TABLE 1 — Densities Compared for Some Candidate Materials

Material	Areal Density kg/cm^2	Density kg/m^3
Aluminum		2900
Silicon Carbide	32.0	3200
CFRP	5.5	1600
ULE Glass	10.0	2210
Zerodur Glass	45.0	2530

CFRP is similar to that of steel. This property, coupled with the lower mass of a CFRP structure, means that telescope components rapidly equilibrate to their surroundings, which reduces image distortion. Table 1 compares areal densities and volumetric densities for some candidate materials.

Candidate Implementation: When we combine coatings, composites, and unique shared apertures, we can eliminate one primary at a minimum. We further decrease mass by replacing the antenna/telescope assembly structure with composites and enable larger, lighter weight optical telescopes that can also support K_a radio transmissions. When applied to the 3-m high gain antenna (HGA) used in the Mars Reconnaissance Orbiter, we can offer an efficient 3-m telescope at each end of a Mars-Earth orbiting link that can support the commensurate increase in data rates for both K_a and SWIR-based communications. Top-level link analysis suggests that the data throughput can be increased over baselined X/K_a-band of 6 Mbps (near Earth) and 500 kbps (furthest from Earth) to multiple Gbps to ~300 Mbps using SWIR. Further efficiencies can be achieved with smart modems that would sense when the optical link is in play and put the RF link in hibernate mode to reduce power draws.[2] This scenario will result in a smaller power loading requirement, further reducing SWAP for the system. Pointing, acquisition, and tracking is assumed to be ideal for these trades. A more sophisticated treatment of the links can be found in Ref. 3.

Implications: The results of this study have been combined with the analysis of other national labs and are being considered by NASA for a Pre-Flight Phase A Study for future Deep Space Missions.

[Sponsored by NASA]

References

[1] G.C. Gilbreath, M.A. Rupar, C.O. Font, and R.C. Romeo, "Combined RF and EO Aperture Study," NRL Report to NASA, September 2010; NRL Formal Report in preparation.

[2] D.S. Kim, G.C. Gilbreath, J. Doffoh, C.O. Font, and M. Suite, "Hybrid Free-space Optical and Radio Frequency Switching," *Proceedings of the SPIE* **7091**, 70910X–70910X-9 (2008).

[3] *Near-Earth Laser Communications*, H. Hemmati, ed., CRC Press (2009).

■

Goal-Driven Autonomy

D.W. Aha,[1] M.L. Molineaux,[2] and M. Klenk[3]
[1]*Information Technology Division*
[2]*Knexus Research Corporation*
[3]*Palo Alto Research Center*

Introduction: The Navy needs autonomous intelligent agents to help dominate the battlespace in several types of missions. For example, agents could reduce the amount of warfighter oversight required to operate unmanned vehicles, serve as proactive decision support assistants in C2 systems, and increase the realism of decision-making behaviors for simulated units in training and wargaming systems. However, current generation agents are severely limited; they must be told how to behave in all situations that they will encounter, or be supervised continuously. Unfortunately, such knowledge engineering is impossible for their complex military environments, which are characterized by incomplete knowledge of the situation, dynamic situation updates, multiple adversaries, and stochastic agent actions. Thus, current agent deployments are highly constrained (e.g., they are programmed to withdraw and report to their operators when any unexpected situations occur).

To achieve full autonomy, agents should dynamically reason about what goals they should pursue to optimize mission performance measures. This will allow them to respond intelligently to unexpected situations as the battlespace situation evolves. We refer to this process as goal reasoning, and next describe our progress on defining a model for it, an agent implementation, and its evaluation.

Model: Figure 3 displays our novel computational model of goal reasoning, also called goal-directed autonomy (GDA). GDA permits agents to select their goals throughout their deployment (under mission constraints). It extends the controller of the existing model for online planning, which envisions (1) a plan generator, (2) a controller that feeds the plan's actions to the environment, and (3) the environment itself, which uses a function to compute state updates from executing specific actions. A GDA agent monitors a plan's execution, and continually compares the current and expected situations (or states). If a discrepancy is detected, then it will generate an explanation of its cause. This is given to a discrepancy resolution module, which may decide to formulate a new goal in response. Finally, a goal manager prioritizes and selects a set of (potentially new) goals to pursue, which are given to the planner to create new plans to execute.

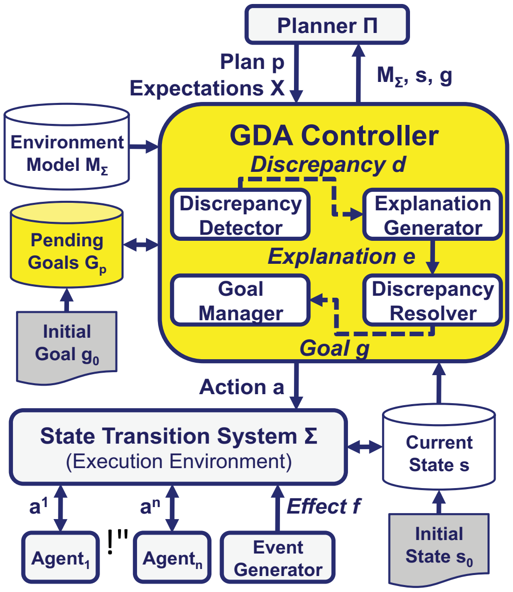

FIGURE 3

Conceptual model for Goal-Driven Autonomy.

GDA is a general model; many inferencing algorithms could be used for each GDA step. Below we describe one of our instantiations and its empirical investigation.

Implementation and Evaluation: Can the GDA model's additional complexity yield quantifiable benefits? Our hypothesis was that when scenarios produce states that require a GDA agent to change its goals to perform well, it should outperform agents that cannot change their goals. We assessed this by implementing this model in the ARTUE (Autonomous Response to Unexpected Events) agent and comparing its performance vs ablated versions in scenarios defined using complex simulation environments.[1] ARTUE is a simple GDA instantiation; it defines discrepancies as any differences between the expected and actual states, it uses an augmented truth-maintenance system to generate explanations, a set of simple rules to trigger the formulation of goals with specified priorities, and always directs planning to the goal with the highest priority.

For some of our studies, we tested ARTUE using the Tactical Action Officer (TAO) Sandbox, which is a comprehensive littoral simulator used to train TAOs at the Surface Warfare Officers School. We modified the Sandbox such that ARTUE assumes the trainee's interactive role and controls a set of assets (e.g., ship, air, sensors) in antisubmarine and related mission scenarios. For example, the SubHunt scenario involves a search for an enemy submarine that has been spotted nearby. A ship is dispatched to locate and engage it. However, the submarine has been laying mines that can incapacitate the searching ship. As with each scenario, we defined a scenario-specific performance measure, in this case one whose value increases as a function of finding and destroying the submarine, as well as sweeping the mines. A non-GDA agent would ignore the mines because it would focus solely on its initial goal, whereas ARTUE would formulate a goal to sweep them. Table 2 displays our results for ARTUE and its ablations (i.e., which perform no GDA steps, only discrepancy detection, or only the first two steps, respectively). Analysis reveals significant performance differences between ARTUE and its ablations for all these scenarios.

TABLE 2 — Mean Performance Scores for the TAO Sandbox Scenarios

N = 25	Scouting	Iceberg	SubHunt
PLAN1	0.33	0.35	0.35
REPLAN	0.40	0.48	0.48
EXPLAIN	0.58	0.64	0.74
ARTUE	0.74	0.73	0.98

Discussion: Our evaluations on the TAO Sandbox and other simulators involving multiple GDA agents have been encouraging.[2] However, many open research issues exist. For example, ARTUE requires a set of rules for triggering goal formulation during discrepancy resolution. We recently extended ARTUE to interactively acquire these rules and will soon study methods for automatically identifying events of interest to discrepancy detection. In an upcoming project, we will transition extensions of ARTUE for use in ICODES, a deployed system for ship cargo loading/unloading. Future plans will also include examining its utility for controlling unmanned systems in complex maritime environments.

[Sponsored by NRL]