

Using Continuous-Planning Techniques to Achieve Autonomy and Coordination Among Multiple Unmanned Aerial Vehicles

Gerald M. Powell, Ph.D.

U.S. Army Communications-Electronics Command (CECOM)
Software Engineering Center
Fort Monmouth, NJ 07703
732/532-3154
gerald.powell@mail1.monmouth.army.mil

Abstract

Existing Unmanned Aerial Vehicle (UAV) systems exhibit shortcomings in providing continuous, responsive, timely, and detailed information and targeting support to Army tactical commander's combat operations in an Army XXI battlespace. To synchronize Tactical UAV missions with supported operations, time is the critical element. Anything that can reduce TUAV planning time, while maintaining plan effectiveness, will expedite execution of a TUAV's mission. Autonomous flight with some ability to avoid and evade certain threats would increase survivability further. This paper begins by presenting some of the recent successes achieved by artificial intelligence (AI) planners and schedulers on complex real-world problems. It then attempts to show how NASA's demonstrated utility of a dynamic AI planning system prototype for conducting autonomous distributed planning and execution for a team of rovers engaged in missions to achieve science goals during planetary operations can be generalized and applied to a team of TUAVs. Last, it discusses some data collection opportunities that should appear due to the ability to place increasingly more processing and data storage capabilities onboard TUAVs and some of the key challenges to utilize those capabilities to produce more timely and immediately usable interpretations.

Introduction

In recent years we have become aware of numerous applications of artificial intelligence (AI) technology to real-world problems in the areas of planning and scheduling, real-time monitoring and control, and scientific data analysis. In the area of space applications, NASA has had multiple successes in applications that require planning and scheduling systems that need to represent and reason about complex activities, resources and interactions, [e.g., 1, 17]. For example, DATA-CHASER, which collected data in the far and extreme ultraviolet wavelengths, was on board the Space Shuttle Discovery on mission STS-85 in 1997 [2]. NASA reported this mission made use of automated planning and scheduling techniques to reduce mission commanding effort by 80 percent and increase science return (efficiency of instrument utilization) by 40 percent (compared to sequence generation done manually). The

planning and scheduling system managed the shuttle resources which successfully carried out the mission. NASA also reported the Modified Antarctic Mapping Mission (MAMM) used a synthetic aperture radar satellite to gather interferometry information covering the Antarctic continent from September to December 2000. NASA's ASPEN (Automated Scheduling and Planning ENvironment) [10, 4] was used to generate and verify the MAMM mission plan which resulted in a decrease from one year of planning effort for the first Antarctic Mapping Mission to approximately eight work weeks for MAMM. NASA reported the mission plan was carried out flawlessly onboard the satellite during the operation, and that the most difficult planning issue for MAMM was to guarantee all images were taken within the operational constraints and all data were down-linked successfully while adhering to the downlink constraints. Resource availability could change during the information collection cycle; this made rapid replanning critical in the event of such changes.

NASA also has demonstrated the utility of a dynamic AI planning system prototype for conducting autonomous distributed planning and execution for a team of rovers engaged in missions to achieve science goals during planetary operations. The system in this case is called CASPER (Continuous Activity Scheduling Planning Execution and Replanning) which is a soft real-time version of ASPEN.

There are striking similarities between the tasks required for data collection and analysis by distributed rovers to achieve science goals and the tasks required for data collection and analysis operations involving unmanned vehicles such as the Army's Tactical Unmanned Aerial Vehicle (TUAV) [16].

Given these similarities between tasks, and the successes by NASA such as those described above, the purpose of this paper is to try to show how the architectures, techniques and technologies developed and utilized by NASA may be applied to design a capability for cooperative, autonomous operation of TUAVs. Among many factors, the operational tempo and lethality characterizing the modern battlespace present significant challenges to having critical information provided to decision-makers and shooters in a timely manner. Also, if it desired to use TUAVs in larger numbers, or as members of a team of TUAVs, there will need to be capabilities to manage them effectively. The rest of this paper will make the case that onboard planning that is continuous, and distributed across vehicles, could play a significant role in meeting these challenges.

The system implementing these capabilities is CASPER. Since CASPER is the soft real-time version of ASPEN, it is necessary to first discuss ASPEN. The remainder of this paper begins by highlighting key capabilities of ASPEN as well as many of the techniques and technologies it implements that are relevant to this problem domain. Next, we discuss continuous planning in terms of what it is; how this type of planning supports the development of autonomous operations of TUAVs and coordination among them; and in particular attempt to show how the CASPER system, which implements continuous planning, could be utilized to provide the capabilities for

battlespace collection activities. This is followed by a discussion of a proposed system architecture for multiple TUAVs, the integration of planning and execution for multi-TUAV operations, and distributed planning. We conclude with a discussion of possible future directions for attempting to develop sophisticated team-based capabilities for autonomous TUAVs. In particular, we take the view that the distributed cooperative problem-solving paradigm is an approach worth considering.

ASPEN

NASA has reported that ASPEN, an object-oriented system, provides a reusable set of software components that implement the elements typically found in complex planning/scheduling systems [4]. These consist of:

- an expressive constraint modeling language to permit the user to create a natural definition of the application domain
- a system for managing constraints: representing and maintaining spacecraft operability and resource constraints, as well as requirements associated with activities
- a set of search strategies for generating and repairing plans to satisfy hard constraints
- a language for representing preferences in plans and optimizing such preferences
- a soft, real-time capability for replanning
- a temporal reasoning system for representing and maintaining temporal constraints
- a graphical user interface for visualizing plans/schedules (for use in mixed-initiative systems in which the process of problem solving is interactive)

ASPEN Modeling Language

Spacecraft models are represented in the ASPEN modeling language [15, 14]. The models get parsed into data structures that provide efficient reasoning capabilities for planning and scheduling. Spacecraft knowledge is represented using seven core model classes: activities, parameters, parameter dependencies, temporal constraints, reservations, resources and state variables. An activity, which is the primary construct in ASPEN, is an occurrence over a time interval that in some manner influences the spacecraft. An activity can represent anything such as a high-level goal or request, or a low-level event or command. These classes can be used to define spacecraft components, procedures, rules and constraints to provide a basis for manual or automatic generation of valid sequences of activities, also called plans or schedules. These classes, together, specify what the spacecraft can and cannot do during operations. All of these classes should be directly usable for developing TUAV models because spacecraft knowledge and TUAV knowledge have many similarities and the representation and reasoning requirements associated with data collection and analysis activities of orbiters, rovers and other NASA vehicles are congruent with those of TUAVs.

The ASPEN language is designed for use by domain experts who have no knowledge of automated planning technology. Knowledge of spacecraft operations is expressible in ways that are natural to operations personnel. Providing a means for subject matter

experts to directly express domain knowledge in the representation language helps alleviate the knowledge acquisition bottleneck and facilitates maintaining an accurate body of knowledge [15].

Iterative Repair Algorithm and Search

The main algorithm for automated planning and scheduling in ASPEN is based on an approach called iterative repair [18]. ASPEN also has the flexibility to support the other major class of AI scheduling algorithm which is called constructive [9]. During iterative repair, the conflicts in the schedule are detected and addressed one at a time until none exist, or a user-specified time limit has been reached. A conflict is a violation of a reservation, parameter dependency or temporal constraint. Conflicts are repaired by using several predefined methods such as moving an activity, deleting an activity, changing a parameter value, etc. [10, 12]. Each conflict gives information about the particular objects involved and how to repair the conflict. The search space for plans and schedules in ASPEN consists of all possible repair methods applied to all possible conflicts in all possible orders. The iterative repair algorithm searches this space of possible schedules by making decisions at certain choice points such as selecting a conflict or selecting a repair method, and revising the schedule in accordance with these decisions. ASPEN uses search heuristics to guide the search. ASPEN currently has a number of domain-independent heuristics that can be used to repair conflicts. For example, there is one for sorting conflicts by their type; another for selecting the repair method for a given conflict; and a third for selecting start-time intervals for activities that are moved or created. The iterative repair algorithm and search heuristics should be directly usable for UAVs.

Continuous Planning

A major success reported by NASA, the Mars Pathfinder mission demonstrated the ability to send rovers to other planets. Missions are underway, or being planned, to send additional robotic vehicles to Mars (e.g., Mars Odyssey) as well as to outer planets and to collect pieces of the sun [11]. To increase science return, future missions will need larger groups of rovers to collect the desired data. These rovers will need to operate in a coordinated manner where each one achieves a subset of the overall mission goals and shares information it acquires. Moreover, it will be beneficial to have highly autonomous rovers needing little communication with scientists and engineers to carry out the rovers' tasks. A rover with autonomy will be more capable of making decisions regarding how to best accomplish science goals as well as being able to react to its environment and deal with unforeseen events while attaining these goals.

An autonomous rover (or team of rovers) must be able to respond in a timely manner to a dynamic and unpredictable environment. Plans used by rovers often need to be modified in the case of fortuitous events such as science observations completing early and in the case of setbacks such as traverses requiring more time than expected or device failures [6].

This type of situation where a plan must be continually updated in light of a changing operating context is sometimes called continuous planning. In this mode of operation, a planner would be continuously updating the plan (e.g., every few seconds) based on sensor and other feedback, and then revising the existing plan to accommodate any new information.

There are probably situations where unmanned vehicles in general, not just TUAVs, would be better able to perform their missions if they could operate in an autonomous mode in addition to being controllable by humans via remotely located control stations. An ability to perform continuous planning is one element supporting autonomous operation. A planner with this type of planning capability can:

- Be more responsive to unexpected changes in the environment. These changes could be related to the status of activities being carried out, as well as updates to state (e.g., illumination in the collection area) or resource values (e.g., fuel level).
- Reduce reliance on predictive models of the environment because it will be updating its plans continuously. Errors in models or uncertainties in the environment can be dealt with without causing plans to fail and without explicitly specifying all contingencies in the planning model.
- Have TUAV fault-protection and execution layers address controlling the AV over a shorter time horizon because the planner will replan in a much shorter time span. (Note: For readability, TUAV and AV are used synonymously in this paper.)

In the traditional mode of automated planning, planning is considered a batch process. The system operates on a relatively long-term planning horizon, and the plan is completely generated prior to the start of execution. In the case of TUAVs, no automated planning capability exists. Rather, military intelligence personnel identify collection goals and develop the maneuver and collection plans for each vehicle; an executable version of the plan is pre-programmed and loaded onto the vehicle prior to launch. Revisions to collection goals (e.g., specific areas of interest) are identified by military personnel; as necessary, they then revise the plan, or develop a new plan, for achieving them.

Changes to plans (new locations, flight routes, etc.) are uplinked to the AV via data link with the ground control station. The ground station monitors AV location and flight instrumentation data and receives collection data as it is collected. If the AV enters a state where it cannot proceed with the mission, it must execute an appropriate pre-programmed contingency response such as flying to a specific location. There is no onboard capability to autonomously analyze the situation and replan in an attempt to execute a response that will allow it to continue with its mission. In addition, if an unpredictable fortuitous event occurs, the plan will not necessarily be modified in a timely manner to take advantage of the situation. Data representing the fortuitous event would have to be downlinked, interpreted by military personnel, a revision to the plan would need to be developed and then uplinked to the TUAV.

This paper presents a continuous planning approach to TUAV operations. This type of planning is intended to achieve a higher degree of responsiveness in situations where replanning is necessary or desired.

This paper will attempt to show the applicability of the CASPER [3] planning system to control a collection of distributed TUAVs for battlespace collection operations. Knowledge relevant to rovers and their operations would need to be replaced by knowledge relevant to TUAVs and their operations. That should be achievable within the customizable framework provided by CASPER. Based on a set of collection goals as input and each TUAV's initial conditions, this planning system could generate a sequence of activities satisfying the goals while obeying each of the TUAV's resource constraints and operations rules. Plans can be generated using an "iterative repair" algorithm which classifies conflicts and resolves them one at a time by performing one or more modifications to the plan. After a valid command sequence is generated, commands would be submitted to the vehicles low-level control software for execution. Execution updates are provided by this control software to a monitoring element within the planning system. As information arrives with respect to command status and actual resource usage, the planner can update projections for a future plan. These updates may cause new conflicts and/or opportunities requiring the planner to replan in order to accommodate the unexpected events. Planning activities would be distributed between the individual TUAVs; each TUAV would be responsible for planning its own activities. One possible architecture for organizing the TUAVs is to have one of the AVs serve as a central planner and be responsible for receiving new goals from the ground station and allocating them appropriately to individual TUAVs on the team.

NASA reported the CASPER planning system has been integrated with other software components to form a multi-rover execution architecture [7,8]. The components include a machine learning science analysis tool which analyzes planetary data and generates a set of goals for new science observations, a simulation environment that models multiple-rover science operations in a Mars-like terrain, a real-time multi-rover hardware and kinematics simulator, and control software from the NASA JPL Rocky 7 rover. An attempt is made in this paper to show how this multi-rover execution architecture could be generalized for multi-TUAV operations.

The remainder of this paper begins by characterizing the multi-TUAV application domain and describes activities characterizing collection behavior among the AVs and the ground station. Next, we describe a proposed multi-TUAV execution architecture which controls and coordinates operations for a team of TUAVs. Then we focus on the planning elements of this architecture including a presentation of a candidate approach to distributed planning, the generalizability of the CASPER continuous planning system, and a possible approach to plan optimization for this domain. The final section discusses possible future directions in which this work could be taken.

Cooperating TUAVs for Battlespace Collection

Using a team of TUAVs for collection has important advantages. Multiple TUAVs working cooperatively on a collection goal could focus on the same target, or area of interest (e.g., a segment of road), for collection from different perspectives simultaneously thereby increasing the overall collection rate. They could also employ different collection disciplines (e.g., imagery intelligence (IMINT) and signals intelligence (SIGINT)) to gather data of different types. The TUAV team should be designed to behave in a coordinated manner decomposing and allocating goals appropriately among the team and sharing acquired information. These approaches have the potential of increasing collection accuracy, and speed of achieving collection goals.

Coordinating multiple distributed agents raises issues pertaining to communication, control and the allocation of capabilities to place onboard each individual agent. In the present paper, a configuration of three TUAVs is employed in which each one has a planning and intelligence analysis capability onboard. Each TUAV can plan for the goals assigned to it, collect data against those goals, and perform data analysis onboard which will be used to develop future goals. Moreover, each AV can monitor the execution of its own plans and carry out replanning as required. Central planner and data analysis modules are assumed to be located on one of the AVs; this is used to coordinate goals and overall intelligence analysis.

This system could be evaluated by testing its performance in terms of time required for, and accuracy of, detecting, locating, tracking and identifying targets. Collection goals could consist of, for example, requests to collect data at certain locations and by different collection disciplines. Goals could be prioritized on each TUAV thereby focusing the planner on the highest priority goals when conditions require it (such as low fuel level).

Data analysis goals are allocated among the three TUAVs. It is assumed initially each AV has only one type of collection discipline onboard (e.g., IMINT or SIGINT) although the model could be expanded such that each AV has the same set of two or more collection disciplines. The Future Directions section of this paper will provide further discussion of distributed processing/problem-solving. It is also assumed that collected data is immediately transmitted to the central planner where it is stored in memory until it can be interpreted. Downlinking of interpreted and/or uninterpreted data can occur whenever an AV is in communication contact with the ground station. An ability to decide when data can be downlinked to ground versus when it needs to be retained in the air for processing would exist on each AV. Downloading data would free up memory.

System Architecture for Multiple TUAVs

The model of distributed planning by TUAVs described here is based on part of a multi-rover execution architecture used to coordinate multi-rover behavior and to

provide autonomous rover operations as reported by NASA [8]. This architecture utilizes a framework for autonomously generating and achieving planetary science goals [7]. The present paper describes how this architectural framework could be used in achieving intelligence collection goals. The overall execution architecture provided by NASA [8] is generalized and applied here as shown in Figure 1.

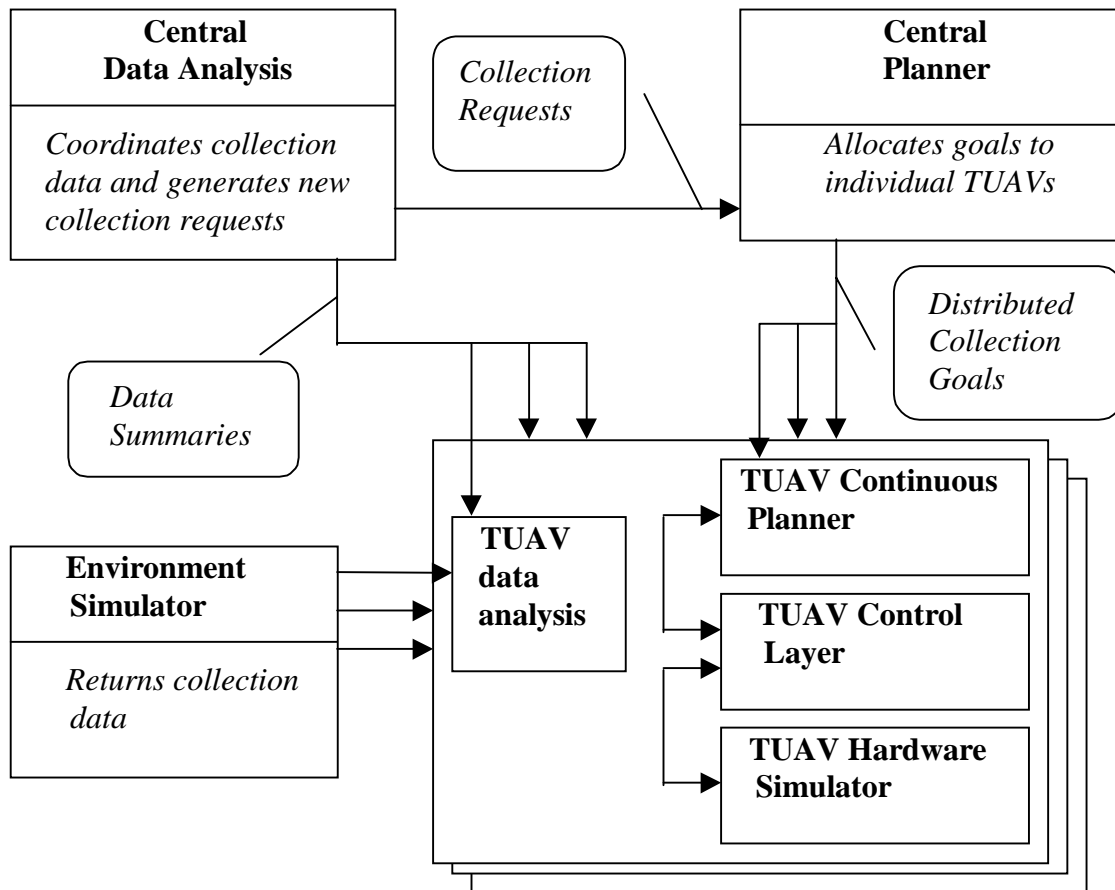


Figure 1. Multi-TUAV Execution Architecture

The following major components comprise the system:

Planning: A dynamic, distributed planning system generates TUAV-operation plans to accomplish collection goals submitted to it. Planning is distributed between a central planner, which divides and allocates collection goals between AVs, and a distributed set of planners that plan the operations on an individual AV. Each

individual AV has the capabilities to carry out execution monitoring and replanning whereby plans get updated as required in response to unanticipated events.

Data Analysis: There are many ways, in principle, to design a team of entities for collaborative data analysis. The space of solutions will be constrained by such factors as the amount of processing power and data storage available as well as the degree to which communication between entities needs to be minimized. One proposal for a design would give each entity in the team a capability to develop an interpretation of the situation data strictly from its local perspective based on the information gathering and analysis capability it has available. For example, one member of the team may be able to obtain location information on a given observable based on one of its characteristics. Another team member could also provide location information, from its local perspective, based on other characteristics of the observable; in this design, neither entity would share data or models with the other, but would pass this information to a central data analyzer for integration. Clearly, many alternative designs can be considered. What specific abilities to give each entity and how the team is set up to work collaboratively on an overall team analysis goal would need to be determined.

TUAV Control Software: Each AV has control software that handles execution of low-level AV commands in the areas of navigation and instrument manipulation. This software carries out low-level monitoring and control of an AV's sub-systems.

TUAV Hardware Simulator (Multi-TUAV Real-time Simulator): This is a multi-TUAV simulation environment simulating the 3-D nature of the battlespace and AV physical operations within that environment. It models AV kinematics and generates sensor feedback which is passed back to the continuous planner of each AV.

Environment Simulator (Collection Simulator): This is a multiple TUAV simulator modeling different collection environments and AV collection activities within them. The simulator manages collection data for each environment, tracks AV operations within that environment, and reflects readings by AV collection instruments.

The overall system operates in a closed-loop manner where the data analysis system can be thought of as an intelligence analyst guiding the collection process. Data from different intelligence disciplines are input to the onboard data analysis algorithms which broadcast their interpretations (situation models) to the central analysis module. This module constructs a global model of the situation and develops a new set of collection goals that are intended to increase the accuracy of the global model. These goals are submitted to a central planner which allocates them appropriately to individual AVs in a manner that provides a basis for achieving each of them. Each AV then generates a set of actions for that AV which will accomplish as many of the assigned goals as possible. Execution of these action sequences is carried out by the AV low-level control software and a multi-TUAV simulation environment that pass action and state updates to each onboard planner. The continuous planner on each

AV can carry out replanning when unexpected events or failures occur. Actions are also executed within the environment simulator and collected data are relayed back to the AV data analysis module. This cycle continues, within operational constraints (e.g., fuel levels and time), until all collection goals are achieved.

Integrating Planning and Execution for Multi-TUAV Operations

To generate individual rover plans for a team of rovers, NASA developed a distributed planning environment using the CASPER continuous planning system [3]. CASPER is an extended version of the ASPEN system which was constructed to address dynamic planning and scheduling applications. NASA reports CASPER provides a generic planning/scheduling application framework that can be tailored to specific domains [6].

CASPER uses techniques from planning and scheduling to automatically produce the required rover-activity sequence to achieve the goals input to it. As described earlier in this paper, an iterative repair algorithm generating this sequence classifies conflicts and addresses them one at a time. Conflicts arise when a plan constraint is violated. The constraint may be temporal in nature, or involve a state, resource or activity parameter. To resolve conflicts, one or more schedule modifications are performed such as moving, adding or deleting activities.

This iterative repair algorithm could be applied to generate and repair plans on individual TUAVs. An AV outside its required location for a scheduled collection activity represents one type of conflict. Adding a movement command to send the AV to the proper location would resolve this type of conflict. An example of another type of conflict is having too many AVs communicating with the central planner AV at one time. The iterative repair algorithm executes until no conflicts remain in the schedule, or a user-defined timeout is reached.

Distributed Planning

NASA developed a distributed planning environment to support missions involving multiple rovers. This paper describes how this environment can be generalized and applied to collection missions involving multiple TUAVs. The environment assumes each AV has a planner onboard. This permits each AV to plan for itself and/or for other AVs. The nature of intelligence collection presents situations where distributing planning and dynamic replanning capabilities across each AV would be beneficial. For example, a goal of remaining undetected argues for minimizing communications between AVs, and between AVs and the ground station. In addition, collection goals to achieve synchrony with overall maneuver plans might be accomplished more rapidly when collection agents have autonomy to generate their own action sequences and to modify them in real-time independent of a central planner or ground control station. This approach to distributed planning would involve using a CASPER continuous planner on each AV, in addition to a central planner (batch planner) onboard one of the AVs. The central planner produces an abstract plan for all AVs. Each AV then elaborates this plan into a detailed, executable plan for its own

activities. The central planner allocates a global set of goals among the AVs; this configuration has been called distributed planning with central goal allocation [5]. As an example, a goal may require collection on a particular area of interest without specifying which AV collects on it. The central planner may assign this goal to the AV closest to the area of interest in order to minimize traversals, or it may make an assignment based on the type of intelligence most likely to yield the most useful information. This design is one of numerous approaches to distributed planning that can be considered [13].

Continuous Planning for Each TUAV

The continuous planning approach would yield a high degree of responsiveness by the planner onboard each AV. When planning is considered a batch process, the planner is given a goal (or set of goals) and an initial state. Its job is to find a sequence of operators (e.g., actions) that can achieve the goal state; it is assumed execution will not begin until planning (out to the planning horizon) is completed. In continuous planning, each TUAV would have a current goal set, a current state, a current plan, and state projections into the future for the current plan. At any moment an incremental update to the goals or current state may update the current plan. The update could represent an unexpected event or simply a progression of time. Each planner has responsibility for maintaining a plan consistent with the most up-to-date information. The current plan represents the planner's estimation regarding what it anticipates happening in the world if things unfold as expected. Since things seldom go exactly as anticipated, the planner is poised to continually modify the plan. The iterative repair techniques discussed above enable incremental modifications to the goals, initial state or plan and then iteratively resolve conflicts that may occur.

Figure 2 specifies the CASPER continuous planning algorithm. With this approach, the state of an AV would be modeled by a set of plan timelines; these would depict the current and expected state of an AV over time. During each iteration of this loop, the actual state of the AV would drift from the state expected by the timelines; these differences indicate changes in the world. When the AV control software and sensors pass back information representing updates, the system would update the timelines with actual state and resource values, as well as starting and completion times associated with activities. Each update, when synchronized with the current plan, may cause conflicts (Step 3). Recall a conflict occurs when an action in the plan is not appropriate because its required state and/or resource values violate the plan's constraints.

If a conflict occurs, the system records it and makes modifications to the plan to bring the plan back into sync with the current state and future-plan projections. Because the plan is updated with very short intervals (e.g., a few seconds), the plan rarely has a chance to become significantly out of sync. Consequently, the high level actions of the system are more responsive to the actual state of the vehicle.

Initialize P to the null plan

Initialize G to the null set
Initialize S to the current state

Given a current plan P and a current goal set G:

1. Update g to reflect new goals or goals that are no longer needed
2. Update S to the revised current state
3. Compute conflicts on (P,G,S)
4. Apply conflict resolution planning methods to P (within resource bounds)
5. Release relevant near-term activities in P to RTS for execution
6. Goto 1

Figure 2. CASPER Continuous Planning Algorithm

Plan Optimization

Data collection at the tactical level of war typically will be severely time-constrained in order to support the operational tempo of maneuver desired by the commander or required by the battlespace situation (e.g., when in a search and rescue mode, collection of key information characterizing different areas of the battlespace in order to plan the “best” flight route). Under such severe time constraints it is advantageous to generate plans and schedules that are optimal (i.e., efficient) for collecting in the areas that need to be traversed and searched. For tasks that can be cast appropriately as traversal problems, NASA has experienced reduced overall traversal distance and expected execution time by implementing heuristics based on techniques for addressing the multiple traveling salesmen problem (MTSP); an extension of the TSP. For MTSP, at least one member of a sales team has to visit every city such that total traveling time is minimized. As with the multiple rover situation, both the central and individual AV planners would use the MTSP heuristics. They would be used to help identify which AV should be given a particular collection activity and to select a temporal placement for the activity.

Future Directions

The planning, observing, and processing performed by information gatherers (rovers; commercial air traffic controllers; fighters of forest fires; etc.) could be considered a case of real-time distributed situation assessment. Often, the operational characteristics of environments where these entities perform these activities require sensor data (spectrometry of soil samples; images reflecting boundaries of lakes and rivers; etc.) and interpretations to be communicated or developed in real-time for decision-makers and other entities. The task is inherently spatially distributed. If enough processing power and data storage can be placed onboard information gatherers, there will be an opportunity to utilize teams of such entities carrying multiple types of gathering capabilities to carry out assessments yielding descriptions of situations at a higher level.

The team-based approach to information gathering and processing can be viewed as a network of distributed problem-solving nodes. However, unlike traditional distributed processing where tasks can be decomposed such that a node seldom needs the assistance of another, in an environment where communication may be highly limited (deep space planetary operations; etc.) there is an increased requirement for more sophistication (knowledge) and autonomy at each node. The requirement for limited communication in a distributed network has led to the construction of problem-solving architectures that can function with possibly incomplete and inconsistent data and control information, but where the sub-problems that a given node works on are not necessarily independently solvable. Consequently, nodes will produce tentative partial answers based on local information and then pass these results to other nodes. The constraints existing between the sub-problems of different nodes are utilized to resolve the local uncertainties and global inconsistencies that arise from inaccurate, incomplete and out-of-date local information.

Key issues in this paradigm include how to organize local and team-wide problem solving so that the entities can cooperate to generate assessments of sufficient quality under fixed deadlines, while using limited communication bandwidth. They will need to be robust enough that their performance degrades gracefully since their ability to gather, process and communicate may deteriorate and ultimately fail over time.

References

1. Chien, S., DeCoste, D., Doyle, R. and Stolorz, P. Making an Impact: Artificial Intelligence at the Jet Propulsion Laboratory, *AI Magazine*, 18 (1): 103-122, 1997.
2. Chien, S., Rabideau, G., Willis, J. and Mann, T. Automating Planning and Scheduling of Shuttle Payload Operations. *Artificial Intelligence*, 114 (1): 239-255, 1999.
3. Chien, S., Knight, R., Stechert, A., Sherwood, R., and Rabideau, G. Integrated Planning and Execution for Autonomous Spacecraft, *Proceedings of the IEEE Aerospace Conference (IAC)*, Aspen, CO, March 1999
4. Chien, S., Rabideau, G., Knight, R., Sherwood, R., Engelhardt, B., Mutz, D., Estlin, T., Smith, B., Fisher, F., Barrett, T., Stebbins, G., and Tran, D. ASPEN Automating Space Mission Operations Using Automated Planning and Scheduling. Paper presented at *SpaceOps 2000*, 19-23, June, Toulouse, France, 2000.
5. Chien, S. Barrett, A., Estlin, T., and Rabideau, G. Three Coordinated Planning Methods for Cooperating Rovers. *Proceedings of the World Automation Congress*, Wailea, HI 2000.
6. Estlin, T., Rabideau, G., Mutz, D. and Chien, S. Using Continuous Planning Techniques to Coordinate Multiple Rovers. *IJCAI99 Workshop on Scheduling and*

Planning meet Real-time Monitoring in a Dynamic and Uncertain World, Stockholm, Sweden, August 1999.

7. Estlin, T., T. Mann, A. Gray, G. Rabideau, R. Castano, S. Chien and E. Mjolsness, An Integrated System for Multi-Rover Scientific Exploration. *Sixteenth National Conference of Artificial Intelligence (AAAI-99)*, Orlando, FL, July 1999a

8. Estlin, T., Yen, J. et al. An Integrated Architecture for Cooperating Rovers. In *Proceedings of the 1999 International Symposium on Artificial Intelligence, Robotics and Automation for Space*, Noordwijk, The Netherlands, June 1999b.

9. Fox, M. and Zweben, M. Eds., *Intelligent Scheduling*. Morgan-Kaufman Pub., 1994.

10. Fukunaga, A., Rabideau, G., Chien, S., and Yan, D. Towards an Application Framework for Automated Planning and Scheduling. In *Proceedings of the 1997 International Symposium on Artificial Intelligence, Robotics and Automation for Space*, Tokyo, Japan, July 1997.

11. JPL, 2001 <http://www.jpl.nasa.gov/missions>

12. Rabideau, G., Knight, R., Chien, S., Fukunaga, A., and Govindjee, A. Iterative Repair Planning for Spacecraft Operations Using the ASPEN System. *International Symposium on Artificial Intelligence Robotics and Automation in Space (ISAIRAS)*, Noordwijk, The Netherlands, June 1999

13. Rabideau, G., T. Estlin, S. Chien, A. Barrett, A Comparison of Coordinated Planning Methods for Cooperating Rovers. *AIAA 1999 Space Technology Conference*, Albuquerque, NM, September 1999

14. Sherwood, R., A. Govindjee, D. Yan, G. Rabideau, S. Chien, A. Fukunaga, "Using ASPEN to Automate EO-1 Activity Planning", *Proceedings of the 1998 IEEE Aerospace Conference*, Aspen, CO, March 1998.

15. Smith, B., Sherwood, R., Govindjee, a., Yan, D., Rabideau, g., Chien, S., Fukunaga, A. Representing Spacecraft Mission Planning Knowledge in ASPEN. *Artificial Intelligence Planning Systems Workshop on Knowledge Acquisition*, Pittsburgh, PA, 1998.

16. TUAV CONOPs, U.S. Army, TRADOC System Manager UAV, January 15, 2001

17. Wilkens, D. E. and desJardins, M. A Call for Knowledge-Based Planning. *AI Magazine*, 22 (1): 99-115, 2001.

18. Zweben, M., Daun, B., Davis, E., and Deale M. Scheduling and Rescheduling with Iterative Repair. In Monte Zweben and Mark Fox (Eds.), *Intelligent Scheduling*, pp. 241-256, Morgan Kaufman, San Francisco, CA 1994.