NIAC Phase 1 Final Study Report
on
Titan Aerial Daughtercraft

Larry Matthies
Jet Propulsion Laboratory
California Institute of Technology

Contents

1   Introduction

1.1   Motivation and Concept Description

Saturn's giant moon Titan has become one of the most fascinating bodies in the Solar System. Even though it is a billion miles from Earth, data from the Cassini mission reveals that Titan has a very diverse, Earth-like surface, with mountains, fluvial channels, lakes, evaporite basins, plains, dunes, and seas [Lopes 2010] (Figure 1). But unlike Earth, Titan's surface likely is composed of organic chemistry products derived from complex atmospheric photochemistry [Lorenz 2008]. In addition, Titan has an active meteorological system with observed storms and precipitation-induced surface darkening suggesting a hydrocarbon cycle analogous to Earth's water cycle [Turtle 2011].
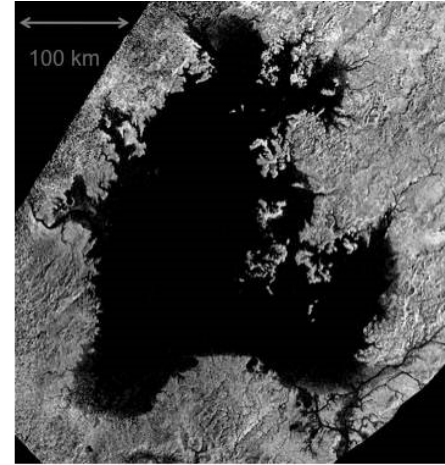

Figure 1. Ligeia Mare sea on Titan, with river channels leading into it.

Titan is the richest laboratory in the solar system for studying prebiotic chemistry, which makes studying its chemistry from the surface and in the atmosphere one of the most important objectives in planetary science [Decadal 2011]. The diversity of surface features on Titan related to organic solids and liquids makes long-range mobility with surface access important [Decadal 2011]. This has not been possible to date, because mission concepts have had either no mobility (landers), no surface access (balloons and airplanes), or low maturity, high risk, and/or high development costs for this environment (e,g. large, self-sufficient, long-duration helicopters). Enabling *in situ* mobility could revolutionize Titan exploration, similarly to the way rovers revolutionized Mars exploration.

Recent progress on several fronts has suggested that small-scale rotorcraft deployed as "daughtercraft" from a lander or balloon "mothercraft" may be an effective, affordable approach to expanding Titan surface access. This includes rapid progress on autonomous navigation capabilities of such aircraft for terrestrial applications and on miniaturization, driven by the consumer mobile electronics market, of high performance of sensors, processors, and other avionics components needed for such aircraft. Chemical analysis, for example with a mass spectrometer, will be important to any Titan surface mission. Anticipating that it may be more practical to host chemical analysis instruments on a mothership than a daughtercraft, we defined system and mission concepts that deploy a small rotorcraft, termed a Titan Aerial Daughtercraft (TAD), from a lander or balloon to perform high-resolution imaging and mapping, potentially land to acquire microscopic images or other *in situ* measurements, and acquire samples to return to analytical instruments on the mothership. In principle, the ability to recharge batteries in TAD from a radioisotope or other long-lived power source on the mothership could enable multiple sorties.

For a lander-based mission, a variety of landing sites is conceivable, including near lake margins, in dry lake beds, or in regions of plains, dunes, or putative cryovolanic or impact melt features. Such missions may require landing with greater precision than in previous missions (Huygens) and mission studies; this could also enhance the ability of TAD to reach interesting terrain from the landing site. Precision descent may also benefit balloon missions, with or without a daughtercraft, by increasing the probability that the balloon will drift over desired terrain early in its mission. Given these potential benefits, the overall concept studied here

includes brief consideration of precision descent for landing or balloon deployment, followed by one or more sorties by a rotorcraft deployed from the mothership, with the ability to return to the mothership.

This concept revolutionizes previous mission concepts in several ways. For a lander mission, it enables detailed studies of a large area around the lander, with the possibility of sampling from more than one location. With precision landing near a lake, it enables acquiring solid and liquid samples with one lander. For a balloon mission, it enables surface investigation and sampling with global reach without requiring that the balloon be brought to the surface. This enables surface science from a balloon without requiring a separate lander or landing and takeoff of the fully instrumented balloon, which has potential for major risk reduction and cost savings. Both scenarios could involve repeated sorties if recharging from the mothership is feasible.

1.2 Prior Work

Prior Flagship mission studies have proposed a passively-navigated, wind-borne balloon or a short-lived lander, with the balloon kept several kilometers from the surface for safety reasons and the lander accessing only one location [TE 2007, TSSM 2009]. Recent research efforts have developed actively navigated airships to enable reaching specific locations, as well as the ability to descend to very near the surface to acquire samples with a harpoon-like device [Badescu 2009, Aboudan 2010]. However, such airships have appeared to be costly and risky, especially for descent to the surface for sampling. For large, independent aircraft (100 to 400 kg) for long-range missions, a variety of fixed wing, rotary wing, or hybrid aircraft have been studied [Lafleur 2006, Young 2001, Gasbarre 2005, Lorenz 2000, Prakash 2006, Barnes 2012]. In [Young 2001], the study includes a brief examination of the aerodynamics, power, and range of a tilt-rotor aircraft in the size range of 10 to 50 kg. The Titan Explorer flagship mission study contemplated an outreach payload on the lander that was a 2 kg unmanned air vehicle [TSSM 2009]; this was conceived as having little autonomy and being unlikely to contribute to science, and no design information was presented. A "Titan Bumblebee" concept examined in [Lorenz 2008] was a fixed-wing, battery-powered UAV of approximately 1 kg, conceived as taking off vertically from a lander, flying a pre-programmed route using inertial and/or optical navigation, and transmitting imagery back to the lander until its battery was exhausted. Aerodynamic, power/energy, and thermal issues were studied in some detail in that paper, but avionics and navigation algorithms were not.

In the past 10 years, there has been an explosion of academic research on autonomous navigation of small rotorcraft for indoor and outdoor operation on Earth, and in the last 5 years there has been a significant expansion of industrial development of such aircraft for commercial and defense applications. Very often these aircraft are multicopters with 4, 6, or 8 rotors, due to the relative simplicity of their mechanical design and control compared to other types of helicopters [Powers 2014]. Avionics architectures for such aircraft have dramatically grown in performance and shrunk in size, due to cross-fertilization from the smartphone industry; this has enabled a corresponding growth in the performance of algorithms for autonomous navigation. The robotics research subfield of "simultaneous localization and mapping" (SLAM), which previously focused largely on ground vehicles, has shifted focus substantially into aerial vehicle applications. Although this is still a very active area of research, there are now a number of relatively mature algorithms that use a smartphone-class processor, a single onboard camera, an inertial measurement unit (IMU), and in some cases an altimeter to do autonomous onboard position, velocity, and attitude estimation for rotorcraft on the order of 1 kg [Li 2013, Weiss

2013, Forster 2015, Leutenegger 2015]. Autonomous safe landing in unknown, hazardous terrain has also been demonstrated by small robotic rotorcraft using only an IMU, an altimeter, and one camera for terrain sensing [Johnson 2005, Brockers 2014, Forster 2015]. A JPL proposal for a payload on the 2020 Mars rover [Golombek 2014] presented a concept for a 1 kg Mars Helicopter Scout that packaged the avionics necessary for autonomous flight in a very small, lightweight, low power module with adequate thermal insulation, leveraging commercial electronics in a fault-tolerant architecture.

Small-scale, long-life motor, bearing, and lubricant technology for cold temperatures has been studied for rovers on Mars and asteroids. Very lightweight liquid sampling is feasible with techniques derived from recent internal JPL research on Titan sampling [Zimmerman 2013].

1.3   Study Objectives and Structure of Report

While there has been related prior work, conceptual design has never been done for the type of autonomous Titan daughtercraft mission addressed here. To maintain a plausible size of the daughtercraft relative to other potential payloads and the total mass of the mothership, we limited the total mass of the daughtercraft to on the order of 10 kg. Sizing studies have never been done at this scale to assess propulsion system characteristics, power, range, endurance, and payload capability of the daughtercraft. There has been no previous work on sample acquisition from a vehicle this light or on how to transfer samples from such a vehicle to instruments on the balloon or lander.

Based on this, our original objectives for this study were to:

1. Develop concepts of operations (CONOPS) for lander-based and balloon-based daughtercraft mission scenarios.

2. Develop and refine a parametric sizing model for both lander-based and balloon-based scenarios to characterize propulsion, power, endurance, range, and payload capability for total daughtercraft masses roughly between 1 and 10 kg.

3. Develop a conceptual design and identify representative components for the entire daughtercraft hardware and software system and the autonomous mobility capabilities needed for both mission scenarios, including estimates of approximate mass, volume, power, energy, and communication budgets and preliminary CAD model.

4. Develop a conceptual design and preliminary CAD model for a science payload on the daughtercraft, including specifying a nominal instrument suite on the balloon or lander, designing a compatible sampling mechanism to acquire solid and/or liquid samples on the daughtercraft, and designing mechanisms and daughtercraft behaviors necessary to transfer the samples to the instruments

This plan was modified after the study began in response to:

- Feedback from NIAC program management that suggested higher prioritization of the required autonomy elements over proposed aspects of the study. This led to reduced emphasis on hardware-related aspects, including the payload.

- Feedback from members of the Titan science community that suggested focus on lander scenarios had more chance of impacting a mission in the next two decades than a balloon scenario, due to relatively immaturity and high cost of balloon systems compared to landers. After the sizing models were complete, this led to curtailing further study of balloon-based scenarios.

- Feedback from members of the Titan science community that emphasized high science value of studying of Titan's potential for prebiotic chemical processes and suggested that liquid/solid interfaces are likely to be especially good places for such study. This led to increased emphasis on technical approaches that might enable access to lake margins.

In an attempt to save mass, our original concept for sampling used one integrated device to sample both solid and liquid material, which required either landing on or hovering within a few centimeters of the surface. The practicality of hovering so close to a lake surface began to appear questionable, whereas other methods of liquid sampling, e.g. lowering a sample container into a lake while hovering higher, appeared more practical but would lead to separate mechanisms for solid and liquid sampling. In keeping with the earlier decision to put more emphasis at this time on showing the feasibility of autonomous navigation capabilities, we report work on the original sampling concept, but we did not study alternatives.

As the study progressed, we perceived increased value of precision landing in this concept and perceived that it might be practical to develop. This led to some study of this issue, where none was originally planned.

As a result, the rest of this report addresses topics in the following order:

- CONOPS for lander-based and balloon-based daughtercraft mission scenarios.
- Parametric daughtercraft sizing models for both lander-based and balloon-based scenarios to assess range versus total mass and diameter for a given payload allocation.
- Conceptual design of the daughtercraft, mainly focused on autonomous navigation in a lander-based scenario, with experimental demonstration of prototype autonomous localization and landing hazard detection capabilities implemented at JPL and a discussion of related work elsewhere.
- Brief summary of the sampling device concept.
- Preliminary assessment of issues and potential approaches to precision landing and automatic, onboard localization of landers and balloons after landing/balloon deployment, as a counterpart to assessments of range versus size of the daughtercraft.
- Summary, conclusions, and recommendations.

2 Concepts of Operation

A common element in all scenarios is that we assume the availability of remote sensing imagery of the landing error ellipse from the Cassini radar, Visual and Infrared Mapping Spectrometer (VIMS), and Imaging Science Subsystem (ISS). These instruments have effective resolutions on the order of 300 m/pixel or larger for radar and 1 km/pixel or larger for VIMS and ISS, after accounting for scattering in Titan's thick atmosphere. This data is critical for a number of reasons, as discussed below.

We assume a level of onboard computing power in both the mothership and the daughtercraft that is at least comparable to the multicore processors now used in smartphones and notebook computers. This level of capability for spaceflight computers is now under development in the form of rad-hard multicore processors in the High Performance Spaceflight Computing (HPSC) project of NASA's Game Changing Development Program, which targets mission infusion

potential in the early 2020s [Doyle 2014]. The benign radiation environment at the surface of Titan raises the question of whether avionics components that are not strictly rad-hard may provide additional options, if sufficient survivability can be shown for the cruise phase of the mission. Evaluating this possibility is beyond our scope.

2.1   Lander-based Scenario

Given a lander with a descent camera, a good inertial measurement unit (IMU), and the level of onboard computing power envisioned above, it is very plausible that the position of the landing site could be determined autonomously onboard, after landing, with an error of a few kilometers relative to Cassini map imagery. This would be achieved by onboard matching of descent imagery to Cassini map imagery, in combination with other techniques; we discuss this further in Section 4. With knowledge of the landing site at this level of accuracy, TAD sorties could be targeted to specific sites of scientific interest, through a combination of prior planning on Earth and onboard planning after landing. In principle, descent imagery and a guided parachute may also enable landing much more precisely than has previously been considered; this may be important to ensure that locations critical to the science objectives are within flight range of the landing site.

Titan landers have been developed or proposed in the past as "short-lived", powered only by a battery, or "long-lived", with a radioisotope power system (RPS). For a short-lived lander, we envision one or two TAD sorties, where the first is planned based on prior human analysis of the scientific potential of the landing ellipse and accurate knowledge of the actual landing site, computed onboard as discussed above. A total flight distance of many 10s of kilometers is plausible (Section 3). Data from TAD would be transmitted to the lander for downlink.

Accurate position, velocity, and heading knowledge of TAD during flight would be needed to reach designated science waypoints and to return to the lander; this would be achieved with a suitable combination of sensors and algorithms (Section 4). Autonomous safe landing for sampling, microscopic imaging, or operating other contact instruments would use onboard depth perception, terrain relative motion estimation, motion planning, and control algorithms for mapping the terrain to find safe landing sites and to achieve safe touchdown. For some specific terrain types, particularly lakes, it will be possible to automatically recognize those using onboard sensors, accurately delineate their boundaries, and navigate accurately relative to them for mapping, sampling, or other purposes. Lake sampling may involve hovering accurately over the lake near the lake margin to lower a sample acquisition vessel into the lake. Returning to the lander would again use a suitable combination of sensors and algorithms, plus visual reference marks on the lander to enable accurate delivery of samples to instruments. Rotor-induced airflow will require high performance state estimation and control when very close to the lander.

If two sorties were possible, the first might be used to obtain high resolution images and maps for downlink to Earth, and the second to revisit specific locations for landing and/or sampling, based on ground-based analysis of telemetry from the first sortie.

TAD missions with multiple sorties may be more viable for long-lived landers, where TAD docks with the lander between sorties to obtain power for survival heating and battery recharging from the RPS on the lander. In this case, multiple sorties could be used for exploration to the maximum radial range in different directions, to acquire samples from different sites, or to use initial sorties for imaging and mapping and later sorties for landing and sampling.

2.2   Balloon-based Scenario

In keeping with past balloon mission studies [TE 2007, TSSM 2009], we assume a balloon hovering at up to 10 km altitude, drifting with the wind at approximately 1 to 2 m/s. Based on past balloon studies, we anticipate a gondola on the order of 1.5 to 2 m in diameter with instruments mounted on the side of the gondola, which allows TAD to dock to the bottom of the gondola. Since a balloon would have a relatively long life and drift over a large distance, multi-sortie capability for TAD would be extremely desirable; thus, while docked TAD would receive power from the balloon RPS to keep its electronics warm and maintain battery charge.

Balloon position knowledge would be maintained by correlating terrain images acquired on the balloon with prior Cassini maps of the surface [Ansar 2009], possibly aided with other sensors; this and atmospheric circulation models would enable prediction of the balloon's approximate path. Candidate surface science targets would be designated on Earth by the science team using prior remote sensing data from Cassini, based on the predicted path of the balloon. These candidates could be approximate locations or they could be classes of terrain, described by appearance in remote sensing instruments and later recognized onboard the balloon by autonomous terrain classification. Due to the latency of round-trip communication with Earth, final target locations would be determined autonomously onboard the balloon using remote sensing data from instruments on the balloon.

TAD operation must be fully autonomous. It descends to the surface, using radio navigation aiding from the balloon and onboard, image-based landmark recognition and terrain classification to reach the target. Once near the surface, imaging, mapping, sampling, and operating other instruments would be conducted similarly to the case of a lander-based mission. Return to the balloon would be similar to the lander-based scenario, except that the TAD power and energy budget must account for the altitude difference and that the balloon will have moved during the sortie, which requires more sophistication in the TAD navigation sensor suite and algorithm. Given the low winds expected within 10 km of the surface and some elapsed time after balloon deployment to damp out transient motions, the balloon is unlikely to have significant swaying motion or rotation, so sample transfer and docking should primarily have to account for steady translational motion of the balloon, plus disturbances causes by the rotor-induced airflow from TAD itself.

In comparison with the lander-based scenario, the balloon-based scenario involves more challenging navigation and control problems, may require more onboard autonomy for designating destinations for each sortie, and derives much more value from the ability to fly multiple sorties.

3   Daughtercraft Sizing Models

3.1   Overall Methodology

In this study, the purpose of the sizing models is to characterize trade-offs between range, size, and payload capacity of the daughtercraft. We limited this to one type of rotor configuration, and chose a multicopter with four rotors (i.e. quadrotor) for its mechanical and control simplicity. For fault tolerance, a flight system might benefit from having more rotors; this was out of scope here, but it would be an easy generalization of the models already developed.

The models were developed by AeroVironment, based on their extensive experience with air vehicles, and captured as Excel spreadsheets that enable changing parameters to vary

assumptions. Since there are potentially a large number of variables to experiment with, and since our goal was to get initial insight rather than to do an extensive exploration of the parameter space, we defined a methodology that fixed most parameters, leaving one independent and one dependent variable at a time that we could examine easily with 2-D plots. Before this study started, JPL had defined a conceptual

Table 1. Key parameters of rotorcraft performance model

Disk loading	50 Pa
Tip Mach number	0.2 to 0.25
Rotor figure of merit	0.75
Motor efficiency	80%
Controller efficiency	95%

avionics architecture for a Mars helicopter, with mass and power breakdowns, that appeared to be applicable in broad terms to our Titan scenarios. Therefore, for this study we made a fixed mass and power allocation of 0.5 kg and 20 W for avionics, based on the conceptual design of the Mars version. Similarly, based on our initial concept for a sampling system, we made a fixed mass and energy allocation of 2 kg and 1.7 Wh for the payload. Then, for a given total mass, we:

- Sized the rotors and motors based on thrust required for flight, using nominal values that experience suggested were appropriate for disk loading, rotor figure of merit, tip Mach number, and other parameters;

- Estimated overall aircraft diameter based on rotor diameter;

- Estimated structural mass as a function of total mass (20%);

- Allocated all remaining mass to the battery, assuming a lithium ion battery with specific energy of 100 Wh/kg. This also assumed that the battery is keep sufficiently warm that this specific energy value is viable, which is the same approach taken for Mars helicopter studies.

With these quantities specified, we used total aircraft system mass as the independent variable to examine possible flight radius for a lander-based scenario and available payload mass fraction for a balloon-based scenario, as described below. Aerodynamic models and parameters used in this analysis were described briefly in [Tokumaru 2014] and are captured in our spreadsheets. Table 1 lists key parameters that were common to both the lander-based and balloon-based scenarios. In short, with about 1/7th the gravity and almost 5 times the atmospheric density of Earth at sea level, Titan is much more favorable to flight than Earth. Other discussions of the aerodynamics of flying vehicles for Titan (and Mars and Venus) are given in [Lorenz 2001, Young 2001, Gasbarre 2005, Lorenz 2008a].

3.2 Lander-based Scenario

This case modeled forward velocity as 8 m/s to limit drag losses. The maximum mission flight radius from the lander was estimated using the procedure described above, assuming that all mission time was allocated to forward flight at a fixed velocity and that the battery was allocated all mass not accounted for in the procedure described above. This simple model gives an initial idea of the order of magnitude flight radius possible, and can be de-rated to make allowances for science activities, additional payload, mass margin, or more conservative estimates of component masses. Figure 2 shows mission radius as a function of total system masses between 4 and 12 kg. The corresponding rotor radius is shown across the top of the graph; total vehicle radius would be about three times the rotor radius. Even if much more conservative model inputs were used, it appears that the flight radius for a 10 to 12 kg rotorcraft can be several 10s of kilometers.

Another way to use this graph is to pick a desired flight radius, then use the horizontal axis to get estimates of the minimum total mass and rotor diameter that would enable that flight radius.

Larger values for total mass would translate into mass available to change other aspects of the model, such as to add mass to components, add payload, or to increase the battery size to provide more energy for science operations. For example, for a flight radius of 60 km, the graph implies that a 12 kg system would have room to double the mass over the estimated minimum value.

3.3 Balloon-based Scenario

For this case, we assumed that the balloon drifts at constant velocity of 2 m/s at an altitude of 10 km. We used the same mass and power allocations for avionics and the same mass and energy allocations for payload (the sampler) as in the lander-based scenario. We assumed that the rotorcraft descends from the balloon to the surface by auto-rotating, so that the descent does not consume power for propulsion. A rotorcraft velocity of 8 m/s was assumed for the descent to the surface and the ascent back to the balloon. To complete the mission timeline, we allocated 5 minutes for hovering at the surface, 5 minutes for sampling activities, and 5 minutes for hovering while docking at the balloon; this gave a total sortie duration of about 57 minutes. For any given total system mass, we estimated the energy required for the entire scenario and allocated a corresponding battery mass. For the model assumptions and the system mass range we used, this left some unallocated mass, which we treated as mass margin. Figure 3 shows the mass margin as a function of total system mass, again with the corresponding rotor diameter across the top of the graph. In this case, a 12 kg system would have 100% mass margin over the values used in the model. Lower balloon altitudes would give correspondingly more latitude to the design of the rotorcraft and its mission, which could be used in many ways: to provide more mass margin, increased sortie duration or lateral range from the balloon flight path, or to make the rotorcraft smaller.
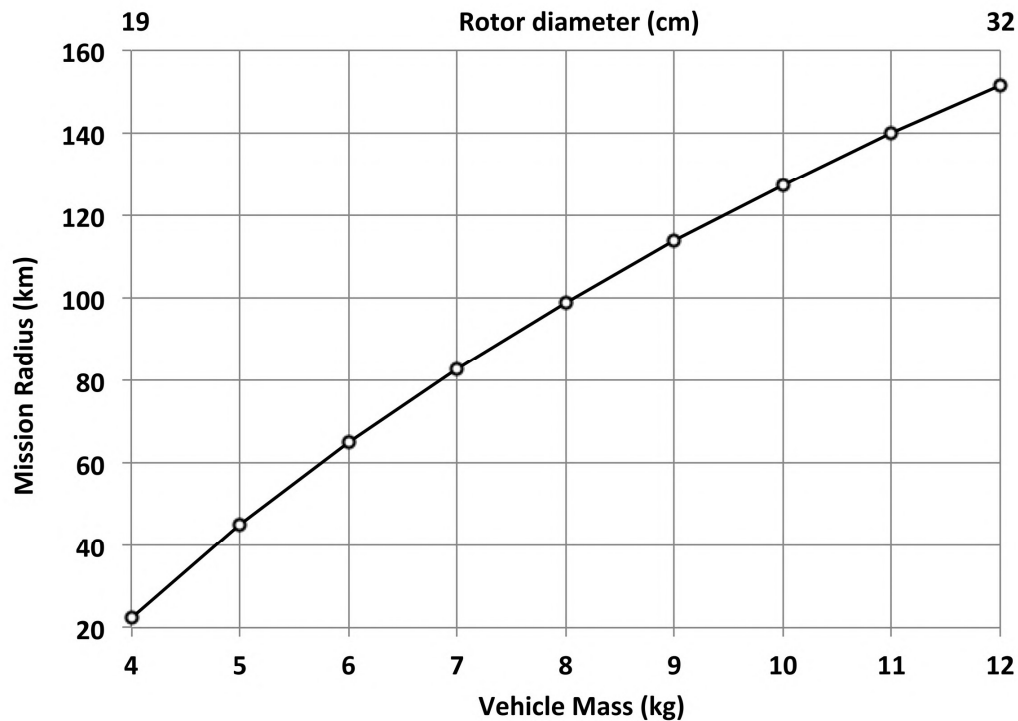

Figure 2. Maximum rotorcraft flight radius from a lander as a function of total system mass. The range of rotor diameters used in the model is shown across the top of the graph.

4   Conceptual Design of the Autonomous Daughtercraft

As already noted, we used a four-rotor multicopter (quadrotor) concept for the mobility mechanism and sizing models, and focused most attention on autonomous navigation and a conceptual architecture for the avionics. Based on informal assessments conducted during the preparation of JPL's Mars Helicopter Scout proposal [Golombek 2014], there appears to potential that commercial grade components could be used for key aspects of the electronics, particularly processors and inertial sensors that can be housed inside an insulated enclosure. Aside from that, environmental survivability of electronics is out of scope for this study.

4.1   Autonomous Navigation

The term "autonomous navigation" is used here in the broad sense that is common in the robotics literature, and includes state estimation, terrain perception, and motion planning and control. We address these separately.

4.1.1   State Estimation

The following elements of traditional navigation state must be estimated onboard:

- Attitude (pitch/roll)
- Absolute bearing and bearing from TAD to the mothership
- Altitude above ground level, including above lake or sea level when appropriate
- Absolute and mothership-relative position
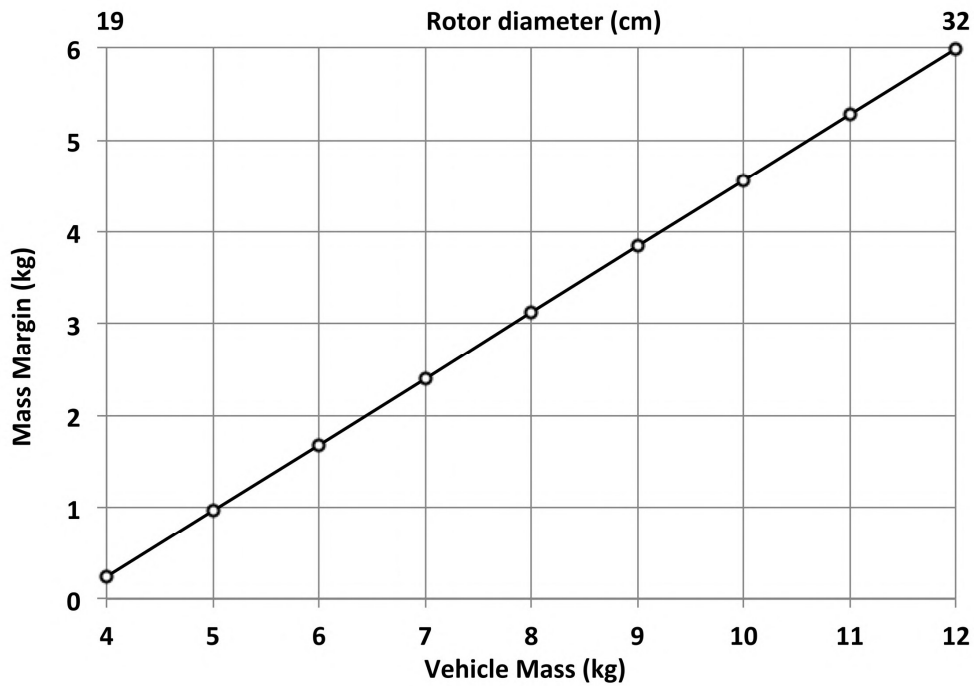- Absolute and mothership-relative velocity

Figure 3. Rotorcraft mass margin as a function of total system mass, for deployment from and return to a balloon at 10 km altitude

Attitude estimation is routinely done for small terrestrial drones with MEMS gyros and accelerometers. The same solution should apply for Titan.

Bearing estimation is necessary to reach specific designations, including returning to the mothership. To some extent, this can be subsumed within visual simultaneous localization and mapping (VSLAM) capabilities, which enable navigation by selecting and tracking visual features of the environment as landmarks, as discussed below. However, bearing knowledge with VSLAM degrades with distance traveled, so reaching specific points at large distances from the mothership may require other methods for bearing estimation. Using a radio beacon on the mothership is possible, though its useful range would be limited by the curvature of Titan's surface. The Huygens probe used a photodiode-based sun sensor to monitor its azimuth angle relative to the sun during descent; this was designed to operate down to the surface [Tomasko 2002], but it only operated to an altitude of 30 km due to radiation-induced loss of sensitivity during cruise and other factors [Karkoschka 2007, DISR Archive 2013]. It may be possible to design a sun sensor for bearing estimation that would work at the surface. If sorties from a lander record visual landmark on the outward flight segment, then onboard visual landmark recognition would enable return to the lander without other bearing sensors.

Altitude estimation for terrestrial drones is typically done with optical or acoustic range sensors, or with barometers to measure changes in altitude over short flights. Optical range sensors on the order of 100 grams with ranges of several kilometers are available commercially from several sources, including FLIR Systems and Voxtel. In addition to challenges of surviving the cold, the optical reflectance of Titan's surface at the wavelength of these sensors would have to be studied to estimate their performance. For altitude estimation over lakes or seas, for example for sampling liquids from TAD, when the aircraft is very near the shore visual feature tracking of features on the ground can provide altitude knowledge as a byproduct of VSLAM. Given the thick atmosphere, acoustic altimetry over lakes and seas is plausible, provided that the sensor operates at Titan temperature.

Position and velocity estimation would use inertial sensors aided by all of the above sensors as well as by onboard visual feature tracking (i.e. VSLAM). Outside the area seen in the Huygens probe images of the surface, it is unknown whether and where the surface of Titan is sufficiently rich with visual features to support visual feature tracking; however, since visual feature tracking as been possible for Mars rovers in most places on Mars, it is plausible that it would be possible for landing sites of interest on Titan. Sensor suites and operations can also be defined conservatively so that minimal mission success criteria could be achieved by relying on non-visual sensors, and more expansive objectives could be addressed by autonomous decision-making onboard, once the system understands how well it can navigate.

Docking with the mothership is a special case of navigation, where a combination of a radio beacon and visual fiducial marks on the mothership could aid accurate position and velocity estimation relative to the mothership. As noted earlier, for liquid sampling by hovering over a lake or sea very close to the shore, position and velocity knowledge could be maintained by visually tracking features on the shore or the shoreline.

Knowledge of the absolute position of the mothership on Titan would be important to enable TAD to reach specific destinations selected by mission planning with prior remote sensing imagery from Cassini or a future orbiter. In principle, this could be achieved onboard the mothership with a precision on the order of 1 km by automatically registering descent images to remote sensing images, even where there are significant differences in the spectral bands of the

image sets [Ansar 2009]. Alternatively, science targets for TAD could be designated with descent images, then TAD could navigate relative to visual landmarks in the descent images with requiring knowledge of the absolute position of a lander.

Proof of concept demonstration. Accurate visual-inertial navigation of small drones is becoming a very mature capability on Earth, using onboard cameras, inertial sensors, and processors from the smartphone industry, sometimes aided with compact altimeters [Weiss 2013, Li 2013, Forster 2015, Leutenneger 2015]. We conducted a simple proof of concept demonstration of this with a 235 m loop flight using a testbed quadrotor flying in the Arroyo Seco dry river wash next to JPL. This used an onboard IMU and a single nadir-pointed camera for state estimation with the algorithm described in [Weiss 2013]. Figure 4 illustrates the terrain and the state estimation performance. In this case, there was an error of about 7% in estimating the scale of the landmark map due to an initialization error; however, this cancels out in flights that return to the origin, as seen in the figure. Such scale errors can be eliminated in several ways.
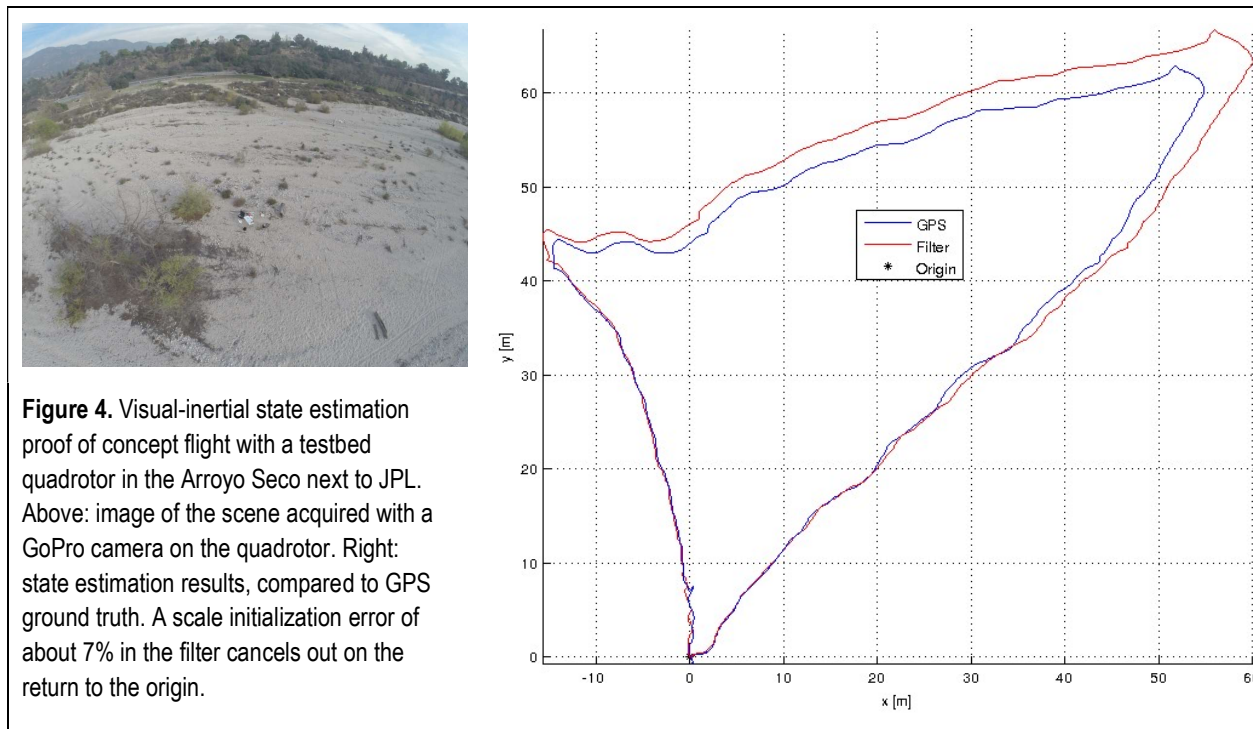

Figure 4. Visual-inertial state estimation proof of concept flight with a testbed quadrotor in the Arroyo Seco next to JPL. Above: image of the scene acquired with a GoPro camera on the quadrotor. Right: state estimation results, compared to GPS ground truth. A scale initialization error of about 7% in the filter cancels out on the return to the origin.
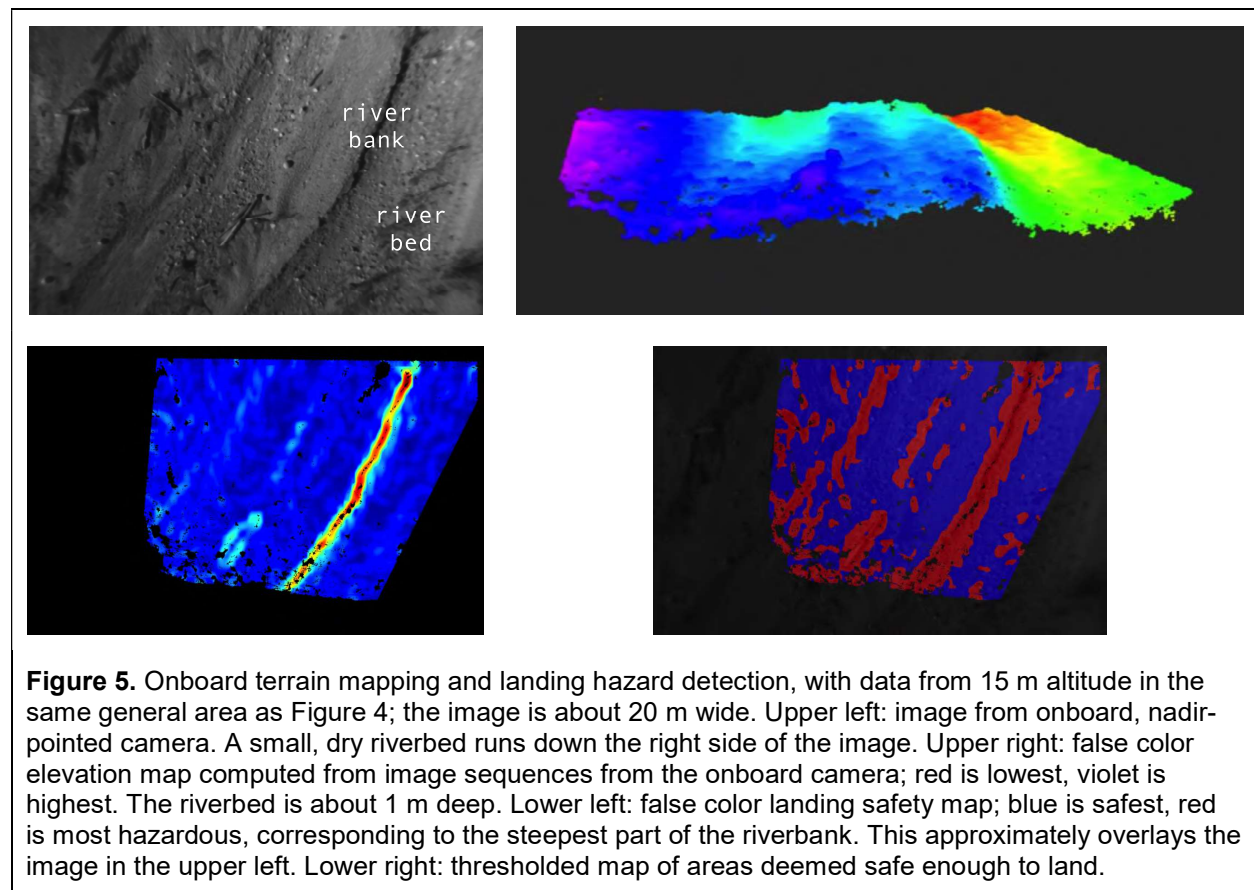
4.1.2 Terrain Perception

Onboard terrain perception capabilities of significance here are:

- Obstacle detection during flight
- Landing hazard detection
- Onboard science target recognition

Obstacle detection during flight is required to avoid collisions with large terrain features while flying. Obstacle detection is normally done with onboard 3-D perception, using range sensors such as lidar, radar, or passive ranging by triangulation with multiple images. This has been done on Earth with full-size, unmanned helicopters that use lidar to create onboard elevation maps [Whalley 2013], with small multicopter drones that use lidar and stereo vision to

create elevation maps and 3-D voxel representations of obstacles [Droeschel 2015], and with small quadrotors that use only a single camera with visual "structure from motion" algorithms for depth perception [Alvarez 2014]. A very small, 230 GHz electronically beam-steered radar is also under development for terrestrial drone applications [Sarabandi 2011]. As a result of this work, algorithms for creating obstacle representations are fairly mature; the main issue for Titan would be developing sensors (camera, lidar, and/or radar) suitable for the environment.

Landing hazard detection focuses on high resolution onboard mapping of the terrain immediately under the aircraft to find a safe landing site. This can be viewed as a special case of in-flight obstacle detection, though landing hazard detection can use just an elevation map, as opposed to 3-D voxel maps needed in terrestrial applications in complex 3-D environments. Real-time landing hazard detection and avoidance has been demonstrated with small helicopters and quadrotors using a single camera for terrain mapping, with algorithms running on smartphone-class processors [Johnson 2005, Brockers 2014, Forster 2015a].

Figure 5. Onboard terrain mapping and landing hazard detection, with data from 15 m altitude in the same general area as Figure 4; the image is about 20 m wide. Upper left: image from onboard, nadir-pointed camera. A small, dry riverbed runs down the right side of the image. Upper right: false color elevation map computed from image sequences from the onboard camera; red is lowest, violet is highest. The riverbed is about 1 m deep. Lower left: false color landing safety map; blue is safest, red is most hazardous, corresponding to the steepest part of the riverbank. This approximately overlays the image in the upper left. Lower right: thresholded map of areas deemed safe enough to land.

Onboard science target recognition is necessary for several reasons. First, the very low resolution of currently available remote sensing images, with effective pixel sizes of at least 300 m/pixel for radar and at least 1 km/pixel for near infrared images, limits the ability to accurately designate locations of science targets via mission planning on Earth. Second, the navigation performance of the rotorcraft may not be accurate enough to reach some targets. Third, the round-trip communication latency from Titan to Earth and back to Titan may preclude downlinking images acquired *in situ* to plan the next sortie on Earth. Specific examples of such science targets include distinguishing precipitated organics from bedrock in inter-dune areas,

following river channels, and discriminating lakes and seas from ground for flights to shorelines or following shorelines. Such capabilities have been demonstrated in analogous terrain on Earth [Matthies 2003, Helmick 2009, Rankin 2011, Nuske 2015]. Enabling these capabilities for Titan would require analyzing Titan sensor phenomenology and observable features to determine appropriate recognition methods.

Proof of concept demonstration of onboard landing hazard detection. Landing hazard detection and avoidance has been shown on a testbed quadrotor at JPL, with an onboard sensor suite consisting of an IMU, an altimeter, and a single nadir-pointed camera that is used to create terrain elevation maps onboard from the sequence of images acquired as the aircraft flies over terrain [Brockers 2014]. This capability runs on a smartphone-class processor at about 1 image/second, processing images with 320x240 pixels. Figure 5 illustrates the onboard mapping capability with a data set acquired while flying at about 15 m over terrain the Arroyo Seco, in the same general area as shown in Figure 4.
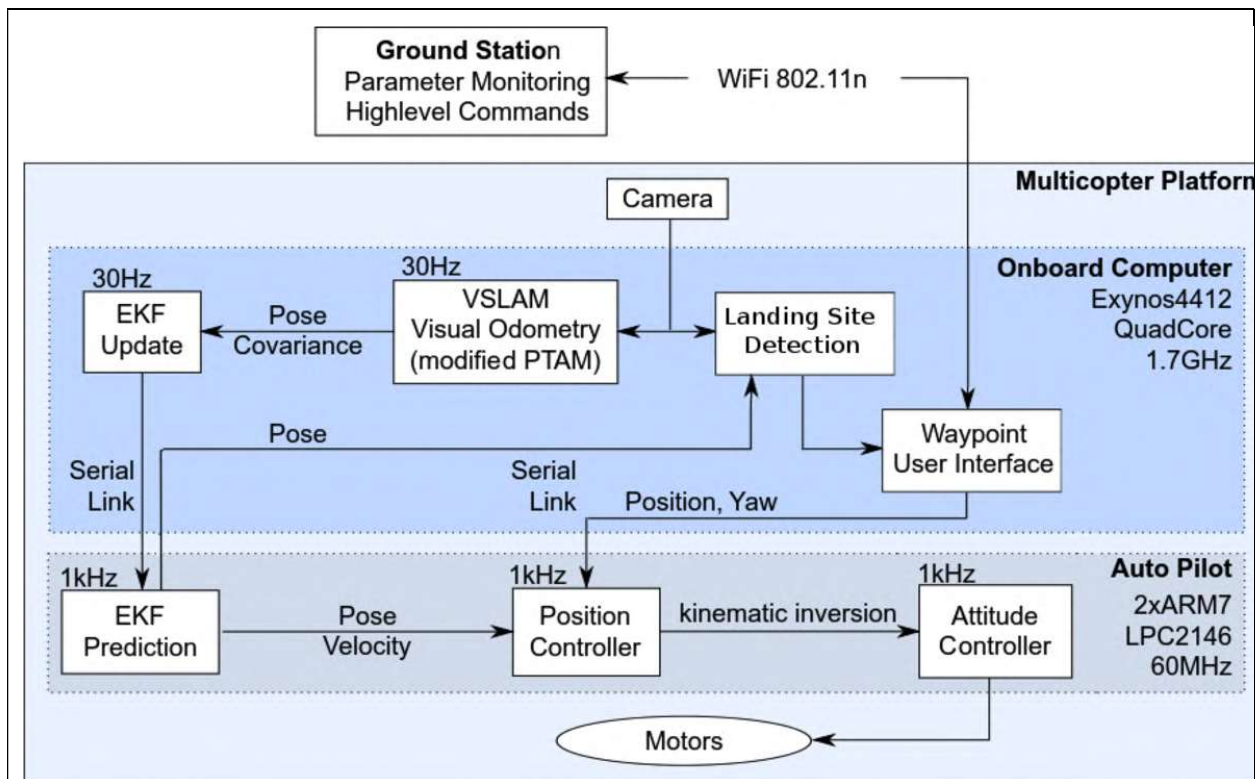

Figure 6. Two-level processor architecture on a 500 gram JPL autonomous quadrotor.testbed aircraft, which is representative of typical drone processor architectures. A low-level microcontroller performs real-time autopilot functions, including attitude estimation, motor control, and state vector propagation, in this case at an update rate of 1 kHz. A high-level computer with a multicore ARM Cortex-based processor performs visual feature tracking for localization (at a frame rate of 30 Hz) and terrain mapping for landing site determination (at a rate of 1 Hz). In this case, navigation is done with one nadir-pointed camera, an IMU, and a barometric altimeter [Brockers 2014].

4.1.3 Motion Planning and Control

For a lander-based mission, motion planning and control for TAD is relatively straightforward while flying point to point, since it largely boils down to maintaining altitude; most of the challenge is in state estimation and terrain perception. The low winds at the surface of Titan (~ 1

m/s) should make motion planning and control relatively straightforward for landing on terrain. Docking with the mothership requires greater precision and tighter control, which may require further research beyond capabilities already demonstrated.

4.2 Avionics

As described above, the essential autonomous navigation functions needed for TAD -- state estimation and safe landing -- can now be implemented with the level of computing capability available in current smartphones, which use multicore system-on-a-chip (SoC) processors with four or more ARM Cortex A9 or A15 or equivalent processor cores. Terrestrial drones typically use a two-level processor architecture, in which a microcontroller performs critical real-time "autopilot" functions, especially attitude estimation and motor control, and a higher performance, multicore processor performs compute-intensive navigation functions, including image processing, position estimation, mapping, and motion planning. Figure 6 illustrates how these functions are supported in recent JPL work with a 500 gram quadrotor testbed aircraft carrying a 15 gram high-level processor [Brockers 2014].

Concurrent with this study, JPL has been doing system design and prototype testing toward a potential technology demonstration of a roughly 1 kg helicopter for Mars [Golombek 2014]. The avionics concept for the Mars helicopter provides a conceptual model of avionics for a Titan rotorcraft. The Mars system uses a number of commercial off-the-shelf (COTS) components, including processors, cameras, and its radio, due to many constraints facing that system. These are kept sufficiently warm in an electronics box

Figure 7. Conceptual model of avionics assembly for a 1 kg Mars helicopter, as originally conceived in [Golombek 2014]. This provides a notional model of avionics for a Titan rotorcraft.

insulated with approximately 3 cm of aerogel and are expected to survive the relatively benign radiation environment on the Mars surface. The processor architecture has the same two-level structure described above, with a real-time flight computer for autopilot functions and a multicore smartphone processor for navigation functions. Although single event effects (SEE) from radiation are expected to be rare, the flight computer is protected against upsets by using a pair of dual-redundant microcontroller chips and a rad-hard FPGA that provide for fault detection and a hot spare in the event of an upset. Software-based fault tolerance is expected to be adequate protection for the navigation computer. The processor architecture, IMU, and other electronics are mounted on printed circuit cards that surround a rechargeable lithium ion battery. All of this is surrounded by the aerogel and enclosed in a lightweight structure that holds the avionics module. The navigation sensor suite for Mars has an IMU, a nadir-pointed camera, and a time-of-flight optical rangefinder as an altimeter. At the time it was proposed, this whole
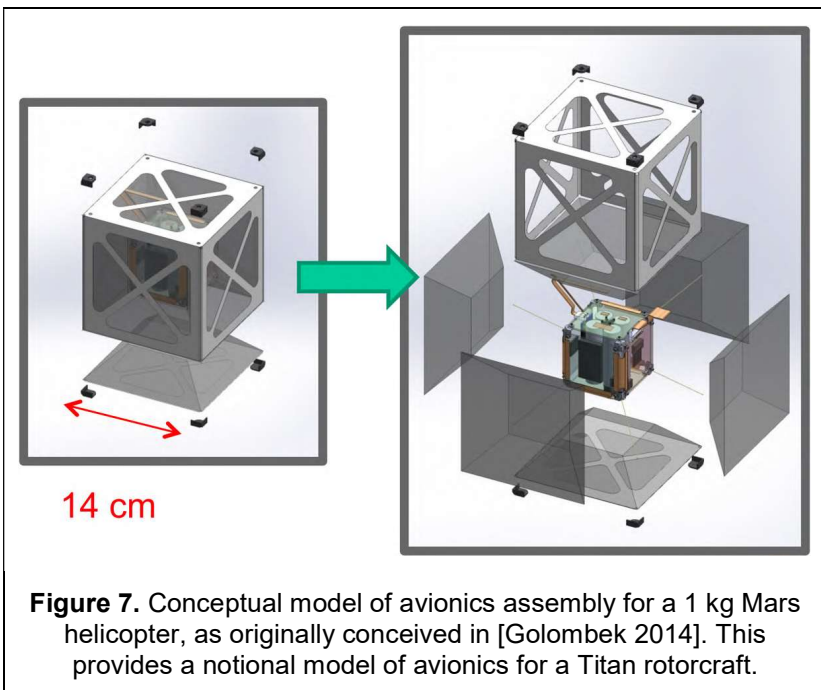
assembly was designed as a 14 cm cube with a mass of about 500 grams (Figure 7); ongoing work may modify the aspect ratio and change the mass.

Preliminary study of thermal issues for Titan suggests that slightly thicker aerogel insulation could keep the avionics assembly at an acceptable temperature for battery health during sorties. Thermal management was not studied for intervals while docked to the mothership.

NASA's High Performance Space Computing project may provide a rad-hard alternative to the COTS processor architecture of the Mars helicopter by the time a Titan rotorcraft could be proposed for a mission.

4.3   Sampling System

This study originally conceived sampling of liquid and solid material by the rotorcraft as a key aspect of the concept. The initial concept for a sampling device sought to minimize the number of actuators required by designing one integrated device that could acquire both liquid and solid samples (Figure 8). This would require bringing the device within a few centimeters of the surface to acquire samples. When the decision was made to focus on autonomy in the context of a lander-based mission and to reduce emphasis on hardware aspects of the study, work on sampling concepts was halted at this point.
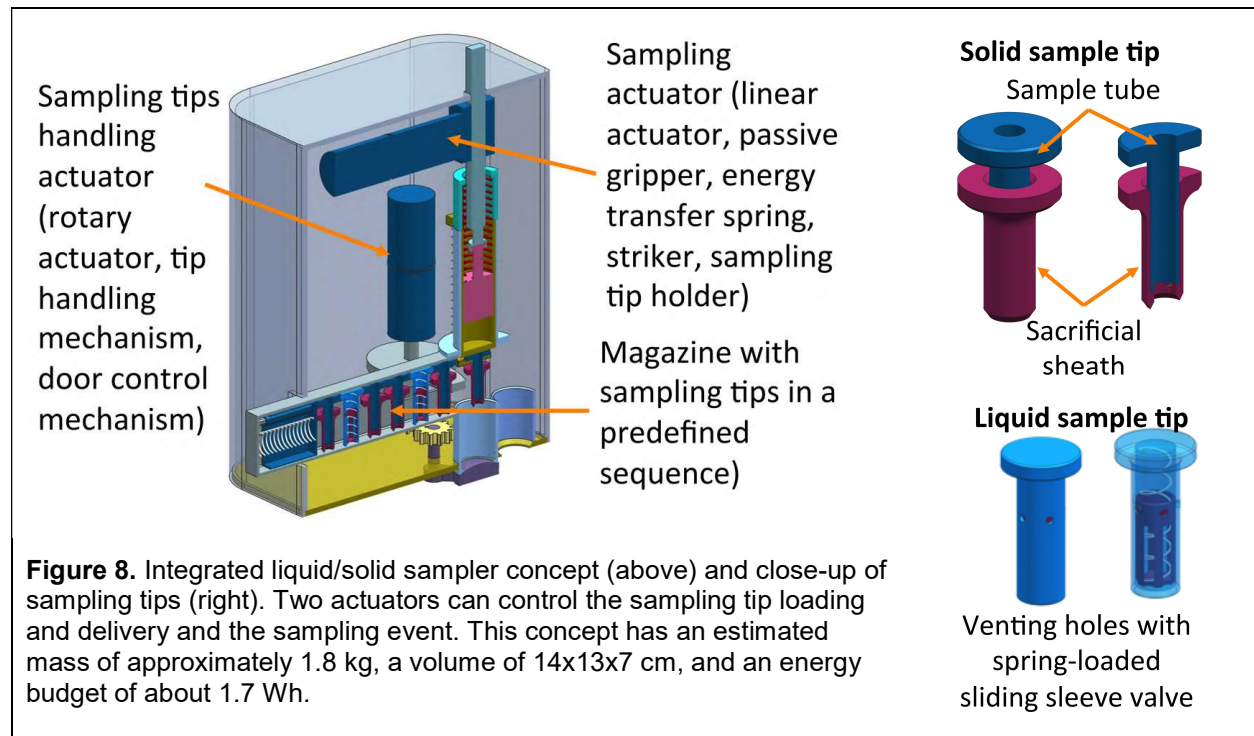

Figure 8. Integrated liquid/solid sampler concept (above) and close-up of sampling tips (right). Two actuators can control the sampling tip loading and delivery and the sampling event. This concept has an estimated mass of approximately 1.8 kg, a volume of 14x13x7 cm, and an energy budget of about 1.7 Wh.

5   Mothership Position Estimation and Precision Landing

As discussed above, if TAD mission planning is done with remote sensing images from an orbiter, knowledge of the position of the mothership relative to the reference frame of the remote sensing data is essential. In principle, this can be done by registering descent images (for a lander) or downward-looking images from a balloon with the remote sensing images. An analogous capability has been under development for many years for Mars precision landing [Mourikis 2009], where the capability to automatically register descent images to remote sensing

images onboard in real-time is now at TRL 6 [Johnson 2015] and is part of the baseline design for the Mars 2020 rover mission. For TAD with a long-lived lander, such registration could be done on Earth with downlinked images; for a short-lived lander or a balloon, it probably would have to be done onboard, but could take much more computing time than for Mars precision landing. The registration problem is harder in several ways for Titan than for Mars, because available remote sensing imagery has far lower effective resolution than for Mars (300 m/pixel for Cassini radar and 1 km or more for Cassini VIMS and ISS images, versus 30 cm/pixel for Mars) and may or may not be in the same spectral band as onboard imagery. Nevertheless, proof of concept demonstrations have been done of automatic registration of descent imagery from the Huygens probe to remote sensing imagery from the Cassini orbiter, using multi-modal image similarity metrics like mutual information, motivated by applications to Titan balloon navigation [Ansar 2009]. Such techniques could be further developed for the scenarios here.

Another important question is how the TAD flight radius model of Section 3 compares to the potential distance from the mothership to terrain of high scientific interest. Of course, this depends on the mothership delivery error, the type of terrain the mothership is deployed in, and the overall scientific objectives of the mission. Current interest in exploration of Ocean Worlds is driving more in-depth consideration of these issues, including methods that might reduce delivery error compared to that for the Huygens probe and prior mission studies. Reduced delivery error could put scientifically important terrain more easily within reach of TAD. We briefly assessed some of these issues at the close of this study.

Landing error ellipses were studied in [TLP 2010] and [Lorenz 2015]. As shown in [Lorenz 2015], the dominant factor in delivery error for Huygens probe-like, unguided entry, descent, and landing (EDL) architectures is the effect of high winds during long (~ 2.5 hour) parachute descents. According to global circulation models, Titan winds are predominantly zonal (east-west), with small meridional (north-south) components, so error ellipses tend to be highly eccentric, with the long axis oriented east-west. Zonal winds depend strongly on season and latitude; meridional winds do not vary substantially. The narrow ellipse axis is dominated by error at atmospheric entry. Taking these factors together, 99% lander delivery error ellipses for Huygens-like EDL architectures considered vary approximately from 200 to 500 km in the east-west direction and 50 to 100 km in the north-south direction. From our sizing model, it is plausible that a rotorcraft of 12 kg or more could fly most of the way across the north-south axis of such ellipses, or about halfway across the east-west axis of ellipses at the small end of the range. Landing in or near a lake or sea is currently of considerable scientific interest, and near-shore ground may host a revealing array of organic molecules. If such error ellipses were positioned with the long axis along a shoreline, the TAD flight radius might be enough to reach the shore with current EDL architectures, no matter where the lander came down in the ellipse, thereby providing options to survey or sample solids and fluids.

The key to reducing delivery error for Titan is to reduce wind drift during the parachute descent phase. This might be achieved by reducing the time spent in the parachute phase, for example by opening the parachute at a lower altitude, and/or by adding some form of control authority. We briefly examined the possibility of using a guided parafoil for control authority, which would need real-time position and velocity knowledge for parafoil guidance. It appears to be possible to provide such knowledge with techniques similar to what we discussed above for mothership position estimation; that is, onboard, real-time registration of descent images to remote sensing images, together with high-quality inertial sensing. Images from the DISR descent camera, which had a broadband spectral response from about 600 to 1000 nm [Tomasko

2002], saw the surface of Titan clearly from an altitude of about 40 km [Lebreton 2005]. Guided parafoils for terrestrial cargo delivery applications achieve a glide ratio of about 3:1 and a delivery error of about 100 m with GPS-based navigation [Airborne 2016]. If a similar glide ratio was obtained on Titan starting at 40 km altitude, the wind-relative divert range could be on the order of 100 km. This would benefit lander and balloon missions with or without an aerial daughtercraft, and could put interesting terrain within much shorter TAD flight radius from the mothership. The low resolution of Titan remote sensing images implies that registering descent and remote sensing images probably would not be possible for descent images acquired below an altitude of several kilometers; below that point, VSLAM algorithms, much like those useful for TAD navigation, would be needed to maintain position and velocity knowledge for parafoil guidance. Thus, autonomous navigation capabilities needed for a Titan guided parafoil would have a great deal in common with that needed for TAD, so technology development for either one could benefit both.

6   Summary and Conclusions

Titan is the richest laboratory in the solar system for studying prebiotic chemistry, with potential to inform us about how life originated on Earth. *In situ* mobility has potential to revolutionize Titan exploration similarly to the way rovers have revolutionized Mars exploration. Technology for small autonomous rotorcraft has matured dramatically in the last decade and a 1 kg autonomous rotorcraft is receiving intensive study for a potential technology demonstration on Mars. This suggests that rotorcraft may have great potential for exploration of Titan, which has a much denser atmosphere and much weaker gravity than Earth and Mars, making it very favorable for aerial mobility from an aerodynamics perspective.

Accordingly, we considered concepts of operation (CONOPS) for Titan aerial daughtercraft (TAD) in scenarios involving deployment from a lander or a balloon, developed sizing models to estimate TAD flight radius and payload mass fraction as a function of total TAD system mass, and presented a conceptual design of key elements of an autonomous rotorcraft, with main focus on assessing the maturity of needed autonomous navigation capabilities and showing the existence of a plausibly feasible avionics architecture. We also showed a preliminary concept for an integrated sampling device for acquiring solid and liquid samples and sketched a potential approach to estimate the position of the mothership autonomously onboard, so that specific TAD destinations could be designated via Earth-based mission planning with prior remote sensing imagery. Finally, we compared our model of TAD flight radius to the size of current Titan delivery error ellipses and discussed the possibility of reducing the sizes of error ellipses by using a guided parafoil as part of the entry, descent, and landing architecture.

Key conclusions are as follows:

- The CONOPS discussions illustrated potential scientific benefit, including the ability to visit lakes and solid ground with a signal landing, and underscored the need for a high degree of autonomy and the potential benefit of the ability to recharge TAD at the mothership to enable multiple sorties, especially for balloon-based scenarios.
- The sizing models showed that, even with very conservative model assumptions, (1) a flight radius of several 10s of kilometers from a lander should be possible with a rotorcraft total system mass on the order of 10 kg, including about 2 kg of payload, and (2) the same rotorcraft, with somewhat different autonomous navigation capabilities,

should be able to descend to the surface from a balloon at 10 km altitude and return to the balloon.

- In research for terrestrial applications, multicopters with a mass on the order of 1 kg are now able to perform visual-inertial navigation with onboard sensors and smartphone-class processors, at a maturity level comparable to Technology Readiness Level (TRL) of about 5 -- that is, integrated system demonstrations in relevant environments, but with somewhat limited performance evaluation. Autonomous landing hazard detection and avoidance is at a similar level of maturity.

- Onboard estimation of the position of the mothership to an accuracy on the order of one kilometer should be possible by automatic registration of descent imagery or downward-pointed balloon imagery with prior remote sensing imagery.

- With current Titan EDL architectures, 99% confidence landing error ellipses have been estimated in the literature as varying from about 200 to 500 km in the east-west direction and about 50 to 100 km in the north-south direction. A rotorcraft with total mass on the order of 10 kg or larger probably could fly most of the way across the north-south axis of such ellipses. If the ellipse was positioned to overlap a shoreline, the rotorcraft may be able to reach the shoreline from anywhere in the ellipse.

- Guided parafoils with a 3:1 glide ratio are used routinely for terrestrial cargo delivery. It is conceivable that similar technology could be used to reduce delivery error on Titan by up to about 100 km, using visual-inertial navigation sensors and algorithms that are very closely related to those needed for navigating a Titan rotorcraft.

7 References

[Aboudan 2010] A. Aboudan, et al, "Aerobot autonomous navigation and mapping for planetary exploration", *International Planetary Probes Workshop*, Barcelona, Spain, June 2010

[Airborne 2016] http://www.airborne-sys.com/pages/view/guided-precision-aerial-delivery-systems

[Alvarez 2014] H. Alvarez, L. M. Paz, J. Sturm, and D. Cremers, "Collision avoidance for quadrotors with a monocular camera", *International Symposium of Experimental Robotics*, 2014

[Ansar 2009] A. Ansar, L. Matthies, "Multi-modal image registration for localization in Titan's atmosphere", *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009

[Badescu 2009] M. Badescu, S. Sherritt, J. Jones, J. Hall, "Self-contained harpoon and sample handling device for a remote platform", *Proc. SPIE Vol. 7292: Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, 2009

[Barnes 2012] J.W. Barnes, et al., "AVIATR:Aerial Vehicle for In-situ and Airborne Titan Reconnaissance", *Experimental Astronomy*, 33, 55-127. doi: 10.1007/s10686-011-9275-9

[Brockers 2014] R. Brockers, M. Humenberger, Y. Kuwata, L. Matthies, S. Weiss, "Computer vision for micro air vehicles", in *Advances in Embedded Computer Vision*, pp. 55-127, Springer, 2014

[Decadal] *Vision and Voyages for Planetary Science in the Decade 2013-2022*, National Academies Press, 2011

[DISR Archive 2013] Data Archive Users' Guide for the Descent Imager and Spectral Radiometer (DISR) on the Huygens Probe, 2013

[Doyle 2014] R. Doyle, et al., "High Performance Spaceflight Computing (HPSC)", poster at *OPAG Technology Forum,* 2014

[Droeschel 2015] D. Droeschel, et al., "Multi-layered mapping and navigation for autonomous micro aerial vehicles", *Journal of Field Robotics*, 33(4), June 2016

[Forster 2015] C. Forster, L. Carlone, F. Dellaert, D. Scaramuzza, "IMU preintegration on manifold for efficient visual-inertial maximum-a-priori estimation", *Robotics Science and Systems Symposium,* 2015

[Forster 2015a] C. Forster, M. Faessler, F. Fontana, M. Werlberger, D. Scaramuzza, "Continuous on-board monocular-vision-based elevation mapping applied to autonomous landing of micro air vehicles", *IEEE International Conference on Robotics and Automation,* 2015

[Gasbarre 2005] J. F. Gasbarre, H. S. Wright, "A design comparison of atmospheric flight vehicles for the exploration of Titan", *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, 2005

[Golombek 2014] *Mars Helicopter Scout*, proposal submitted to NASA AO NNH13ZDA0180, 2014

[Helmick 2009] D. Helmick, A. Angelova, and L. Matthies, "Terrain adaptive navigation for planetary rovers", Journal of Field Robotics, 26(4), April 2009

[Johnson 2005] A. E. Johnson, J. Montgomery, L. Matthies, "Vision guided landing of an autonomous helicopter in hazardous terrain", *IEEE International Conference on Robotics and Automation*, 2005

[Johnson 2015] A. Johnson, Y. Cheng, J. Montgomery, N. Trawny, B. Tweddle, "Real-time relative navigation test results from a relevant environment for Mars landing", *AIAA Guidance, Navigation, and Control Conference,* 2015

[Karkoschka 2007] E. Karkoschka, M. G. Tomasko, L. R. Doose, C. See, E. A. McFarlane, S. E. Schroder, and B. Rizk, "DISR imaging and the geometry of the descent of the Huygens probe within Titan's atmosphere", *Planetary and Space Science*, 55:1896-1935.

[Lafleur 2006] J. M. Lafleur, "Derivation and application of a method for first-order estimation of planetary aerial vehicle power requirements", *4th Interplanetary Probe Workshop*, Pasadena, CA, 2006

[Lebreton 2005] J.-P. Lebreton, et al., "An overview of the descent and landing of the Huygens probe on Titan", *Nature,* 438, 758-764, 8 December 2005

[Leutenegger 2015] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, P. Furgale, "Keyframe-based visual-inertial odometry using nonlinear optimization", *International Journal of Robotics Research*, 34(3): 314-334, 2015

[Li 2013] M. Li, B.H. Kim, A.I. Mourikis, "Real-time motion tracking on a cellphone using inertial sensing and a rolling-shutter camera", *IEEE International Conference on Robotics and Automation*, 2013

[Lopes 2010] R.M.C. Lopes, et al., "Distribution and interplay of geologic processes on titan from Casssini radar data", *Icarus*, 205, 540–558. doi: 10.1016/j.icarus.2009.08.010.

[Lorenz 2000] R. D. Lorenz, "Optimizing science return from Titan aerial explorers", *IEEE Aerospace Conference*, 2000

[Lorenz 2001] R. D. Lorenz, "Flight power scaling of airplanes, airships, and helicopters: application to planetary exploration", *Journal of Aircraft*, 38(2): 208-214, 2001

[Lorenz 2008] R.D. Lorenz, et al., "Titan's inventory of organic surface materials", *Geophysical Research Letters,* 35, L02206. doi: 10.1029/2007GL032118

[Lorenz 2008a] R.D. Lorenz, "Titan Bumblebee: a 1 kg lander-launched UAV concept", *Journal of the British Interplanetary Society,* 61, 118-124, 2008

[Lorenz 2015] R. D. Lorenz and C. E. Newman, "Twilight on Ligeia: implications of communications geometry and seasonal winds for exploring Titan's seas 2020-2040," *Advances in Space Research*, 56, 190-204, 2015

[Matthies 2003] L. Matthies, P. Bellutta, and M. McHenry, "Detecting water hazards for autonomous off-road navigation", *Proc. SPIE Vol. 5083: Unmanned Ground Vehicles V*, 2003

[Mourikis 2009] A. I. Mourikis, N. Trawny, S. I. Roumeliotis, A. E. Johnson, A. Ansar, L. Matthies, "Vision-aided inertial navigation for spacecraft entry, descent, and landing", *IEEE Transactions on Robotics*, 25(2): 264-279, 2009

[Nuske 2015] S. Nuske, et al., "Autonomous exploration and motion planning for an unmanned aerial vehicle navigating rivers", Journal of Field Robotics, 32(8), December 2015

[Powers 2014] C. Powers, D. Mellinger, V. Kumar, "Quadrotor kinematics and dynamics", *Handbook of Unmanned Aerial Vehicles*, pp. 307-328, Springer, 2014

[Prakash 2006] R. Prakash, et al., "Design of a long endurance Titan VTOL vehicle", *IEEE Aerospace Conference*, 2006

[Rankin 2011] A. Rankin., L. Matthies, and P. Bellutta, "Daytime water detection based on sky reflections", *IEEE International Conference on Robotics and Automation*, 2011

[Sarabandi 2011] K. Sarabandi, M. Vahidpour, M. Moallem, J. East, "Compact beam scanning 240 GHz radar for navigation and collision avoidance", *Proc. SPIE Vol 8031: Micro- and Nanotechnology Sensors, Systems, and Applications III*, 2011

[TE 2007] *Titan Explorer Flagship Mission Study*, 2007

[TLP 2010] *JPL Team X Titan Lake Probe Study Final Report*, 2010

[Tokumaru 2014] P. Tokumaru, J. Balaram, J. A. Cutts, L. H. Matthies, "Fixed and rotary wing flight of small air vehicles on Mars, Venus, and Titan", *11th International Planetary Probe Workshop,* 2014

[Tomasko 2002] M. G. Tomasko, et al., "The Descent Imager/Spectral Radiometer (DISR) experiment on the Huygens entry probe of Titan", *Space Science Reviews,* 104, 469-551, 2002

[TSSM 2009] *Titan Saturn System Mission: NASA/ESA Joint Summary Report*, 2009

[Turtle 2011] E.P. Turtle, et al., "Rapid and Extensive Surface Changes Near Titan's Equator: Evidence of April Showers", *Science,* 331, 1414-1417. doi: 10.1126/science.1201063

[Weiss 2013] S. Weiss et al., "Monocular vision for long-term micro aerial vehicle state estimation: a compendium", *Journal of Field Robotics*, 30(5), 803-831, 2013

[Whalley 2013] M. S. Whalley, et al., "Flight test results for autonomous obstacle field navigation and landing site selection on the RASCAL JUH-60A", *American Helicopter Society 69th Annual Forum*, 2013.

[Young 2001] L. A. Young, E. W. Aiken, "Vertical lift planetary aerial vehicles: three planetary bodies and four conceptual design cases", *27th Rotorcraft Design Forum*, 2001

[Zimmerman 2013] W. F. Zimmerman, et al., "Development of technologies for cryogenic in-situ chemical analysis on Titan", *Annual Report for JPL Research and Technology Development Program*, 2013