High Level Risk Analysis and Decision Making Regarding the Prediction of Thermal Lift Locations for an Autonomous Mars Glider

Jen Jen Chung, Salah Sukkarieh

Australian Centre for Field Robotics
School of Aerospace, Mechanical and Mechatronic Engineering
University of Sydney
NSW Australia

Summary: The ARES (Aerial Regional-scale Environmental Survey of Mars) glider project is a current NASA program that aims to send a fully autonomous glider system to survey the Martian landscape from a gliding altitude of 1.5km. The design of the glider propulsion system is complicated by issues of low atmospheric density and the lack of appreciable amounts of O_2 in the Martian environment, thus it is desirable to make use of the glider's soaring capabilities to maximise gliding endurance by exploiting thermal energy sources in the region whilst also considering the external mission goal of exploring the environment. This paper discusses the application of a Bayesian Forecast-Decision System (BFS) to the prediction of thermal motion, the detection and management of system uncertainties, and high level risk assessment of available flight paths through the environment. The results of this paper show that the forecast decision system requires very little a priori information to effectively gauge the system uncertainties and manage platform energy to efficiently survey the allocated search region. Furthermore, it is also shown that the system can run in real-time on an autonomous sensing and reacting system operating in a dynamic environment.

Keywords: Autonomous glider, Bayesian decision making, risk assessment.

Introduction

ARES Scout Mission

The primary science objectives of the ARES Mars Scout Mission are to sample the Martian atmosphere to characterise the structure and dynamics of the atmospheric boundary layer over regional scales, measure water-equivalent hydrogen abundance and ice burial depth, and investigate the crustal magnetism on Mars [1-3]. The mission design requires the ARES glider to explore up to 610km of the Southern Highlands region from a gliding altitude of 1.5km [4]. However, low atmospheric density and the lack of appreciable amounts of O_2 in the Martian environment pose a number of issues to the design of the glider propulsion system; thus it is desirable to make use of the glider's soaring capabilities. To maintain the level of autonomy necessary to successfully carry out the mission, the glider system must be capable of seeking out energy sources within the environment to extend gliding endurance whilst maintaining the nominal pre-planned trajectory to ensure that sufficient data are collected.

The hardware system aboard the ARES glider is shown in Fig. 1. On board sensors are capable of sampling the local atmosphere to determine air temperature and pressure to characterise nearby thermals. Given this information, it is proposed that a Bayesian forecast-decision system can quantify the uncertainties in the sensor measurements and forecast models to plan a path through the environment that accounts for the risks associated with each waypoint decision.

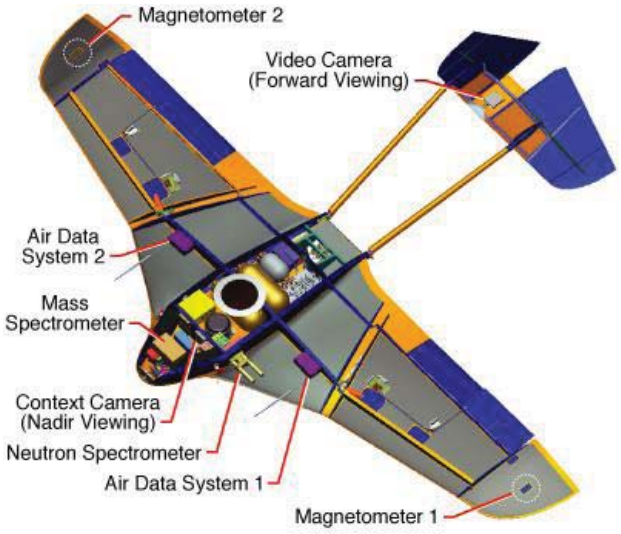

Fig. 1: The NASA ARES platform is equipped with sensors to measure the crustal magnetisation of Mars as well as sample the Martian atmosphere to determine sources of biogenic, volcanic, and chemically active gases [4].

The Bayesian Forecast-Decision System for Time Series Data

The Bayesian forecast-decision system (BFS) was first presented by Krzysztofowicz for the purpose of identifying and quantifying sensor observation and forecast model uncertainties and was applied to meteorology forecasting and sensor fusion problems [5]. The BFS is able to predict thermal motion according to a user defined forecast model, and applies the Bayesian Processor of Forecast (BPF) to detect and manage the contribution of model and input uncertainties associated with glider sensor data and forecast data to produce a posterior distribution over future thermal positions. In general, the glider decision system receives and processes information regarding the state of nature via the timing schematic shown in Fig. 2.

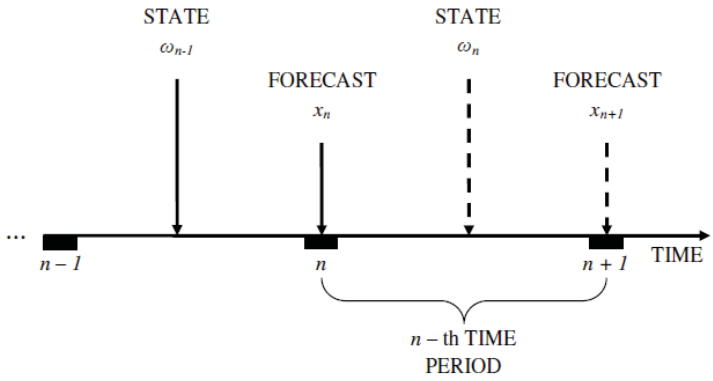

Fig. 2: Schematic timing of information as processed by the BFS. At the beginning of the n-th time period, the system has observations of the previous states and knowledge of all the forecasts including the current forecast, x_n, upon which a manoeuvre decision is made. The state of nature during the n-th time period, ω_n, is observed by the system at the end of this time period and is used to formulate the next forecast x_{n+1}.

The role of the decision model of the BFS is to select the next waypoint by assessing the risk associated with traversing available paths through the environment, where risk is computed according to goal point reachability and the map uncertainty as quantified by the BPF.

Algorithm Design

Forecast Model

The purpose of the initial stage of the BFS is to gather the available raw data and generate a forecast of the state of nature for the next period of time. In this case, a simplified forecasting model is used, which assumes that over relatively short time intervals (within the space of one minute), thermals travel at a constant velocity. Examination of the previous thermal locations can provide a measure and spread of the thermal velocities, however the number of prior states to use should be large enough to account for outliers but also small enough pick up changes in thermal velocities over time. Additionally, an initial estimation of the thermal velocities is required as no prior states have been recorded. The spread of the estimated velocities is later used by the BPF to quantify the uncertainties stemming from this modelling method.

Bayesian Processor of Forecasts

The BPF judges the performance of the forecasting model by comparing previous forecasts against known thermal locations to quantify the current forecast uncertainty. For time series data, the BPF updates the prior distribution of the uncertain state with the new forecast x_n, and employs Bayes' theorem to produce a posterior distribution η_n over the possible future thermal positions ω_n. Namely:

$$\eta_n(\omega_n \mid \omega_1,...,\omega_{n-1}, x_1,...,x_n) = \frac{f_n(x_1,...,x_n \mid \omega_1,...,\omega_n) g_n(\omega_1,...,\omega_n)}{\xi_n(\omega_1,...,\omega_{n-1}, x_1,...,x_n)} \tag{1}$$

Where ξ_n is the predictive density of the forecast:

$$\xi_n(\omega_w,...,\omega_{n-1}, x_1,...,x_n) = \int f_n(x_1,...,x_n \mid \omega_w,...,\omega_n) g_n(\omega_w,...,\omega_n) d\omega_n \tag{2}$$

Densities f_n and g_n are the known joint and prior densities of the relevant variables, respectively. The prior density describes the natural uncertainty regarding the process (in this case, thermal motion) that exists before receiving the forecast. The joint density for fixed forecasts represents the likelihood of the states, that is, the function calculates the forecast uncertainty by judging how well the forecasting model has predicted the state in the past.

Krzysztofowicz [5] suggests a number of assumptions that can be made about the structure of the likelihood function and the relationship of the states to simplify the posterior distribution calculation.

Assumption 1: Conditional upon $\omega_1,...,\omega_n$, forecasts $x_1,...,x_n$ are stochastically independent of each other.

$$f_n(x_1,...,x_n \mid \omega_1,...,\omega_n) = \prod_{i=1}^{n} f_i(x_i \mid \omega_1,...,\omega_n) \tag{3}$$

Assumption 2: Conditional upon ω_i, forecast x_i is stochastically independent of ω_j for all $j \neq i$. Thus, for any $i \leq n$,

Assumption 3: The time series of thermal locations, states $\omega_1,...,\omega_n$, form a first-order Markov process.

$$g_n(\omega_1,...,\omega_n) = \prod_{i=1}^{n} g_i(\omega_i \mid \omega_{i-1}) \tag{5}$$

$g_1(\omega_1 \mid \omega_0) = g_1(\omega_1)$ is the initial state density, and $g_2,...,g_n$ are one-stage transition densities.

Given Eqn 3, Eqn 4 and Eqn 5, Eqn 1 takes the form:

$$\begin{aligned}\eta_n(\omega_n \mid \omega_1,...,\omega_{n-1},x_1,...,x_n) &= \frac{f_n(x_n \mid \omega_n)g_n(\omega_n \mid \omega_{n-1})}{\int f_n(x_n \mid \omega_n)g_n(\omega_n \mid \omega_{n-1})d\omega_n} \\ &= \eta_n(\omega_n \mid \omega_{n-1},x_n)\end{aligned} \tag{6}$$

The full derivation can be found in [6]. According to Eqn 6, given the thermal positions during the previous time period and the current forecast for the thermal positions at the next time period, the posterior state transition density $\eta_n(\cdot \mid \omega_{n-1},x_n)$ over the possible thermal locations in the next time period is a combination of the natural uncertainty about the state and the uncertainty in the forecast x_n of ω_n, the prior state transition density $g_n(\cdot \mid \omega_{n-1})$ and the likelihood function $f_n(x_n \mid \cdot)$, respectively.

The natural uncertainty can be estimated based on on-site historical records, regional information and any scientific knowledge or experience relevant to ω_n. For the purpose of locating thermals, this could include the tendency for thermals to diminish within certain periods of time (particularly if they are likely to disappear altogether in a space of time shorter than the data sampling time), and other known thermal movement patterns and thermal interaction behaviour.

The uncertainty in the forecast x_n of ω_n is estimated from historical records; by observing J joint observations of forecasted thermal locations and actual thermal locations $\{(x_j,\omega_j) : j = 1,...,J\}$, it is possible to determine the performance of the forecasting model. At the start of operation, no adequate historical record of the forecast and actual states will exist. Krzysztofowicz suggests that under this circumstance, a simulated record can be used; however, care must be taken to ensure that the simulated model accounts for all uncertainties and assumptions of the actual forecasting model. As the glider begins to receive sensor observations, the significance of the historical records on the state estimate will decrease as a sliding window is used to regulate the number of observation entries used in the estimation.

Given the estimates of f_n and g_n, the predictive density of the forecast can be calculated as:

$$\xi_n(x_n \mid \omega_{n-1}) = \int f_n(x_n \mid \omega_n)g_n(\omega_n \mid \omega_{n-1})d\omega_n \tag{7}$$

With $\xi_1(\omega_1 \mid \omega_0) = \xi_1(\omega_1)$. Combining Eqn 6 with Eqn 7 gives:

$$\eta_n(\omega_n \mid \omega_{n-1},x_n) = \frac{f_n(x_n \mid \omega_n)g_n(\omega_n \mid \omega_{n-1})}{\xi_n(x_n \mid \omega_{n-1})} \tag{8}$$

Fig. 3 shows the layout of the BPF algorithm.

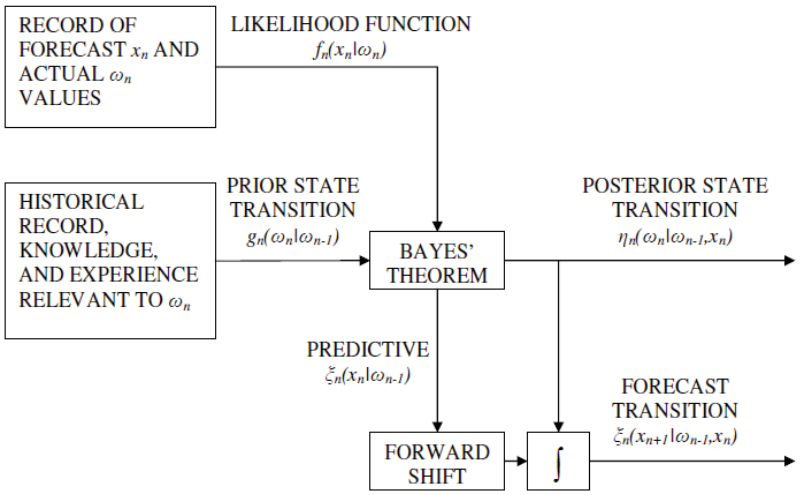

Fig. 3: Simplified block diagram showing the steps involved in the BPF algorithm for a first-order Markov Process [5].

With respect to the overall risk analysis system, the variables of interest from the BPF are the posterior mean and variance of η_n. Since the likelihood and prior densities are both assumed to be normal, then it can be seen that g_n is a conjugate prior of f_n. Thus it is possible to derive a closed form representation of the posterior distribution from which the two hyperparameters of interest can be obtained.

Supposing that the relationship between the forecast x_n and the actual state ω_n can be modelled using a linear forecast equation:

$$x_n = a_n \omega_n + b_n + \theta_n \tag{9}$$

where a_n and b_n are fixed parameters and θ_n represents the residual noise with a normal density distribution $N(0,\zeta_n^2)$. In the case of perfect forecasts, $a_n = 1$ and $b_n = \zeta_n = 0$. Therefore it is expected that the calculated values of these parameters converge to this ideal condition. Computation of the parameters necessitates that b_n and θ_n are measured as one variable with a distribution $N(b_n,\zeta_n^2)$ from which the required parameter values can be extracted.

From Eqn 9, the likelihood function can be represented as a normal distribution:

$$f_n(x_n \mid \omega_n) = N\left(a_n \omega_n + b_n, \zeta_n^2\right) \tag{10}$$

Or in terms of the forecast:

$$\begin{aligned} f_n(x_n \mid \omega_n) &= N\left(\mu_f, \sigma_f^2\right) \\ &= N\left(\frac{x_n}{a_n} - \frac{b_n}{a_n}, \frac{\zeta_n^2}{a_n^2}\right) \end{aligned} \tag{11}$$

In a similar fashion, since the states form a first-order Markov process, then:

$$\omega_n = c_n \omega_{n-1} + d_n + \nu_n \tag{12}$$

where c_n and d_n are fixed parameters and ν_n represents the residual noise with a normal density distribution $N(0,\tau_n^2)$. Once again, if perfect thermal data are received from the sensors and the thermals behaved exactly according to the linear motion assumption made in the forecast model, then $c_n = 1$, $d_n = \Delta t \times v$ and $\tau_n = 0$, where $\Delta t \times v$ is the time step multiplied by the thermal velocity. Further, computation of d_n and ν_n must also combine the two parameters into a single distribution $N(d_n,\tau_n^2)$ containing both values of interest.

Eqn 12 gives:

$$\begin{aligned} g_n(\omega_n \mid \omega_{n-1}) &= N(\mu_g, \sigma_g^2) \\ &= N(c_n \omega_{n-1} + d_n, \tau_n^2) \end{aligned} \tag{13}$$

Applying Eqn 11 and Eqn 13 into Eqn 7 and Eqn 8, the mean and variance of the posterior distribution can be calculated using the posterior hyperparameter relations derived by Raiffa and Schlaifer [7]. That is, for:

$$\eta_n = N(\mu_n, \sigma_n^2) \tag{14}$$

$$\begin{aligned} \mu_n &= \frac{\left(\frac{\mu_g}{\sigma_g^2} + \frac{\mu_f}{\sigma_f^2}\right)}{\left(\frac{1}{\sigma_g^2} + \frac{1}{\sigma_f^2}\right)} \\ &= \frac{a_n \tau_n^2 x_n + c_n \zeta_n^2 \omega_{n-1} + d_n \zeta_n^2 - a_n b_n \tau_n^2}{\zeta_n^2 + a_n^2 \tau_n^2} \end{aligned} \tag{15}$$

$$\begin{aligned} \sigma_n^2 &= \left(\frac{1}{\sigma_g^2} + \frac{1}{\sigma_f^2}\right)^{-1} \\ &= \frac{\tau_n^2 \zeta_n^2}{\zeta_n^2 + a_n^2 \tau_n^2} \end{aligned} \tag{16}$$

Decision Model

The purpose of the decision model is to allocate a discrete action to any given state of nature. In the BFS, the decision model receives the probability distributions calculated in the BPF to create a map of the world including the glider's own position and where it expects thermals to be located. Given the estimated state and the respective probabilities associated with each thermal's existence, the model is able to choose a suitable target for the glider. The decision of which option represents the 'best' choice is based on the calculated expected utility resulting from each manoeuvre.

For the glider, utility is determined as a weighted combination of three factors:

- Distance of detour as a fraction of the nominal path
- Deviation of glider heading angle from nominal path
- Energy/altitude gain available

The weighting is an operator defined heuristic and should reflect the mission objectives. For the given exploration and survey missions outlined in this paper, a weighting ratio of 20:1:10 is applied, respectively. The low weighting for the glider heading angle deviation reflects the manoeuvrability of the platform while the distance factor is weighted highest to drive the glider to its goal location.

In addition to these factors, the decision model is also able to safely direct the glider when it is operating below an operator defined critical altitude. Under such circumstances, the decision model instructs the glider to apply powered flight to the nearest sensed thermal.

Glider Simulation Design

The glider simulation is set up such that for each trial, a random set of thermals are generated in the field of operation. Each thermal has an associated 3-dimensional velocity and position, which specifies the potential energy/altitude gain available to the glider (defined as the effective maximum height) and the thermal radius of effect (estimated as one-tenth of the effective maximum height). The glider parameters of interest are outlined in Table 1.

Table 1: Glider platform and sensor parameters

Parameter	Value
Nominal gliding speed	25 m/s
Descent rate	1 m/s
Maximum banking angle	55°
Sensor range	1000m
Sensor sample interval	60s
Camera FOV	20° – 30°

The sensor range of 1000m is an assumption based on the work in [8] which presents a method for state estimation of non-linear wind fields. The time step used in [8] was 1.2 hours and dealt with lower frequency dynamics of the wind field, however for our purposes, we are concerned with the high frequency dynamics and thus have chosen a 40 second horizon for the sensor range. The reader is referred to [6] for further simulation results that show the effect of varying sensor ranges and platform gliding dynamics.

The waypoints for the glider mission are predefined and loaded at the start of the simulation. Two main missions were examined in simulation:

1. Local area survey using a boustrophedon search pattern.
2. Cross-country exploration following a predominantly straight path.

These two paths best represent the ARES scout mission objectives of sampling regional atmospheric and crustal properties, and following water trails to test for hydrogen content as outlined in the introduction.

Fig. 4 gives a snapshot of the Matlab simulation. In the instance shown, the glider is adjusting its course to intercept the thermal at (1020, 2024, -3400)[1].

.

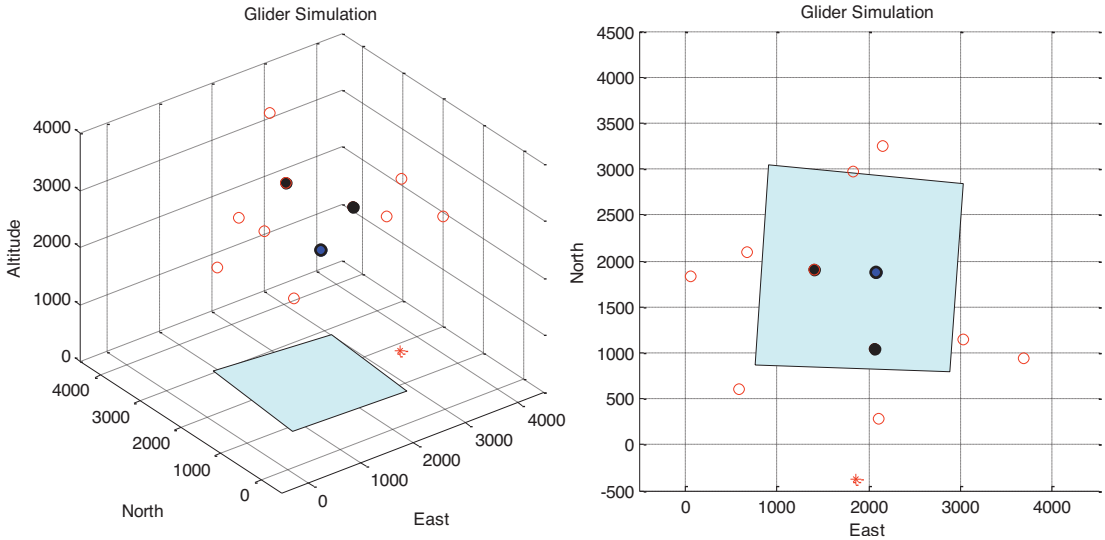

(a) (b)

Fig. 4: Matlab simulation of the glider mission; (a) 3D view, (b) aerial view. The glider is represented as the blue point mass with a camera footprint shown in light-blue; red circles represent effective thermal heights and are located at the centre of the thermal columns; black points mark assumed thermal locations; red stars represent the flight plan waypoint.

Simulation Results

Localised Survey

The flight paths taken by the glider under the localised survey mission mode are shown in Fig. 5. In both flights, the BFS decision model selects intuitive thermal waypoints to maintain altitude and navigate the surveyed region. It is evident from Fig. 5b that the system is capable of predicting the movement of the thermals as previously visited thermals travel significant distances (up to 500m) between survey loops.

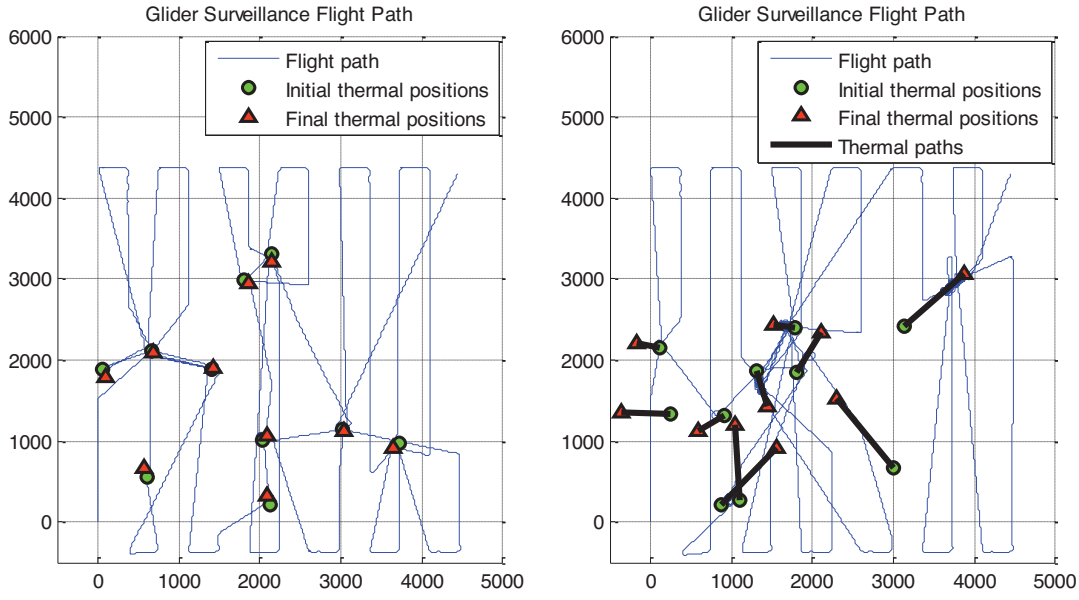

(a) (b)

Fig. 5: Flight path taken by the glider to survey a region of slow-moving thermals (a)*; flight path taken by the glider to survey a region of fast-moving thermals* (b).

The flight paths show that the basic boustrophedon search pattern is maintained while the glider exploits thermal sources for altitude gain. Fig. 6 shows the glider altitude over the flight. For both flights, the glider remains near the nominal gliding altitude of 1.5km, thus maintaining image resolution whilst correcting when the critical flight altitude of 750m was violated.

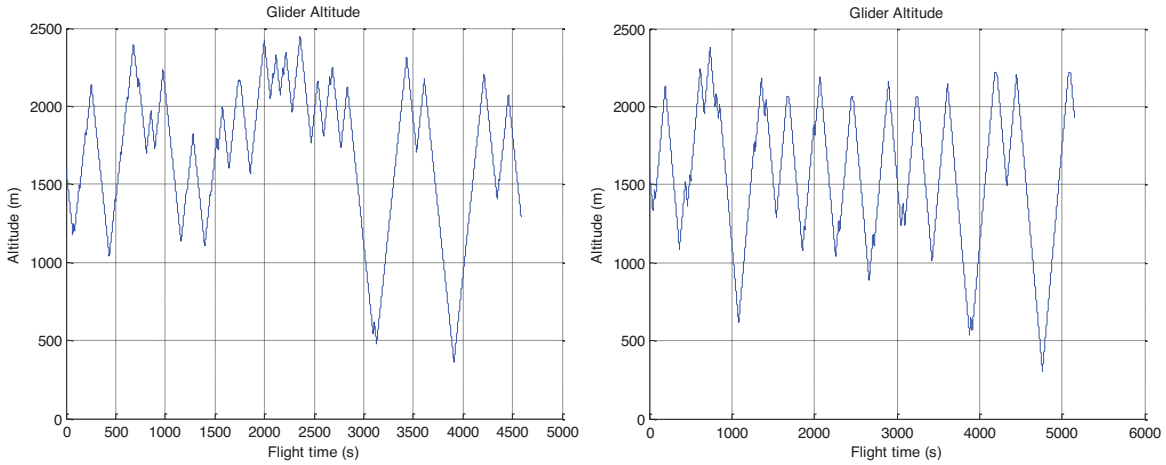

(a) (b)

Fig. 6: Glider altitude over the flight paths shown in Fig. 5a and Fig. 5b, respectively.

The 20° FOV sensor coverage over the flight path is given in Fig. 7. Fig. 7a shows that apart from a few isolated pockets, the entire 4000m×4000m survey region is captured. In Fig. 7b, the gaps in the sensor coverage are due to the placement of thermals within the field causing the glider to deviate further from the nominal search path to gain energy from nearby thermals. Despite these path deviations, over 90% of the search area was captured.

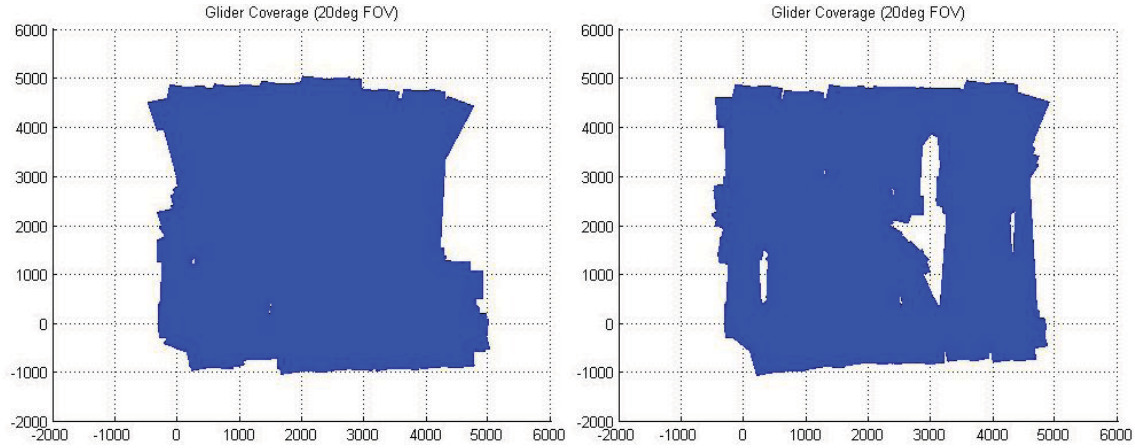

(a) (b)

Fig. 7: 20° FOV sensor footprint coverage across the flight paths shown in Fig. 5a and Fig. 5b respectively. Observations taken at a glider bank angle > 10° are discarded due to poor image resolution. Such images would be taken when the glider is turning around waypoints or circling thermals and would be subject to higher levels of noise.

Given the 30° FOV sensor, the coverage across the flight paths is shown in Fig 8. The entire surveyed region is now captured in case (a) and Fig. 8b shows a significant improvement from the 20° FOV camera in case (b). Over 99% of the surveyed region is captured with the wider lens camera. It is also noted the poorer quality sensor observations taken while the glider is circling may still be used to resolve coverage 'black spots' if necessary.

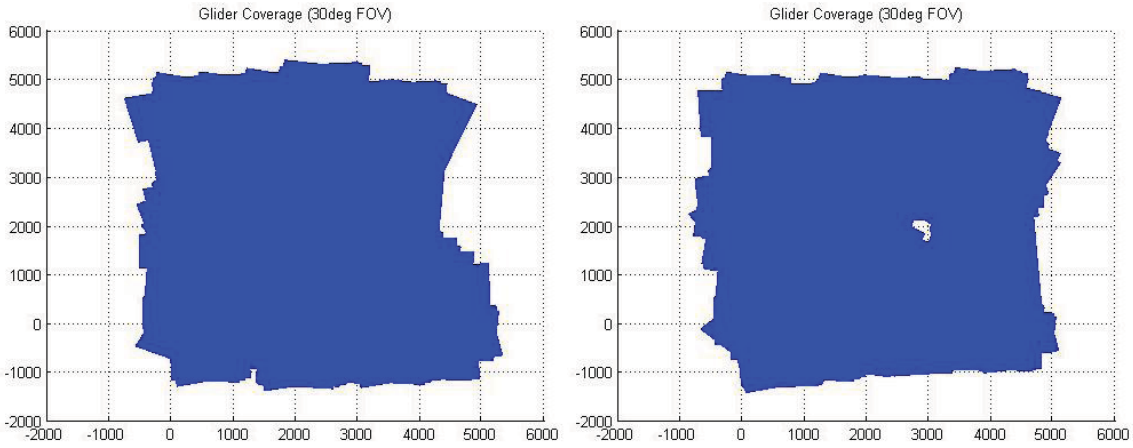

(a) (b)

Fig. 8: 30° FOV visual camera footprint coverage across the flight paths shown in Fig. 5a and Fig. 5b respectively.

Cross-Country Exploration

The performance of the system while the glider is operating under the cross-country exploration mode is shown in Fig. 9. Fig. 9a shows the glider's ability to maintain its search track whilst locating and exploiting nearby thermals to increase gliding endurance. Fig. 9b shows that the 20° FOV sensor footprint provides consistent coverage along the nominal path.

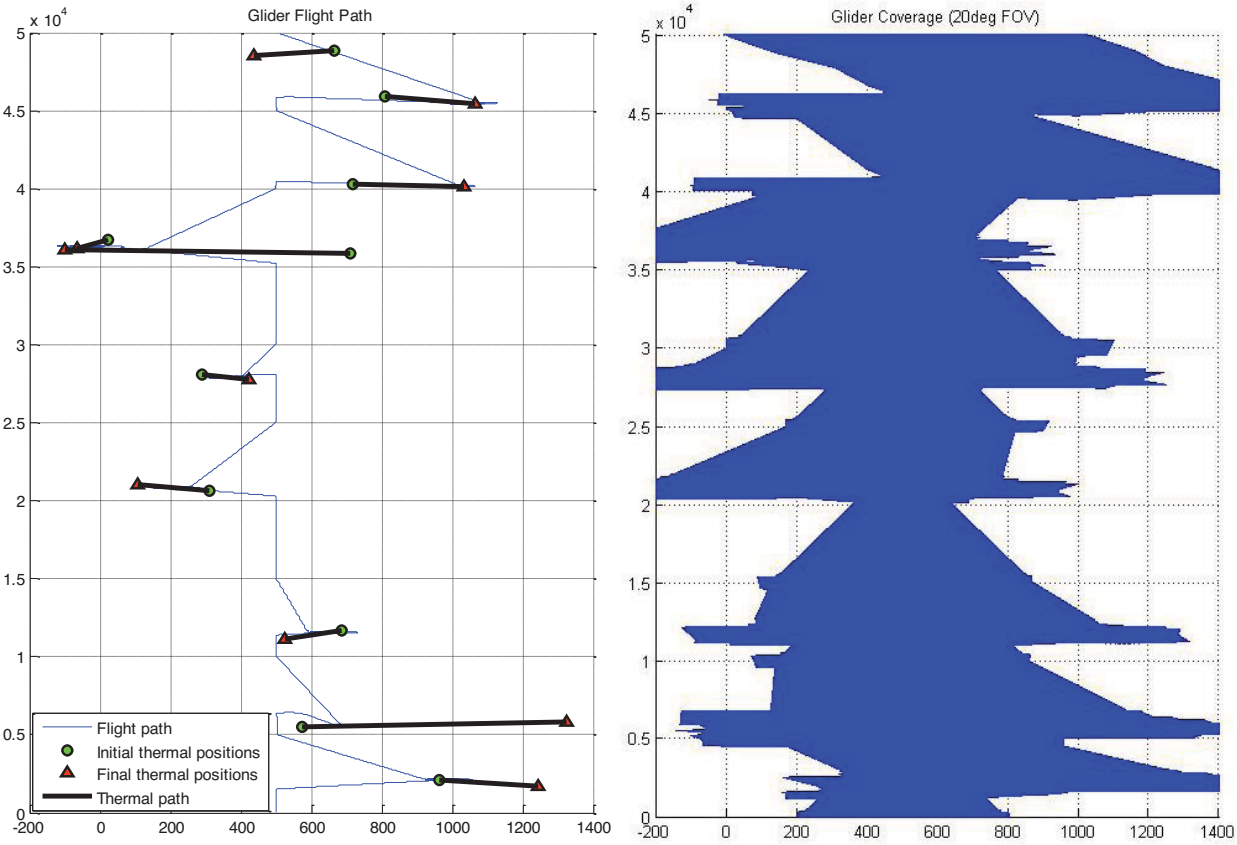

(a) (b)

Fig. 9: Glider performing autonomous exploration traversal over a 50km straight stretch.

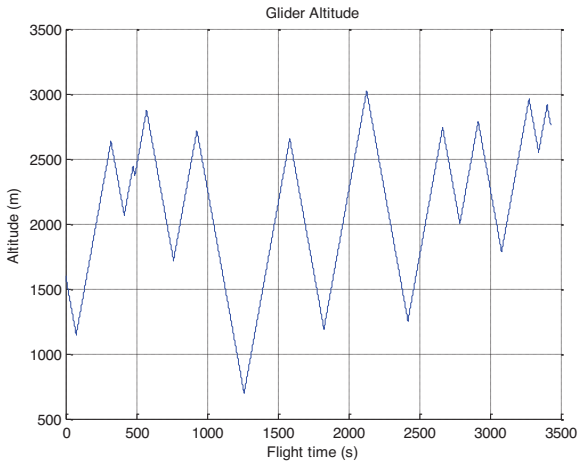

Fig. 10: Glider altitude over the cross-country flight path shown in Fig. 9a.

The glider altitude over the 50km flight is shown in Fig. 10. The maximum acceptable glider altitude was increased to 2500m for this simulation to ensure glider safety; due to this, the glider's altitude varies mainly between 1000m to 2500m over this flight plan. It is noted that over the course of the flight, the glider only drops below the critical altitude once and rarely deviates further than 500m off the allocated track.

Real-Time Computation

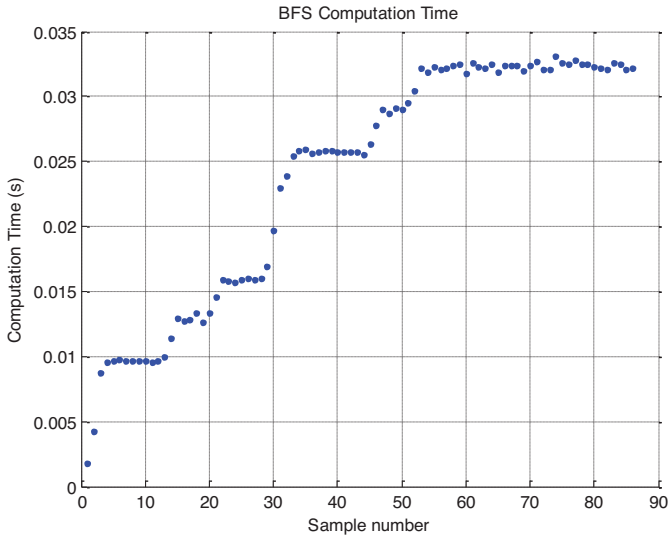

Fig. 11: BFS computation time over the flight plan shown in Fig. 5a. Simulation was run on an Intel® Core™ 2 Duo CPU E8400 @ 3.00GHz 4.00GB RAM.

Fig. 11 shows that the BFS computation time scales proportionally to the number of thermals detected by the glider. It is evident from the graph that although previous sensor observations and forecasts are retained in the BFS, the computational load does not increase over samples when no new thermals are detected. The BFS computation time of 3.25ms/thermal is compatible to the sensor sample rate (once every 60 seconds) and the dynamics of the glider. However, over extended flights, it will be desirable to remove thermals that have not been observed for a period of time to free on board memory space and reduce the computational load. These thermals may have moved out of the region of interest or have dissipated their energy altogether.

Conclusion

The application of a Bayesian forecast-decision system to an autonomous glider operating in an unknown and dynamic environment has been shown in this paper. The BFS is capable of seeking out and exploiting available thermal energy sources in the environment to maintain gliding altitude whilst considering an external mission goal such as the ARES glider missions of regional surveys and cross-country exploration.

Future work in this area may involve the use of terrain modelling from camera observations to enable another layer of probability modelling focusing on the coupling between geographical features such as mountain ranges and ridges and sources of rising air (orographic effect).

References

1. Levine, J.S., et al. 2003, "Science from a Mars Airplane: The Aerial Regional-scale Environmental Survey (ARES) of Mars", 2nd AIAA "Unmanned Unlimited" Systems, Technologies, and Operations – Aerospace, Land, and Sea Conference, Paper Number AIAA-2003-6576.

2. Guynn, M.D., et al. 2003, "Evolution of a Mars Airplane Concept for the ARES Mars Scout Mission", 2nd AIAA "Unmanned Unlimited" Systems, Technologies, and Operations – Aerospace, Land, and Sea Conference, Paper Number AIAA-2003-6578.

3. Sandford, S.P., et al. 2003, "ARES and Beyond: Autonomous Aerial Platforms Provide a Unique Measurement Capability for Earth and Planetary Science", 2nd AIAA "Unmanned Unlimited" Systems, Technologies, and Operations – Aerospace, Land, and Sea Conference, Paper Number AIAA-2003-6579.

4. Levine, J.S. 2010, *ARES Mars Scout Mission Proposal – Science*, Hampton, viewed 27 June 2010, < http://marsairplane.larc.nasa.gov/science.html>.

5. Krzysztofowicz, R. 1985, "Bayesian Models of Forecasted Time Series", *Water Resources Bulletin*, vol. 21, no. 5, pp. 805-814.

6. Chung, J.J. 2009, "High Level Risk Analysis and Decision Making for an Autonomous Glider", BE Hons thesis, University of Sydney.

7. Raiffa, H. & Schlaifer, R. 1961, Applied Statistical Decision Theory, 3rd edition, Harvard University, Boston.

8. Park, S., et al. 2009, "Learning Covariance Dynamics for Path Planning of UAV Sensors