Agent based simulation and support of C3 decisionmaking: Issues and opportunities. Proceedings of the Conference on Computer Generated Forces and Behavior Representation, May 2002.

- Chaiken, S., Elliott, L., Dalrymple, M., & Schiflett, S., Coovert, M., Riddle, D., Gordon, T., Hoffman, K., Miles, D., King, T. (2001). Weapons Director intelligent agent-assist task: Procedure and findings for a validation study. Proceedings of the International Command and Control Research and Technology Symposium, June, 2001.
- Chiara, J. & Stoyen, A. (1997). Intelligent, adaptive research and training platform for individual and team training, with web capabilities and applications. Phase II SBIR Report AF971-018. Brooks AFB, TX: Air Force Research Laboratory, Warfighter Training Research Division.
- Dalrymple, M. A. (1991). Evaluating Airborne Warning and Control System strategy and tactics as they relate to simulated mission events. (AL-TP-1991-0049, AD A242820). Brooks AFB, TX: Armstrong Laboratory.
- Elliott, L. R., Cardenas, R., & Schiflett, S. (1999) Development and Analysis of Team-based Rating Scales Of AWACS Performance in Distributed Mission Training. Proceedings of the Annual Research and Technology Symposium on Command and Control, Newport, RI. June, 1999.
- Elliott, L. R., Hollenbeck, J. R., Schiflett, S. G., & Dalrymple, M. (2001). Investigation of Situation Awareness and Performance in Realistic Command and Control Scenarios. In M. McNeese, E. Salas, & M. Endsley (Eds.) "Group Situational Awareness: New Views of Complex System Dynamics". HFES Press.
- Fahey, R. P., Rowe, A., Dunlap, K. and DeBoom, D.(in review). *Synthetic Task Design (1): Preliminary Cognitive Task Analysis of AWACS Weapons Director Teams*. Technical Report. Brooks AFB, TX: Armstrong Laboratory.
- Hess, S., MacMillan, J., Elliott, L., & Schiflett, S. (1999). Team-in-the-Loop, Synthetic Simulation: Bridging the Gap Between Laboratory and Field Research. Proceedings of the 43rd Annual Meeting of the Human Factors and Ergonomics Society, Houston, TX; 1999.
- Kleinman, D. L. & Serfaty, D. (1989). Team performance assessment in distributed decision-making. Proceedings of the Symposium on Interactive Networked Simulation for Training (pp. 22-27), Orlando, FL.
- MacMillan, J, Serfaty, D., Young, P., Klinger, D., Thordsen, M., Cohen, M., & Freeman, J. (1998). *A system to enhance team decision making performance: Phase I Final Report*. Brooks AFB: Air Force Research Laboratory, Warfighter Training Research Division.
- Petrov, P. V., Stoyen, A. D. (2000). An Intelligent-Agent Based Decision Support System for a Complex Command and Control Application. In Proceedings of Sixth International Conference for Engineering of Complex Computer Systems (ICECCS'2000).

Salas, E., Cannon-Bowers, J., & Kozlowski, S. (1997). The science and practice of training— Current trends and emerging themes. In J. K. Ford, S. Kozlowski, K. Kraiger, E. Salas, & M. Teachout (Eds.) *Improving Training Effectiveness in Work Organizations*.

Schiflett, S. & Elliott, L. (2000). Development of Synthetic Team Training Environments: Application to USAF Command and Control Aircrews. In H. O'Neil & D. Andrews (Eds.) *Aircrew Training: Methods, Technologies, and Assessment*. LEA Press.