

Predictive Control for Soaring of Unpowered Autonomous UAVs

Darren Lee * Stefano Longo ** Eric C. Kerrigan ***

* *Department of Aeronautics, Imperial College London, London SW7 2AZ, United Kingdom (e-mail: dl808@imperial.ac.uk).*

** *Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, United Kingdom (e-mail: s.longo@imperial.ac.uk; tel: +44 (0)20 7594 6118; fax: +44 (0)20 7581 4419)*

*** *Department of Electrical and Electronic Engineering and Department of Aeronautics, Imperial College London, London SW7 2AZ, United Kingdom (e-mail: e.kerrigan@imperial.ac.uk)*

Abstract: We design an autonomous soaring controller for an unpowered UAV in a nonlinear MPC framework. The UAV is controlled with the aim of extracting the maximum amount of potential/kinetic energy from the environment's updrafts. We focus on conceptual feasibility at this stage and make the realistic assumption that the UAV obtains updraft information only along the flight trajectory. The surrounding updraft distribution is then recursively estimated (online) by combining the measurements from the optimal trajectory with a heuristic search, if necessary. A variation of the standard grid search is used such that the grid spacing is altered depending on the updraft information along the UAV's flight path. Results from both standard and adaptive grid search approaches are presented. In abstract terms, this work can be viewed as finding optimal paths in uncertain vector fields.

1. INTRODUCTION

A current area of keen interest is that of endurance of Unmanned Aerial Vehicles (UAVs). Because UAVs do not have the constraints of human endurance in determining their time aloft, the limitations to a UAV's flight time are often determined by its altitude, fuel or batteries. The normal mission purpose of a UAV involves loitering over a target area of interest. Therefore any ability to improve this loiter time by optimally extracting energy from the environment would be highly beneficial. The energy could be used to gain altitude, reduce fuel payload, and even recharge batteries. Rising masses of air, which for many reasons naturally occur within the lower atmosphere, is one of the more accessible forms of energy waiting to be captured. Presently, this energy is only exploited by soaring birds and glider pilots (Akos et al. [2010], Shannon et al. [2002]).

NASA has conducted work regarding control of soaring UAVs (Allen and Lin [2007]) by assuming a model of the so-called 'thermals' (rising masses of air generated from the uneven surface heating). In that study, a set of algorithms used to detect thermals for energy gain were flight tested. Parameters such as updraft velocity, radius and position were estimated. However, no optimal control was used. A significant amount of work done regarding energy extraction from thermals has revolved around thermal centering methods (Fonseka [2007], Wharington [1998]). Often a model of the thermal is assumed (Kahveci et al. [2008], Qi and Zhao [2005]) and a maximum-likelihood picture of the thermal is estimated.

While the use of a typical thermal model could be useful in online estimation of the thermal environment the UAV encounters, it cannot be used to model other forms of atmospheric updrafts such as ridge lift or wave lift (Cutler et al. [2010], Kagabo [2010]). These forms of updrafts do not have the same

distributions or behavior of a thermal and the model is hence not portable.

This study attempts to combine the use of nonlinear Model Predictive Control (MPC) with an online estimation of the updraft environment (without assuming a predefined updraft model) to compute the optimal trajectory to harvest the maximum amount of energy (potential and kinetic) from the atmosphere. In order to obtain information about the surrounding atmosphere, a heuristic search is executed when various conditions are fulfilled in order to improve the chances of the UAV converging to a strong updraft.

The in-house Imperial College London Optimal Control Software (ICLOCS, Falugi et al. [2010]) is used in conjunction with MATLAB. While alternative programs with their respective implementation such as ACADO and CasADi are available, ICLOCS was chosen because of its ease-of-use coupled with its in-house characteristic and code accessibility via MATLAB.

A brief introduction to the layout of the paper is as follows. In Section 2, the problem setup is laid out. In Section 3, the simulation results are presented for three cases: (i) when full updraft knowledge is passed to the UAV, (ii) when full updraft knowledge is withheld from the UAV, and (iii) when a systematic search is used. In Section 4, we undertake a discussion of the results and examine further directions this work can take.

2. PROBLEM SETUP

2.1 Optimal control problem

The optimal control problem given below is used as the framework upon which other model parameters such as the UAV's equations of motion, the constraints and the updraft model are

Parameter (unit)	Value
Mass, m (kg)	300
Wing area, S (m ²)	11
Air density, ρ (kg·m ⁻³)	1.225
Gravitational acceleration, g (m·s ⁻²)	9.81

Table 1. UAV (DG-100 glider) model parameters.

inserted. This form is also used within ICLOCS and will be used for the purposes of this study:

$$\min_{u(t)} J(x(\cdot), u(\cdot), t_f) \quad (1a)$$

$$\text{s.t.: } \dot{x} = f(x(t), u(t)), \quad \forall t \in [t_0, t_f] \quad (1b)$$

$$g_L \leq g(x(t), u(t)) \leq g_U, \quad \forall t \in [t_0, t_f] \quad (1c)$$

$$\phi_L \leq \phi(x_f) \leq \phi_U \quad (1d)$$

$$x_L \leq x(t) \leq x_U, \quad \forall t \in [t_0, t_f] \quad (1e)$$

$$u_L \leq u(t) \leq u_U, \quad \forall t \in [t_0, t_f] \quad (1f)$$

$$u_0 = u(t_0), \quad x_0 = x(t_0), \quad x_f = x(t_f). \quad (1g)$$

The cost function is defined as

$$J(x(\cdot), u(\cdot), t_f) := \int_{t_0}^{t_f} L(x(t), u(t)) dt + E(x_f) \quad (2)$$

where $E(\cdot)$ is the cost associated with the boundary conditions and $L(\cdot)$ is the stage cost function. The cost function is minimised over $u(t)$, the time-varying control inputs, and x_0 , the initial state. Also, $g(\cdot)$ describes the general path constraints and $\phi(\cdot)$ imposes the ending boundary conditions for the phase.

2.2 Aircraft model

We consider a glider-type UAV. A 3 degree-of-freedom (3DOF) representation, which models the aircraft as a point mass in a three-dimensional space, is used. Rate constraints are included to ensure it manoeuvres within the envelope of a real aircraft, and therefore allowing a 3DOF model to be used instead of a more computationally intensive 6DOF model.

The 3DOF model differential equations representing the aircraft are given by

$$\dot{x}_c = v \cos \gamma \cos \chi \quad (3a)$$

$$\dot{y}_c = v \cos \gamma \sin \chi \quad (3b)$$

$$\dot{z}_c = v \sin \gamma \quad (3c)$$

$$\dot{\gamma} = \frac{1}{mv \cos \gamma} (F \cos \mu - mg \cos \gamma) \quad (3d)$$

$$\dot{\chi} = \frac{1}{mv \cos \gamma} (F \sin \mu) \quad (3e)$$

$$\dot{v} = \frac{-1}{m} (D + mg \sin \gamma) \quad (3f)$$

where the horizontal position is given by x_c and y_c , z_c is the altitude, v is the true airspeed, γ is the flight path angle, χ is the heading angle, $F = \frac{1}{2} C_L \rho S v^2$ is the lift and $D = \frac{1}{2} C_D \rho S v^2$ is the drag. The coefficient of lift is $C_L = 0.7 \cdot 2\pi\alpha$ where α is the aircraft Angle Of Attack (AOA) and the factor of 0.7 is used to discount for the three-dimensional flow effects on a finite-span wing. The coefficient of drag is $C_D = 0.01 + 0.02 C_L^2$. Other parameters, corresponding to the ones of a DG-100 glider, are given in Table 1. The input vector used to control the UAV consists of the pitch and roll rates, and is defined as

$$u := (\dot{\alpha}, \dot{\mu}). \quad (4)$$

The state vector consists of the six states in the equations of motion together with the AOA and roll angle. It is defined as:

$$x := (x_c, y_c, z_c, \gamma, \chi, v, \alpha, \mu). \quad (5)$$

States (unit)	Lower bound	Upper bound
x_c (km)	-2	2
y_c (km)	-2	2
z_c (km)	0	∞
γ (°)	-30	30
χ (°)	$-\infty$	$-\infty$
v (m/s)	15	70
α (°)	0	10
μ (°)	-45	45

Table 2. State constraints.

Inputs (unit)	Lower bound	Upper bound
$\dot{\alpha}$ (°/s)	-10	10
$\dot{\mu}$ (°/s)	-30	30

Table 3. Input constraints.

The effect of the updraft (a vertical wind) is incorporated into the equations by adding the relative motion between the air and the ground (Patel [2010]) to the UAV's vertical velocity, i.e.

$$\dot{z}_c = v \sin \gamma + w_z, \quad (6)$$

where w_z is the updraft velocity that needs to be estimated.

2.3 Constraints and Cost Function

Constraints are imposed on some states and these are given in Table 2 (as per ICLOCS methodology, bounds of $\pm\infty$ are used to denote unconstrained states).

The UAV is limited to fly within a 4-by-4 kilometer square space and extending from the ground upwards. In addition to trajectory state bounds, ICLOCS also allows for the specification of a terminal state bound. We have imposed a more restrictive terminal state bound for the UAV's altitude as

$$z_c(t_f) \in [100, \infty]. \quad (7)$$

This ensures that upon termination of control of the flight trajectory, the UAV is not in a precarious position. Constraints are also imposed to the inputs and these are given in Table 3.

We are interested in maximizing the total energy of the UAV at the end of the prediction horizon. Hence, the cost function is selected to have only a boundary cost defined as

$$E(x(t_f)) := - \left(z_c(t_f) + \frac{v(t_f)^2}{2g} \right), \quad (8)$$

and $L(x(\cdot), u(\cdot)) := 0$. The formulation in (8) is also known as the specific energy height of an aircraft which is the total energy (potential and kinetic) divided by the weight of the aircraft (Allen and Lin [2007], Qi and Zhao [2005], Chakrabarty [2010], Zhao [2004]). The minus sign is because problem (1) is defined as a minimization problem.

2.4 Atmospheric Scenario

The updraft scenario chosen for simulation is shown in Fig. 1 (this is a scaled version of MATLAB's peaks function). Unlike other studies that assume single or multiple updraft cores or tubes sufficiently distant and beyond the influence of each other, we have employed a more difficult problem of multiple updraft cores with different magnitudes and radii which have overlapping regions of influence. As a result, there exist multiple local optima. Downdraft areas have also been included to see if the controller can successfully negotiate these areas.




Fig. 1. Contour plot of the atmospheric updraft scenario. The units of the contour (z direction) are m/s.

In addition, in order to examine the performance of the simulation, the global updraft distribution has been chosen such that, on average, a purely random flight will result in the UAV encountering a downdraft as opposed to an updraft (the expected vertical wind is negative). This is on top of the natural rate-of-sink that the unpowered glider-type UAV experiences.

2.5 Updraft Velocity Online Estimation

As the UAV does not have full updraft distribution information, a method of estimating the updraft environment is required. A two-layer Generalized Regression Neural Network (GRNN) is used to regress the gathered updraft velocity data that the UAV obtains along its flight path. A ‘picture’ of the surrounding atmospheric updraft is created and used to solve the optimization problem.

Unlike other studies where the UAV is either assumed to contain special equipment, such as forward-sensing infrared thermal cameras (Kagabo [2010]) or where the UAV is directly fed the actual updraft environment in its accurate entirety (Kowalczyk and Olinski [2011]), we have restricted the UAV’s knowledge to the measured updraft velocities along its flight path at the respective time instants only. This makes the problem more difficult but reflective of reality. Should the project eventually be implemented on a real aircraft, the ability of the algorithm to work without depending on an unrealistic piece of equipment will be desirable. It also helps increase the robustness of the solution. By achieving a method that can work under very restrictive conditions, any subsequent methods that provide additional information would only serve to improve performance.

2.6 Heuristic Search Model

A heuristic search model is employed to ensure that:

- (1) the UAV finds acceptably/sufficiently strong updrafts
- (2) the UAV flies a path to obtain sufficient information for the estimation model to produce accurate results.

The heuristic model works as follows:

- (1) ensure that the UAV is at or higher than a ‘safe altitude’ before considering any heuristic search
- (2) if either the Root Mean Squared Error (RMSE) between the measured and estimated updraft is larger than an

Parameter (unit)	Value
Safe altitude (m)	300
RMSE threshold (m/s)	1
Reasonable updraft strength (m/s)	2.5
Heuristic search time (s)	60

Table 4. Heuristic search parameters.




Fig. 2. UAV trajectory using a control horizon of 60 s. In this plot, as in the following similar ones, the numbers next to the path are seconds along the path.

‘RMSE threshold’ or the maximum measured updraft value is less than a ‘reasonable updraft strength’, then execute the heuristic search

- (3) generate a reference direction for the UAV to fly for a set time (‘heuristic search time’).

Table 4 lists the heuristic search parameters that were used. Therefore, the heuristic search model ensures that the UAV ‘settles’ on a updraft only when it has obtained a reasonably good fit of the updraft environment, and the updraft is of a sufficiently high strength given the average surrounding conditions. Furthermore, the heuristic search ensures that the UAV is in no imminent threat of landing while performing a search by setting a ‘safe altitude’.

3. SIMULATION RESULTS AND DISCUSSION

The simulation that will follow are closed-loop results obtained from implementing the optimal controller (solution to (1)) in a receding horizon strategy with a sampling time of 2 s.

3.1 Full Updraft Knowledge

In the first set of simulations, the UAV is simulated with the full environmental updraft information being available to the controller. The UAV is released adjacent to and pointing in the general direction of the strongest environmental downdraft in the global distribution. In the first case, with a control horizon of 60 s, it is observed that the UAV successfully engages in a bank (turn) and negotiates away relatively quickly from the nearby downdraft area (Fig. 2). However, it converges to the nearby local optimum. This is despite having full knowledge that a stronger, globally optimum updraft exists.

In the next case, a control horizon of 120 s is used in order to see if the UAV would attempt to converge to the strongest updraft.




Fig. 3. UAV trajectory using a control horizon of 120 s.




Fig. 4. Comparison of energies between two simulations with differing control horizons — 60 s (red) and 120 s (blue).

From Fig. 3 we observe that this is indeed the case and the UAV performs the same aggressive avoidance manoeuvre of the downdraft and tracks for the global optimum. This highlights the classic problem in nonlinear MPC of the need to select an appropriate control horizon.

On the way to the global optimum, the UAV does not take the closest Euclidean distance. Instead, it is skewed to pass through the intermediate updraft, thereby increasing its altitude and hence potential energy. In a sense, the UAV has taken the shortest energy distance to the global optimum.

For the first 40 s, the energies and trajectories of both simulations were largely similar. After 40 s however, the UAV with the shorter control horizon proceeded to loiter at low speed within the intermediate updraft. From Fig. 4 we can observe its energy increasing very slowly. This is because the updraft in the region is just over 1 m/s while the UAV's minimum sink rate is approximately 0.62 m/s.

The UAV with the longer control horizon increased its speed by decreasing its AOA. In the process, it traded potential for kinetic energy and lost more altitude than the first UAV. However, this allowed it to reach the strongest updraft core (up to 4 m/s) more quickly (in approximately 100 s). The result is a much higher energy state at the end of 200 s as seen in Fig. 4.




Fig. 5. UAV trajectory overlaid upon the actual updraft environment.




Fig. 6. UAV trajectory overlaid against the estimated updraft environment.

3.2 Online Estimation Without Full Updraft Knowledge

In the next set of simulations, the environmental updraft information is withheld from the UAV and with 120s control horizon. The updraft information is estimated using the GRNN and an optimal trajectory is flown at each step using only the known information up to that point in time.

When the conditions for a heuristic search are met and the heuristic search is enabled, a randomly-generated course is assigned to the UAV for the duration of the random search.

From Fig. 5 we can see the UAV searching for more than 350 s before it converges on the global maximum updraft. By the end of the trajectory, it had estimated an updraft picture of its surroundings as shown in Fig. 6.

In the 600-second simulation horizon, the UAV has gained 220 m of specific energy. The UAV required 210 s (35% of the total simulation time) before it was aware of the large updraft system. Thereafter, it required another 140 s (23% of the time) to probe the bounds and character of the updraft structure. Finally, 560 s after release, the UAV effectively mapped out an accurate structure of the largest updraft system and was circling around the strongest core.




Fig. 7. UAV trajectory with systematic search.

This compares against the 90 s and the 150 m of maximum specific energy loss for the UAV with full updraft knowledge to converge to the global optimum (Fig. 3 and 4).

3.3 Online Estimation Without Full Updraft Knowledge With Systematic Search

Combining the knowledge gained from the results of the random heuristic search and the fact that in the real world the optimal updraft could be within any grid space with almost equal probability, a systematic procedure for searching the environment could prove to be more robust.

In our systematic search simulations, the environment is dissected into grid squares (600-by-600 m) and a systematic search of the space surrounding the UAV is conducted (Lee [2012]). An optimum route for search is computed as the UAV explores its environment and the problem is in many ways similar to the Travelling Salesman Problem (Stone [1975], Castanon and Sandell [1979], Stromquist and Stone [1981], Raphael and Smith [2000]). Fig. 7 shows the case for a typical UAV release.

Because the UAV had lost too much altitude and breached the safe altitude of 300 m at approximately 500 s, the systematic search was overridden at around 520 s. Thereafter, the UAV proceeded to the best updraft it had knowledge of and circled around the suboptimal updraft as seen in Fig. 7. The UAV's limited knowledge of the environment can be observed in Fig. 8.

3.4 Online Estimation Without Full Updraft Knowledge With Systematic Search and Adaptive Grid

With this understanding of the search dynamics of the UAV, an adaptive grid search method was implemented in order to increase the probability of discovering the best updraft whilst at the same time expediting the search in areas where was observed to avoid losing excessive altitude.

The adaptive grid search works by increasing (coarsening) the grid spacing when the UAV's observed updraft gradient over the past 20 s is low or negative. Conversely, the grid space is refined when the trailing 20-second history of the updraft gradient is strong and positive. This allows the UAV to promptly negotiate updraft and downdraft gradients, while not sacrificing detail when mapping out suitably strong updrafts. The adaptive grid




Fig. 8. UAV trajectory with systematic search overlaid against the estimated updraft environment.




Fig. 9. Trajectory with systematic search and adaptive grid.

size used here ranges from between 400-by-400 to 800-by-800 m.

Fig. 9 shows the UAV's trajectory for the same release conditions as in Fig. 7. At 255 s, the UAV traversed and experienced an updraft greater than the heuristic search condition of 2.5 m/s. Thereafter, it proceeded to terminate the heuristic search and relied upon its online estimation of the environment to navigate towards the globally optimal updraft.

Not only has the adaptive grid enabled the UAV to locate the best updraft area, but it has done so more quickly, enabling a much higher ending specific energy state as seen in Fig. 10. Note that, under the same conditions, the UAV with the non-adaptive grid search (Fig. 7) would have been forced to land eventually as it was unable to extract sufficient energy from the environment (Fig. 10).

4. DISCUSSIONS AND FUTURE WORK

In the exploration of the surrounding updraft picture around the UAV, a random search direction could possibly yield good results. However, utilizing a systematic search promises to be more robust in the real world where the probability of the optimum updraft could be anywhere in the absence of any a priori knowledge or forward-sensing equipment.




Fig. 10. Comparing the different specific energy histories.

It is observed that using an adaptive grid spacing for the systematic search expedites exploration of the UAV's surrounding environment whilst retaining adaptability to changing gradients. This in turn promises better results at converging to a stronger updraft within a potentially shorter time.

In the online estimation of the updraft environment, the current implementation equally considers every observed updraft data point along the UAV's trajectory. Realistically however, recently-gathered information will have higher informational value compared to that observed further in the past. A possible method for this could be to implement a schedule of decreasing weights.

Moving forward, an implementation of the code in alternative programs such as ACADO and CasADi could be performed with the intention to ultimately transit toward real-time, hardware-in-the-loop simulations and actual implementation. Current solution times are of the order of the sampling time; a 2-second control step requires approximately 2 s to solve for. However, optimizing the code, alternative implementations and faster methods (e.g. using non-uniform hold constraints as in Longo et al. [2011]) would be expected to decrease this time dramatically and will be the subject of future work.

5. CONCLUSION

This paper has shown the advantages of using constrained optimal control techniques on glider-type UAVs for the exploitation of energy from atmospheric updraft currents (vertical winds). Aircraft can take advantage of this atmospheric energy for improving endurance, loiter time and increasing fuel efficiency. In addition, using an adaptive search grid for exploration of the environment promises to be more successful than fixed-grid spacing or searching in random directions. The results of this study are significant as it lends conceptual support to the feasibility of utilizing glider-type aircraft to harvest the naturally-occurring energy present in atmospheric updrafts.

ACKNOWLEDGEMENTS

This research has been supported by the EPSRC grant number EP/G031576/1, EP/F041004/1 and the European Union Seventh Framework Programme FP7/2007-2013 'EMBOCON' under grant agreement number FP7-ICT-2009-4 248940.

REFERENCES

- Z. Akos, Nagy M., S. Leven, and T. Vicsek. Thermal soaring flight of birds and unmanned aerial vehicles. *Bioinspiration & Biomimetics*, 5(4), 2010.
- M.J. Allen and V. Lin. Guidance and control of an autonomous soaring UAV. Technical Report TM-2007-214611, NASA Technical Manual, 2007.
- D.A. Castanon and N.R. Sandell. Optimal search trajectories for a single platform. In *Proc. of the 18th IEEE Conf. on Decision and Control*, volume 18, pages 574–575, 1979.
- A. Chakrabarty. *Flight path planning of UAV based on atmospheric energy harvesting*. PhD thesis, Pennsylvania State University, USA, 2010.
- M.J. Cutler, T.W. McLain, R.W. Beard, and B. Capozzi. Energy harvesting and mission effectiveness for small unmanned aircraft. In *Proc. of the AIAA Guidance, Navigation and Control Conf.*, volume 15, 2010.
- P. Falugi, E.C. Kerrigan, and E. Van Wyk. Imperial College London Optimal Control Software (ICLOCS) @ONLINE, 2010. URL <http://www.ee.ic.ac.uk/ICLOCS/>.
- E. Fonseca. *Modeling and flying thermal tubes with a UAV*. PhD thesis, ETH Zurich Autonomous Systems Lab, Switzerland, 2007.
- W.B. Kagabo. *Optimal trajectory planning for a UAV glider using atmospheric thermals*. PhD thesis, Rochester Institute of Technology, USA, 2010.
- N.E. Kahveci, P. A. Ioannou, and Mirmirani M.D. Adaptive LQ control with anti-windup augmentation to optimize UAV performance in autonomous soaring applications. *IEEE Trans. Control Systems Technology*, 16(4):691–707, 2008.
- Z. Kowalczyk and K.E. Olinski. Gliding strategy design with the use of discrete optimisation. In *Proc. of the 17th IFAC World Congress*, pages 14011–14016, 2011.
- D. Lee. Control of Energy-Harvesting Gliders. Master's thesis, Imperial College London, UK, (available on request), 2012.
- S. Longo, E.C. Kerrigan, K.-V. Ling, and G.A. Constantinides. A parallel formulation for predictive control with nonuniform hold constraints. *Annual Reviews in Control*, 35(2): 207–214, 2011.
- R.B. Patel. *Prohibited volume avoidance for aircraft*. PhD thesis, Imperial College London, UK, 2010.
- Y.C. Qi and Y.J. Zhao. Energy-efficient trajectories of unmanned aerial vehicles flying through thermals. *J. of Aerospace Engineering*, 18(2):84–92, 2005.
- B. Raphael and I.F.C. Smith. A probabilistic search algorithm for finding optimally directed solutions. In *Proc. of Construction Information Technology*, pages 708–721, 2000.
- H.D. Shannon, G.S. Young, M.A. Yates, M.R. Fuller, and W.S. Seegar. Measurements of thermal updraft intensity over complex terrain using american white pelicans and a simple boundary-layer forecast model. *Boundary Layer Meteorology*, 104:167–199, 2002.
- L.D. Stone. *Theory of optimal search*. Academic Press, Inc., 1975.
- W.R. Stromquist and L.D. Stone. Constrained optimisation of functionals with search theory applications. *Mathematics of Operations Research*, 6(4):518–529, 1981.
- J. Wharington. Control of a high endurance unmanned air vehicle. Technical report, International Council of the Aeronautical Sciences and the American Institute of Aeronautics and Astronautics, 1998.
- Y.J. Zhao. Optimal patterns of glider dynamic soaring. *Optimal Control Applications and Methods*, 25:67–89, 2004.

Predictive control for soaring of unpowered autonomous UAVs

Lee, Darren

2012-08-27

Attribution-NonCommercial-NoDerivatives 4.0 International

Darren Lee, Stefano Longo and Eric C. Kerrigan. Predictive control for soaring of unpowered autonomous UAVs. IFAC Proceedings Volumes, Volume 45, Issue 17, 2012, Pages 194-199

<https://doi.org/10.3182/20120823-5-NL-3013.00021>

Downloaded from CERES Research Repository, Cranfield University