

Autonomous Planning and Execution for a Future Titan Aerobot*

Daniel Gaines, Tara Estlin, Steve Schaffer, Caroline Chouinard and Alberto Elfes

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109

{*firstname.lastname*}@jpl.nasa.gov

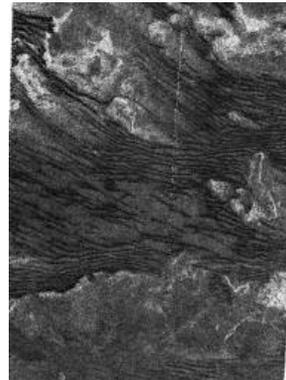
Abstract

We are developing onboard planning and execution technologies to provide robust and opportunistic mission operations for a future Titan aerobot. Aerobots have the potential for collecting a vast amount of high priority science data. However, to be effective, an aerobot must address several challenges including communication constraints, extended periods without contact with Earth, uncertain and changing environmental conditions, maneuverability constraints and potentially short-lived science opportunities. We are developing the AerOASIS system to develop and test technology to support autonomous science operations for a future Titan Aerobot. The planning and execution component of AerOASIS is able to generate mission operations plans that achieve science and engineering objectives while respecting mission and resource constraints as well as adapting the plan to respond to new science opportunities. Our technology leverages prior work on the OASIS system for autonomous rover exploration. In this paper we describe how the OASIS planning component was adapted to address the unique challenges of a Titan Aerobot and we describe a field demonstration of the system with the JPL prototype aerobot.

Introduction

NASA's 2008 Solar System Exploration Roadmap (NASA 2008) highlights the importance of aerial probes as a strategic new technology for Solar System exploration, and outlines missions to Venus and Titan that would use airborne vehicles, such as balloons and airships (blimps). For example, recent data and imagery from Cassini and Huygens show that Titan is a fascinating planetary body with a variety of surface features (See Figure 1). Pictures from the Huygens probe dramatically illustrate the utility of low altitude (< 10 km) aerial imagery at Titan, showing river channels and other striking terrain features not visible from orbit. Airborne vehicles provide a promising means of exploring these areas as they offer traversal capabilities, geographical coverage, and speeds that are orders of magnitude greater than rovers, leading to an enormous science data collection potential.

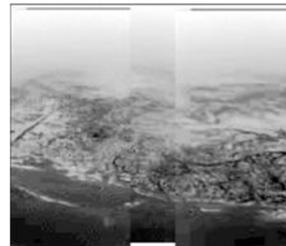
* (c) 2009 California Institute of Technology. Government sponsorship acknowledged.



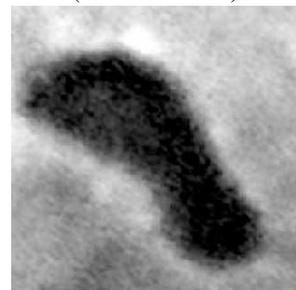
Dune Field
(Huygens)



Possible Cryovolcano
(Cassini VIMS)



Possible Shoreline
(Cassini radar mapper)



Liquid Lake
(Cassini VIMS)

Figure 1: Examples of geographic features of interest on Titan.

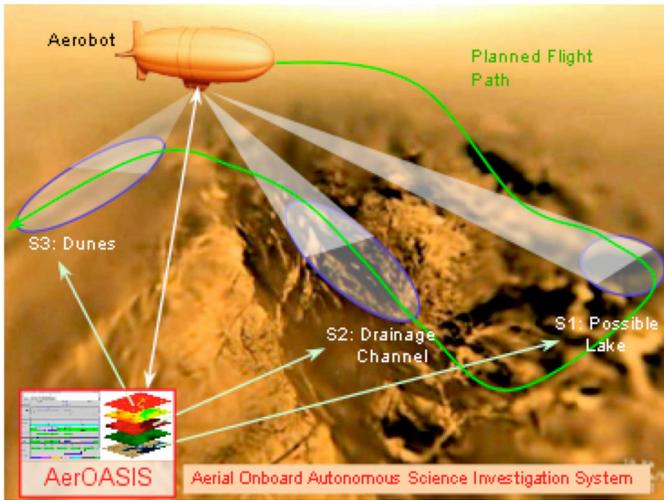


Figure 2: An aerobot science operations plan showing goals that made it into the plan and goals in reserve for potential replanning.

However, bandwidth constraints, communications latencies and blackouts, and flight maneuvering limitations require these vehicles to have onboard autonomous science capabilities. This will allow science data acquisition to be planned and executed in real time, and will consequently maximize the mission science return. Specifically, autonomous science technology would provide the ability to 1) prioritize data for downlink such that most important data is downlinked early, 2) summarize and compress data, for example, by compiling and downlinking statistics on terrain features observed rather than complete image sets, and 3) detect and respond to science opportunities onboard the vehicle before they are passed over.

We are developing the Aerial Onboard Autonomous Science Investigation System (AerOASIS) to provide onboard science capabilities for aerial probes. Figure 2 illustrates the concept. The illustration shows a Titan aerobot (robotic blimp) surveying the Huygens descent region at 6 km altitude. The AerOASIS software will analyze the data obtained from onboard multi-spectral imagers and other sensors, and uses scientist-defined signatures to prioritize data, search for high-value science targets and plan science activities. In this example, AerOASIS has selected three science sites: a possible methane lake, a drainage system, and some distant dunes. The aerobot plans its flight trajectory taking into account local wind conditions, the topography, and its own flight maneuverability and power limitations, and performs close-up surveys of the three selected sites.

In this paper we will focus on the autonomous planning and scheduling component for AerOASIS. The planning component accepts mission goals and constraints from the ground and develops an operations plan that attempts to maximize science while respecting resource and mission constraints. It also supports opportunistic science for targets identified onboard by attempting to alter the plan to collect additional data within available time and resource limita-

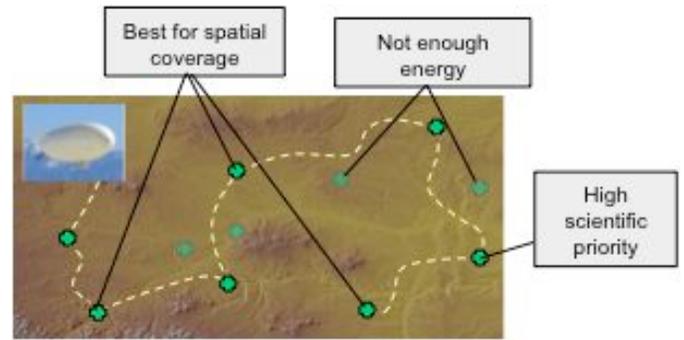


Figure 3: An aerobot science operations plan showing goals that made it into the plan and goals in reserve for potential replanning.

tions.

One of the significant challenges faced by an aerial vehicle is handling the large degree of uncertainty in the environment. For example, the system must respond to unexpected wind conditions as well as take advantage of new science opportunities. Using the CASPER continuous planning and executing system, we have developed autonomous technology that enables an aerobot to generate mission operation plans and adjust the plans to accommodate unexpected events during execution (Figure 3). This enables the system to appropriately respond to unexpected events and to take advantage of new science opportunities.

This technology is based on prior work we have performed for rover operations (Castano *et al.* 2007; Gaines, Estlin, & Chouinard 2008). However, some significant developments were required to address the unique challenges of an aerobot platform. This includes the fact that the aerobot platform is never stationary and is traveling at a significantly greater rate than a rover. In the paper we will describe these challenges and discuss how we addressed them. We will also describe a field demonstration of the planning system integrated with the JPL Aerobot.

AerOASIS: Aerial Onboard Autonomous Science Investigation System

The Aerial Onboard Autonomous Science Investigation System (AerOASIS) (Figure 4) allows an airborne planetary exploration vehicle to:

1. summarize and prioritize the most scientifically relevant data from the various incoming sensor streams for relay to Earth;
2. identify and select high-value science sites for additional investigation by the aerial vehicle (through low-altitude, high resolution surveys; in-situ probe deployment; and/or surface sample acquisition); and
3. dynamically plan, schedule and monitor the various science activities being performed by the aerial vehicle, even during extended communications blackout periods with Earth.

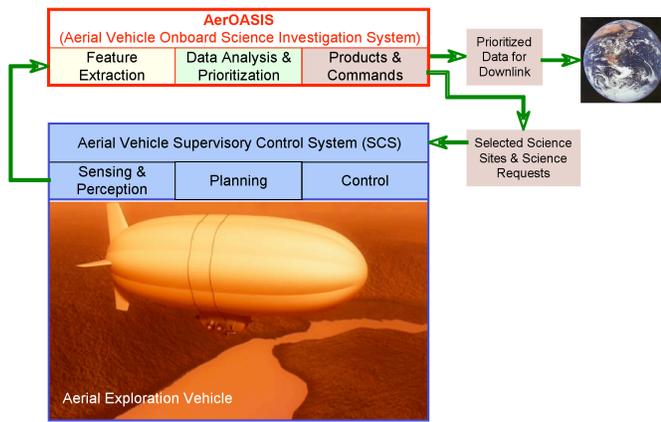


Figure 4: The AerOASIS Aerial Vehicle Autonomous Science System.

The AerOASIS system is composed of three main subsystems:

1. Feature Extraction, which processes sensor imagery and other types of data (such as atmospheric pressure, temperature, wind speeds, etc.) and performs data segmentation and feature extraction.
2. Data Analysis and Prioritization, which matches the extracted feature vectors against scientist-defined signatures. The results are used to 1) detect novelty, i.e. statistically significant new types of information or data correlations; 2) perform science data prioritization and summarization for downlink to Earth; 3) identify and select high-value science sites for in-situ studies to be conducted by the aerial vehicle.
3. Planning and scheduling, which generates operations plans to achieve observation requests submitted from Earth and from onboard data analysis. These science requests can include low-altitude high-resolution surveys, in-situ sonde deployment, and/or surface sample acquisition for onboard analysis.

The AerOASIS system receives images and other sensor data streams from the aerial vehicle, and performs data segmentation and feature extraction. The resulting feature vectors are processed by the data analysis and prioritization subsystem, which performs the identification and selection of high-value science sites, as well as the prioritization and summarization of science data for downlink to Earth. The selected high-priority science sites, as well as additional science requests, are handed as observation requests to the planning and scheduling component which attempts to accomplish these new requests. The planner interfaces with the Aerial Vehicle Supervisory Control System (SCS) to execute the plan and to receive updates on the current state of the vehicle and the world.

The Aerial Vehicle Supervisory Control System (SCS) supervises all sensing/perception, planning, flight navigation and control activities of the aerial exploration vehicle and its deployable sensors and probes (Elfes *et al.* 2004;

2006). SCS provides the lower levels of the onboard autonomy architecture, including sensor and actuator control, vehicle state estimation, power management, the Flight Control System (FCS), and the underlying flight mode controller, as well as intermediate levels of autonomy such as navigation and flight planning, image-based motion estimation (IBME) for GPS-denied vehicle localization, image mosaicking and geographical mapping, and 3D terrain structure estimation.

In the remainder of this paper, we discuss the planning and scheduling component of AerOASIS in more detail.

Planning and Execution in the AerOASIS System

Our objective is to enable onboard planning software to generate correct and high quality operations plans to achieve mission objectives issued from ground operations as well as respond to science opportunities detected onboard the vehicle. In particular, the system considers prioritized observation requests. This will enable the ground team to uplink a larger set of observations and let the aerobot dynamically select among them based on the scientific and engineering merit of the resulting plan and the aerobot's assessment of available resources. During execution, the aerobot will modify the plan based on the current estimate of its resources.

Our approach is implemented within the CASPER system (Estlin *et al.* 2002; Chien *et al.* 2000). CASPER employs a continuous planning technique where the planner continually evaluates the current plan and modifies it when necessary based on new state and resource information. Rather than consider planning a batch process, where planning is performed once for a certain time period and set of goals, the planner has a current goal set, a current aerobot state, and state projections into the future for that plan. At any time an incremental update to the goals or current state may update the current plan. This update may be an unexpected event (such as a new science target) or a current reading for a particular resource level (such as battery charge). The planner is then responsible for maintaining a plan consistent with the most current information.

A plan consists of a set of grounded (i.e., time-tagged) activities that represent different aerobot actions and behaviors. Aerobot state in CASPER is modeled by a set of plan timelines, which contain information on states, such as aerobot position, and resources, such as energy. Timelines are calculated by reasoning about activity effects and represent the past, current and expected state of the aerobot over time. As time progresses, the actual state of the aerobot drifts from the state expected by the timelines, reflecting changes in the world. If an update results in a problem, such as an activity consuming more memory than expected and thereby over-subscribing RAM, CASPER re-plans, using iterative repair (Zweben *et al.* 1994), to address conflicts.

CASPER includes an optimization framework for reasoning about soft constraints such as reducing the distance traversed by the aerobot and increasing the value of science data collected. User-defined preferences are used to compute plan quality based on how well the plan satisfies these

constraints. Optimization proceeds similar to iterative repair. For each preference, an optimization heuristic generates modifications that could potentially improve the plan score.

Figure 5 provides a high level description of the control algorithm used for the aerobot application of CASPER. The algorithm takes as input a set of goals with associated science priorities and a set of time and resource constraints. CASPER’s optimization framework supports a wide-range of user-defined preferences. The main loop of the algorithm interleaves iterative repair and iterative optimization to search for a conflict-free plan of high quality. The loop begins by processing any updates on state and resource timelines or on activity status. It then enters a loop in which it attempts to improve the plan by repairing conflicts or performing optimization steps.

```

Input
A set of science observations (oversubscribed)
Time, resource constraints & Preferences

Repeat:
  Process updates from Executive
  Optimize for n iterations
    If no conflicts
      Select a preference to improve
      Select and apply an improvement method
    Else
      Select conflict to work on
      Select and apply a repair strategy
  Compute plan score
  If current plan is best seen so far, save it
  Reload plan with highest score
  Commit and/or Rescind activities
  If idle, attempt to move up future activities
  
```

Figure 5: CASPER control algorithm for aerobot domain.

If there are no conflicts, CASPER attempts to improve the plan by satisfying an observation. If there are conflicts, it will perform an iteration of repair, selecting one of the available repair methods (e.g. move an activity, add an activity, ...). If deletion of an observation is selected, it will select the observation providing the smallest spatial coverage contribution.

Note that satisfying an observation will likely introduce conflicts as this is where CASPER will evaluate the resource and temporal requirements of an observation. CASPER will use subsequent iterations to try to resolve these conflicts. For example, if the aerobot is not currently at the appropriate location to take an observation, CASPER will identify a state conflict which it will attempt to resolve. One option for fixing this conflict is to add an activity that could move the aerobot from one location to another, i.e. a traverse activity. This is also where CASPER selects an ordering of observations in an attempt to minimize traverse distance. We use a simple traveling salesman heuristic to pick start times for activities to reduce traverse distance.

Figure 6 illustrates the “lifetime” of observations in the system. New observations are placed in a requested bin. When an observation is selected to be satisfied, it moves from requested to pending in which it awaits execution. In the meantime, it may be deleted to resolve conflicts in the plan, in which case it moves back to requested. As it nears time for a pending observation to be executed, it is committed and sent to an executive process for execution. If a problem occurs in the plan before the actual execution time of the activity, the planner has the ability to request a rescind of the observation from the executive. If the executive is able to honor the rescind request, it is as if the observation had been deleted from the plan and it returns to the requested bin.

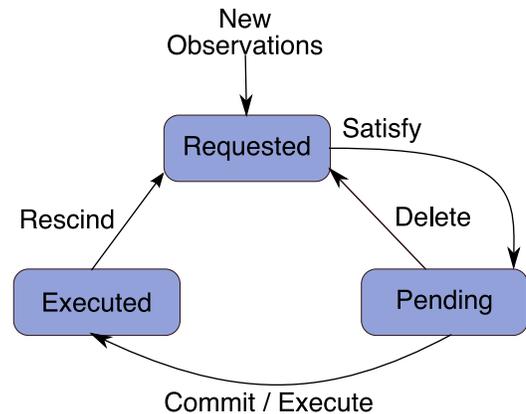


Figure 6: The lifetime of an observation.

A Comparison with Autonomous Planning for Rovers

The planning component of AerOASIS was developed from prior work of the authors on autonomous science operations for rovers (Castano *et al.* 2007; Gaines, Estlin, & Chouinard 2008). However, due to the unique challenges of an aerobot platform, there were some key differences.

The primary challenge with the aerobot deployment is that, unlike a rover, an aerobot is never stationary. It is always moving and subject to wind conditions. In the rover model, if the planner is not commanding rover motion, we can be reasonably assured that the rover is not moving. This gives the system plenty of time to schedule and execute imaging and other science observations. In contrast, the aerobot is always on the move, and depending on the situation, can be moving quite fast. Thus, the planning system had to be tailored to responding quickly changing conditions. Because target locations for imaging can be reached and then passed by so quickly, the planning system for the aerobot delegates the timing of the imaging acquisition to the underlying executive. When a goal location approaches, the planner commits the imaging request to the executive along with the target coordinates and a tolerance about which the image can be acquired. With a much tighter control loop, the executive is able to more closely monitor the aerobot po-

sition and command the image at the appropriate time.

In both the rover and aerobot applications, the planner makes predictions about how long it will take to traverse to a target location. Both systems must be prepared to respond if things do not go as planned. For the rover, loose terrain or hazards can cause the vehicle to take longer than expected to reach the goal. For the aerobot, changing wind conditions can result in dramatic differences between estimated and actual time of arrival. Thus, both systems monitor traverse execution and revise the plan, perhaps shedding lower priority goals, if it appears that the vehicle is behind schedule. Alternatively, both the aerobot and rover could reach their goals ahead of schedule. In both cases, the system is able to update the plan to reflect the early arrival and the potential surplus of system resources that result. This surplus can then be allocated to goals that had not yet been included in the plan. The main difference between the aerobot and the rover is that wind conditions can change quite rapidly whereas terrain conditions and hazards are much more static. Thus, the aerobot system must monitor conditions at a higher frequency in order to respond in a timely manner.

Field Demonstration

The JPL Aerobot (Figure 7) is a robotic airship, i.e., a self-propelled lighter-than-air (LTA) vehicle developed by the JPL Aerobot team. The airship specifications are as follows: length of 11 m, diameter of 2.5 m, total volume of 34 m³, maximum speed of 13 m/s (25 kts), maximum ceiling of 3000 m (although, for safety reasons, it is flown below 1000 m), average mission endurance of 60 minutes, static lift payload of 12 kg ASL (at sea level), and dynamic lift payload of up to 16 kg ASL. The avionics and communication systems are installed in the gondola. The forward and aft ropes in Fig. 7 (a,b) are mooring lines for ground handling, not tether lines; i.e., the Aerobot flies free, not tethered.

The aerobot avionics is built around a PC-104+ computer architecture. The navigation sensors consist of an IMU (angular rates, linear accelerations), inclinometers (yaw, roll and pitch angles), magnetic compass, laser altimeter (surface relative altitude), barometric altimeter (absolute altitude against reference point), differential GPS (absolute 3D position) and ultrasonic anemometer (relative wind speed/direction).

The imaging sensors currently consist of a down-looking navigation camera and a wide-angle science camera. The navigation camera is used for image-based motion estimation (IBME), i.e., the ability to estimate the aerial vehicle trajectory in a GPS-denied environment through registration of sequences of images. The results are fed into an extended Kalman filter, together with other sensor estimates, to provide motion and positions estimates for the Aerobot. Other vision-based products include image mosaicking for mapping of large regions, and motion-based estimates of 3D terrain structure. The science camera is used for science image acquisition, science target selection, and visual go-to-target and stationkeep over target capabilities using a visual servoing approach to flight control. An onboard control switching system allows toggling between autonomous flight control and human pilot control. For safety reasons, the human pilot



Figure 7: The JPL Aerobot.

always has ability to override the SCS system and reassert control over the Aerobot.

In November, 2008, we successfully demonstrated this technology on the JPL aerobot platform at the Southern California Logistics Airport (SCLA) in Victorville, CA (Figure 8). The flight began with two imaging observation requests at different latitude/longitude coordinates submitted to the system from the “ground team.” The planner generated an operations plan to visit these locations after determining that there was sufficient time and energy to do so.

The winds were relatively strong and, for one of the observations, the aerobot reached the goal much sooner than initially predicted due to tail winds. Since the planner continually monitors the current state, this early arrival was detected and the image request was sent down for execution in time. The second goal required the aerobot to fly against the wind resulting in longer traverse time than the first goal. Again, the planner monitored the situation and waited until the aerobot came close to the target location before commanding the imaging.

While the aerobot was traveling toward the second goal, we simulated an onboard opportunistic science request, causing a new imaging goal to be sent to the planner. The planner correctly added the request to the plan and achieved this new goal after completing the second goal.

The flight demonstrated the system’s ability to generate mission operations plans given science goals, to command the aerobot control software and monitor plan execution, and



Figure 8: JPL Aerobot flying during field trials at the Southern California Logistics Airport (SCLA) in Victorville, CA.

to successfully modify the operations plan to respond to dynamic, opportunistic science events.

Conclusions

The AerOASIS system provides autonomous planning and execution capabilities for aerial vehicles. The system is capable of generating high quality operations plans that integrate observation requests from ground planning teams as well as opportunistic science events detected onboard the vehicle while respecting mission and resource constraints. We have successfully demonstrated the planning component in a field demonstration integrated with the JPL Aerobot platform. In future work we plan to extend the field demonstration with integrated, onboard science analysis.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the Jet Propulsion Laboratory Research and Technology Development program.

References

- Castano, R.; Estlin, T.; Anderson, R. C.; Gaines, D. M.; Castano, A.; Bornstein, B.; Chouinard, C.; and Judd, M. 2007. OASIS: Onboard Autonomous Science Investigation System for opportunistic rover science. *Journal of Field Robotics* 24(5):379–397.
- Chien, S.; Knight, R.; Stechert, A.; Sherwood, R.; and Rabideau, G. 2000. Using iterative repair to improve the responsiveness of planning and scheduling. In *Fifth International Conference on Artificial Intelligence Planning and Scheduling*.

Elfes, A.; Hall, J. L.; Montgomery, J. F.; Bergh, C. F.; and Dudik, B. A. 2004. Towards a substantially autonomous aerobot for exploration of titan. In *Proceedings of IEEE International Conference on Robotics and Automation*.

Elfes, A.; Hall, J.; Kulczycki, E.; Morfopoulos, A.; Montgomery, J.; Clouse, D.; Bayard, D.; Cameron, J.; Machuzak, R.; Magnone, L.; Bergh, C.; and Walsh, G. 2006. Autonomy capability development for a titan aerobot. In *Proceedings of the Fourth Annual International Planetary Probe Workshop (IPPW-4)*.

Estlin, T.; Fisher, F.; Gaines, D.; Chouinard, C.; Schaffer, S.; and Nesnas, I. 2002. Continuous planning and execution for a Mars rover. In *Third International NASA Workshop on Planning and Scheduling for Space*.

Gaines, D. M.; Estlin, T.; and Chouinard, C. 2008. Spatial coverage planning for a planetary rover. In *Proceedings of 2008 IEEE International Conference on Robotics and Automation (ICRA-08)*.

NASA. 2008. 2008 solar system exploration roadmap. Technical report, NASA, Washington, DC.

Zweben, M.; Daun, B.; Davis, E.; and Deale, M. 1994. Scheduling and rescheduling with iterative repair. In Fox, M., and Zweben, M., eds., *Intelligent Scheduling*. Morgan Kaufmann Publishers Inc. 241–256.