Action Selection Algorithms for Autonomous Microspacecraft

E. A. Gillies,* A. G. Y. Johnston,† and C. R. McInnes‡
Glasgow University,
Glasgow, Scotland G12 8QQ, United Kingdom

Introduction

THE new generation of microspacecraft currently under development for Earth orbiting missions will require a high degree of autonomy to meet stringent cost and performance goals.[1] The microspacecraft may have to perform multiple, conflicting tasks, such as pointing its solar array toward the sun or communicating with a ground station, to ensure it operates within limits and performs a useful function. The traditional, deliberative methods of artificial intelligence require complex world models to reason, whereas neural network and fuzzy logic approaches are somewhat inscrutable and difficult to validate to ensure spacecraft survival.[2] In contrast, recent concepts for artificial agents borrow heavily from ethology, where the agent responds directly to environmental stimuli.[3]

In this Note such an artificial agent approach is proposed that provides a method for action selection that balances the demands of the satellite users (for example, gathering or communicating data, thus draining the spacecraft batteries) and the actions necessary to ensure spacecraft survival (for example, charging the batteries with a solar array). We have adopted the cue-deficit action selection algorithm of McFarland and Spier[4] because it is directly developed from optimal control theory and is computationally simple to implement on a microspacecraft.

The spacecraft is modeled as a nonlinear dynamic system with a state space consisting of key internal variables (battery charge, memory state, temperature, transmitted data bits, and attitude). The state space has a set of limits that defines the satellite's useful operating domain. A finite repertoire of behaviors is then used to control the internal dynamics of the spacecraft. A cost function is provided that measures the deviation of the spacecraft from its normal state-space operating point. Application of optimal control theory yields the optimum action selection rules. The action selection rules must maintain a position close to this operating point in the presence of perturbations due to the spacecraft's own behavior. For example, transmitting data may be necessary for the satellite to perform a useful function but will drain the batteries.

The proposed action selection rules are found to display a degree of opportunism: For example, we show during a simple satellite simulation that the satellite charges its batteries when the opportunity arises (and not just when the batteries become low on charge). We demonstrate that in the event of major hardware failures the algorithm will resequence the spacecraft actions to ensure survival and to continue to achieve its goals, albeit in a degraded manner.

Satellite Model

We consider a simplified model of an Earth orbiting satellite and its subsystems. The satellite operates in a low circular or elliptical orbit, coplanar to the ecliptic, and is considered to have a single rotational degree of freedom about the normal to the orbit plane that can be controlled by a reaction wheel.

The electrical power system model consists of a solar array, battery (with upper and lower charge limits), and various electrical loads. The payload consists of a camera that records a steady rate of data when activated, a solid state memory and a radio transmitter to broadcast data to a ground station. The model also calculates the satellite's equilibrium temperature (based on external heat inputs from the sun and Earth albedo and internal energy dissipation). The individual subsystems are coupled together: for example, switching on the transmitter drains the satellite's battery and reduces the quantity of data stored in the onboard memory.

The satellite is controlled by switching the camera, the transmitter, and an internal heater on or off and commanding an attitude control subsystem to track one of three targets by moving the reaction wheel (point solar array to sun, point camera to target, or point transmitter to ground station) or to drift. To provide pointing constraints, the solar array, camera, and antenna are on different faces of the cube-shaped satellite. The internal heater may be switched on or off independently of what other task the satellite performs; the heater subsystem activates automatically when the temperature drops below a certain limit and is not commanded by the action selection algorithm. The heater, however, drains the batteries and, therefore, influences the action selection indirectly. The model is constructed using the SIMULINK package.

Autonomous Action Selection

The algorithm for satellite action selection presented here borrows heavily from the work of Spier and McFarland[3] and McFarland and Spier,[4] who have suggested and implemented action selection mechanisms for mobile robotic[4] and artificial agent platforms.[3] Because the method is rooted in optimal control theory and also because of its computational simplicity (the method only requires calculation and comparison of products and involves minimal internal or environmental modeling), we suggest the method for microsatellite action selection.

To ensure survivability, our model satellite must never drain its batteries below maximum depth of discharge. The satellite energy deficit d_{energy} is a measure of how much the batteries have discharged. This deficit may be lowered by pointing the solar array toward the sun. Alternatively, the energy deficit may be minimized by powering down unessential systems and drifting. When charging up or drifting, the spacecraft is not able to perform other useful tasks such as recording data with its payload and transmitting this data to Earth via a ground station. The users of the satellite may define a work deficit d_{work} for the spacecraft that increases all of the time that the satellite is not recording or transmitting data. We have defined a work deficit that is dependent on the average rate that the satellite manages data, either from recording with its camera or by transmitting to the ground station. The work deficit is the difference between the satellite's actual rate and a target data rate. Essentially our model is a two-resource problem of energy and work. The satellite control algorithm must ensure that neither deficit becomes too large and also that the satellite switches from one behavior to the other at an appropriate point. The deficits are normalized such that $0 < d_i < 1$.

The satellite is equipped with sensors or an internal model that determines the availability r_i of any resource i in the environment. For example, $r_{\text{solar}} = 1$ when the satellite detects (via a sun sensor) that it is in sunlight, and $r_{\text{solar}} = 0$ when the satellite is in eclipse. The availability $r_{\text{transmit}} = 1$ when the satellite has data stored in memory and the ground station is present [the ground station proximity is ascertained using a global positioning system (GPS) or by detecting a ground signal]. Otherwise, $r_{\text{transmit}} = 0$. Similarly, $r_{\text{record}} = 1$ when there is space in the spacecraft memory, and it is zero otherwise. The satellite is always able to power down and drift and so $r_{\text{drift}} = 1$.

The rate at which the spacecraft can gain energy or perform work by any task is modeled by parameters k_i. For example, the rate at which the energy deficit can be reduced by charging, k_{solar}, is the maximum array power output (normalized with the value at launch). The increase in energy deficit may alternatively be minimized by powering down and drifting; k_{drift} is, therefore, selected to be a small fraction of k_{solar}, reflecting that it is better for the satellite to prolong its operational life by recharging than by powering down and drifting. If the satellite detects that the solar array is damaged (by monitoring previous power outputs during charging cycles), then k_{solar} is lowered, for example, if 50% of the array fails at time $t = t_f$, then $k_{\text{solar}}(t_f) = 0.5k_{\text{solar}}(t_{\text{launch}})$. This allows reprioritization of tasks after equipment failure. The optimal transmit data rate k_{transmit} is similarly selected to be unity for an undamaged transmitter, and it

Received 31 July 1998; revision received 1 June 1999; accepted for publication 9 June 1999.

*Lecturer, Department of Aerospace Engineering.
†Research Student, Department of Aerospace Engineering.
‡Reader, Department of Aerospace Engineering.

is lowered should the actual transmission baud rates degrade during the satellite lifetime. The recording data rate k_{record} is selected to be a value less than the transmit rate for our simulation (it is important for the satellite to communicate any data it records).

By formulating the problem in this way, Spier and McFarland[3] have shown that (on robot platforms) a quadratic cost function (the cost of possessing any particular deficit increases at a greater than linear rate, a quadratic cost is also mathematically simple) and Pontyagrin's maximization principle can be used to find the optimal control. They found that the optimal control could be summarized by performing a behavior associated with the highest deficit $\times$ availability $\times$ rate (drk) product.

The spacecraft selects the optimum behavior at any time by computing the energy and work deficits, taking environmental cues and in-built models to compute the availabilities of solar energy and cues for doing work in the environment, and finally computing drk products for each behavior. In our model, the energy deficit is measured directly from the battery state, and the work deficit is calculated by a simple internal model. The availability of sunlight is detected by a sensor, and the ground station is detected by either a sensor or a GPS model. The rates at which the satellite can reduce its work or energy deficits are simply measured from previously attained values. The action selection process, therefore, involves little environment or internal modeling/planning and mainly relies on sensed or measured quantities and is consequently computationally simple. Uncertainty about measured cues will mean that the spacecraft will be maintained close to a point in state space that is displaced from the optimum point. The behavior of the algorithm should not be affected, but the satellite will operate suboptimally. Failure of the sun sensor or uncertainty in GPS position is more critical as the spacecraft would then receive no cues for charging or communication. If these failures were detected, then the spacecraft would not be able to continue autonomous operation.

We have found that this simple algorithm readily manages the spacecraft tasks and has advantages over a controller that merely selects the task with the most pressing deficit. Figure 1 shows typical cue-deficit products for the spacecraft during a mission simulation and compares them to the deficits in work and energy. There are several points where the satellite selects a behavior that does not have the highest (or even high) deficit. For example, if the satellite work deficit is high and the battery deficit low, the satellite may, for instance, still charge up its batteries if sunlight is available and cues for doing work (e.g., visibility of the ground station) are low.

Results for a simulation where half of the solar array fails during the mission are shown in Fig. 2. The array failure occurs after 47 orbits (91 orbits are shown). The battery charge for the spacecraft operating with its full solar array settles down to a predictable pattern after an initial small transient of a few orbits (resulting from a difference between the satellite's launch conditions and the normal operating point in orbit). The satellite also begins to record and

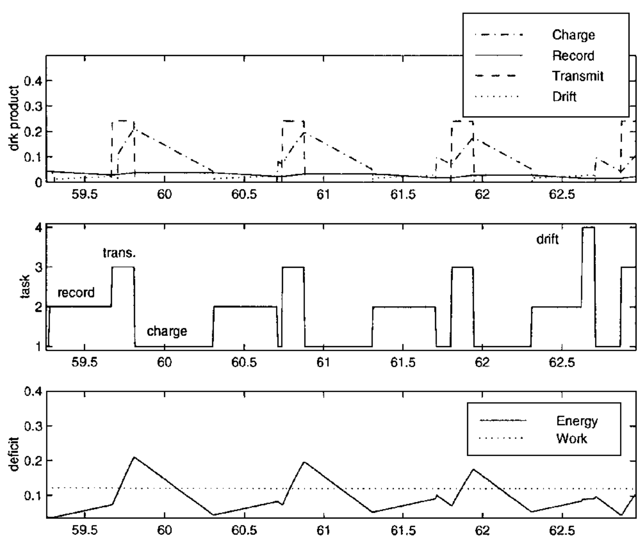

Fig. 1 Cue-deficit products compared to deficits for spacecraft.

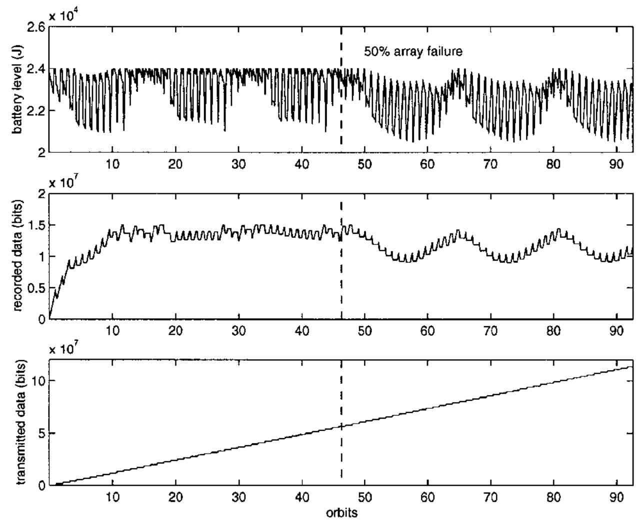

Fig. 2 Performance of spacecraft during array failure.

transmit data. After approximately 47 orbits, half of the solar array is failed to examine the robustness of the behavior sequencing algorithm. The rate parameter for charging changes, so that the pattern of task selection for the spacecraft changes. The battery charge, after an initial transient of a few orbits, settles down to another pattern (which has a slightly lower average charge, but nevertheless ensures satellite survivability). The satellite now spends more time drifting in a power down mode than prior to the array failure. Consequently, the satellite spends less time recording data. The total data throughput is lower, but the satellite behavior is robust to this array perturbation. Similar resequencing of satellite tasks is observed after damage to the payload or transmitter.

When two or more mutually exclusive tasks have a very similar drk product and when the resources associated with the tasks are tightly coupled (e.g., when transmitting data is expensive in terms of battery charge), then oscillatory behavior may emerge. We solved this oscillation between behaviors by introducing a threshold value such that the satellite behavior is only switched when the drk products differ by the threshold value.

Scalability of the Approach

Our simulation is somewhat simplified, and an actual spacecraft/microsatellite may have many more operational tasks that may be autonomously controlled or be scheduled or commanded by ground control. Unlike some other action selection algorithms, such as subsumption architecture, the cue-deficit model may easily incorporate scheduled tasks because it does not lock out any behavior. A scheduled task may be built in via a simple internal model by setting the deficit and cue to equal unity at a particular time.

McFarland and Spier[4] have argued that adding extra tasks to an autonomous agent using the cue-deficit action-selection model is straightforward; each new behavior is simply given an associated deficit, availability, and accessibility. The selected behavior is still the one with the highest drk product. Scaling the action selection algorithm up to the complexity of an actual spacecraft problem should, therefore, be relatively straightforward. McFarland and Spier were able to scale up their action-selection algorithms from simple two-dimensional artificial agent simulations to actual robotic experiments in a complex environment.

One great advantage of the cue-deficit method for action selection is that the satellite tasks are selected using measurement of environmental cues (such as the presence of sunlight or the ground station) and measured internal parameters (such as battery charge and memory level). Complex internal models of the environment are not required to select the appropriate behavior. Considering the battery charge as an example, the particular battery model used in our simulation is not directly relevant to the performance of the algorithm (other than increasing the fidelity of the simulation); the algorithm uses the measured battery charge rather than using a model of the battery charge. Therefore, it is expected that the modeling of more complex, and more numerous, spacecraft subsystems in the simulation would not change the qualitative behavior of the algorithm.

The authors and co-workers are currently extending the procedure to manage cooperation between constellations of satellites.[5]

Conclusions

We have outlined a scheme for sequencing of tasks on a microspacecraft platform. The scheme is easily implemented by virtue of its computational simplicity. Moreover, the strategy is derived from optimal control theory. Action selection with the cue-deficit algorithm is seen to perform well on our microsatellite simulation and is demonstrably robust to significant failures of the satellite solar array. Our work has also demonstrated robustness to failures of payload, transmitters, and injection into nonoptimum orbits (such as elliptical transfer orbits, rather than the mission orbit) although these results are not shown here for brevity.

Acknowledgment

This work was partially supported by a Grant from the Leverhulme Trust.

References

[1]Creasey, R., "Implementing Spacecraft Autonomy Using Discrete Control Theory," *Space Technology*, Vol. 16, No. 2, 1996, pp. 109–116.

[2]Horvath, J., "Spacecraft Autonomy Issues: Present and Future," *Journal of the British Interplanetary Society*, Vol. 49, No. 6, 1996, pp. 215, 216.

[3]Spier, E., and McFarland, D., "Finer-Grained Motivational Model of Behaviour Sequencing," *4th International Conference on Simulation of Adaptive Behaviour*, MIT Press, Cambridge, MA, 1996, pp. 255–263.

[4]McFarland, D., and Spier, E., "Basic Cycles, Utility and Opportunism in Self-Sufficient Robots," *Robotics and Autonomous Systems*, Vol. 20, Nos. 2–4, 1997, pp. 179–190.

[5]Radice, G., Gillies, E., Johnstone, G., and McInnes, C., "Autonomous Action Selection for Micro-Satellite Constellations," International Astronautical Foundation, Paper ST-98-W.1.05, 1998.

Effects of Planar Thrust Misalignments on Rigid Body Motion

Daniel Levy de F. Rodrigues* and
Marcelo Lopes de O. e Souza†
National Institute for Space Research,
12227-010 São José dos Campos, Brazil

Introduction

SPACE vehicles are frequently designed to have propulsive systems capable of generating longitudinal forces with the thrust axis passing through the c.m. of the vehicle. However, in many practical cases this does not occur, and a disturbance torque appears. The misalignments can be angular or linear. Such misalignments usually have low magnitude and are mainly caused by the structural design of the vehicle, vibrations, the motion of movable parts (especially liquids), and displacements of the c.m. caused by fuel consumption.[1] Schwende and Strobl[2] mention as reasonable an angular misalignment $\delta < 0.002$ rad for a typical apogee motor.

In the case of orbital transfers, the misalignment torque—even with small magnitude—is very important because it alters the force orientation with respect to the specified strategy and, as a consequence, produces an error with respect to the desired final orbit, thus generating the necessity of implementing a fast and precise active attitude control during the orbital transfer. Longuski et al.[3] present a control proposal that eliminates the damage caused by this misalignment during propulsive maneuvers. Their proposal consists of splitting it in two parts, intercalated by a time interval without propulsion. The spin stabilization is a strategy that has also shown very much use in apogee motors, by canceling the torque of the transverse misalignment.

In this work we intend to verify analytically the effects of planar thrust misalignments on rigid body motion, supposing the absence of a fast and precise active attitude control during the propulsion. A numerical example illustrates the situation for a hypothetical vehicle. Details are in Chap. 5 of Rodrigues.[4]

Simplified Analytical Study

Consider a model constituted by a rigid and constant-mass vehicle subject only to a unique force with a constant orientation with respect to the vehicle's principal inertia axes and with its application axis not intercepting the c.m. (Fig. 1). In this way, the trajectory described by the c.m. remains in the plane that contains this force and the c.m. Thus, we can choose a reference system $OXYZ$ whose plane XY contains the referred trajectory.

The equations of motion of the attitude and of the c.m. motion with respect to $OXYZ$ can be obtained as

$$I\ddot{\theta}(t) = F\varepsilon \qquad \text{where} \qquad \varepsilon = h\sin\delta + \nu\cos\delta \tag{1}$$

$$m\ddot{\boldsymbol{R}}(t) = \boldsymbol{F}(t) \tag{2}$$

where m is the body mass (supposed constant), $\boldsymbol{R}$ is the c.m. position vector with respect to the inertial reference system $OXYZ$, $\boldsymbol{F}$ is the resulting force written in terms of the $OXYZ$ vectors, h is the longitudinal distance of the ideal thrust position to the c.m., I is the inertia moment with respect to the rotation axis Z, θ is the attitude angle, ν and δ are the linear and the angular misalignments of the thrust vector, and t is the time. From Eq. (1) we get

$$\dot{\theta}(t) = (F\varepsilon/I)t + \dot{\theta}(0) \tag{3}$$

$$\theta(t) = (F\varepsilon/2I)t^2 + \dot{\theta}(0)t + \theta(0) \tag{4}$$

To avoid some mathematical difficulties, we consider zero initial conditions: $R(0) = 0$, $\dot{R}(0) = 0$, $\theta(0) = 0$, and $\dot{\theta}(0) = 0$. For the components of the c.m. position vector $X(t)$ and $Y(t)$, Eq. (2) becomes

$$m\ddot{X}(t) = F\sin[(F\varepsilon/2I)t^2 - \delta] \tag{5}$$

$$m\ddot{Y}(t) = F\cos[(F\varepsilon/2I)t^2 - \delta] \tag{6}$$

Approximate Analytical Solution

It is possible to obtain approximate solutions to Eqs. (5) and (6) for t near zero and for t growing without bound. Consider initially the case where $t \to 0$. In this case the relations $\cos\theta(t) \cong 1$ and $\sin\theta(t) \cong \theta(t)$ hold. Normally δ is very small, thus we can write $\cos\delta \cong 1$ and $\sin\delta \cong \delta$. Then Eqs. (5) and (6) can be rewritten as

$$m\ddot{X}(t) \cong (F^2\varepsilon/2I)t^2 - F\delta \tag{7}$$

$$m\ddot{Y}(t) \cong F \tag{8}$$

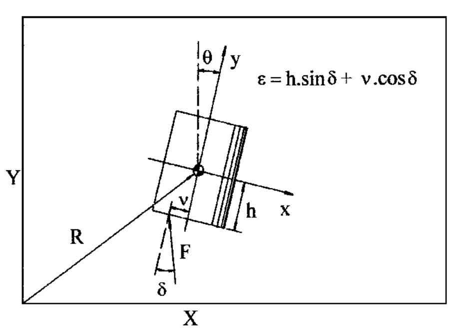

Fig. 1 Linear and angular misalignments of a simplified system in planar motion.

Received 2 November 1998; revision received 1 April 1999; accepted for publication 12 May 1999.

*Corresponding Researcher, Division of Space Mechanics and Control-DMC, INPE, Av. dos Astronautas, 1758-CP 515.

†Senior Researcher, Division of Space Mechanics and Control-DMC, INPE, Av. dos Astronautas, 1758-CP 515; marcelo@dem.inpe.br.