ACHIEVING IMPROVED MISSION ROBUSTNESS *

John Hobbs[1] and Philip Husbands[2] and Inman Harvey[2]
[1] Matra Marconi Space, Bristol, UK
[2]COGS, University of Sussex, UK

Abstract

This paper gives background to and outlines a proposed programme of work aimed at exploring biologically inspired approaches to developing a potentially spaceworthy planetary exploration vehicle. A major consideration is in achieving a very high high degree of robustness and reliability. In order to develop an improved system level mission design, an integration of evolutionary robotics and 'traditional' space mission design methodologies is proposed.

1. OVERALL CONCEPT

The success of currently planned near/medium term planetary exploration missions is almost invariably based on unmanned autonomous or semi-autonomous activities requiring high levels of reliability and mission robustness. Studies into such mission scenarios show that conventional European space technology, even with optimistic predictions of near-term development, is unlikely to achieve the necessary system performance. Using traditional design techniques, current robotic vehicles need to exhibit virtually fault-free behaviour of complex mechanism architectures in order to achieve mission objectives. Mechanisms in question range from launch lock-out clamps and payload deployment devices, to locomotion subsystem drives, and manipulators.

Major improvements in the capability and robustness of such missions are promised by work currently in hand at the University of Sussex in the UK. This has concentrated on the evolutionary design of control systems in noisy and uncertain conditions. The evolutionary strategy ensures that highly-rated designs have been tested under many conditions, typically using real hardware as well as in simulation. Robustness in the face of sensor faults, electronic hardware faults, and mechanical faults, can naturally be incorporated into the evaluation criteria driving the design process. Failure, or partial failure of individual mechanisms within a conventional design typically threatens curtailment of mission objectives, but the evolutionary approach under hostile testing can promote designs that still operate in altered fashions when damaged. Finding hard test-cases is itself a difficult problem, but can be achieved through co-evolution of test-cases with solutions.

The object of the proposed programme of work described here is to develop the design for a potentially spaceworthy planetary exploration vehicle incorporating evolutionary robotics principles; to progress this into appropriate breadboarding stages, and plan the target mission configuration and Engineering Model phases. The target mission chosen is the ROSA/M programme (subject of the recent ESA Invitation to Tender AO/3094), which is aimed at robotic deployment of instruments on the surface of Mars, in about 2003.

2. BACKGROUND TO EVOLUTIONARY AND BIOLOGICALLY INSPIRED TECHNIQUES

2.1 Introduction

Control systems for an autonomous robot need to coordinate action with perception, and be robust to a (relatively) wide range of conditions never previously met.

Traditional design approaches to such complex control systems have been "Divide and Conquer"; a complex problem is decomposed into separate, less daunting, sub-problems. However, the interactions between such sub-problems must be few in number; for robot control systems there are at least four major problems.

- It is not clear *how* a robot control system should be decomposed.
- Interactions between separate sub-systems are not limited to directly visible connecting links between them, but also include interactions mediated *via the environment.*
- As system complexity grows, the number of potential interactions between sub-parts of the system grows *exponentially.*
- Interactions that have not been anticipated can lead to catastrophic failures under unforeseen circumstances.

Classical approaches to robotics have often assumed a primary decomposition into Perception, Planning and Action modules; these in turn to be divided into further sub-modules, e.g. for vision, route planning, motor coordination. Many people now see this as a basic error [2]. Brooks' alternative subsumption architecture

---

approach is incremental in a way inspired by natural evolution. Initially simple behaviours are 'wired into' a robot, and thoroughly debugged, before adding the next behaviour. Brooks still requires design by humans, and this approach still suffers from the problems listed above.

So an obvious alternative approach is to explicitly use evolutionary techniques to incrementally evolve increasingly complex robot control systems, rather than attempt to figure out each evolutionary step by hand design. Unanticipated and elusive interactions between sub-systems need not directly bother an evolutionary process where the only benchmark is the behaviour of the whole system. This is the approach taken in the Evolutionary Robotics research at Sussex, within the Evolutionary and Adaptive Systems group. Related work on an evolutionary approach to agent control using neural networks has been done by Beer and Gallagher [1], Colombetti and Dorigo [4], Floreano and Mondada [5], Nolfi [10].

Genetic Algorithms (GAs) are the most common form of algorithm which uses evolutionary ideas for search, optimisation and machine learning — the fields covered in [6]. Evolutionary Robotics typically needs adaptive improvement techniques, rather than optimisation techniques, and this critical but little-understood distinction needs to be made clear. The Species Adapted Genetic Algorithm (SAGA) system [7] was developed to provide this. To maintain a selection/mutation balance in SAGA requires maintaining a constant selective pressure, which normal fitness-proportionate selection does not provide; and mutation rates should be maintained at a rate of the order of 1 mutation per genotype [8].

In contrast to classical AI views of rule-based control systems, we see the robot as a whole — body, sensors, motors and control system or 'nervous system' — as a dynamical system coupled (via the sensors and motors) with a dynamic environment [1]. This coupled interaction generates the robot behaviour which is to be evaluated. The control system should itself be a dynamical system, and hence the genetic specification of this should be at the level of the primitives of a dynamical system. One convenient form of dynamical system is an (artificial) neural net (ANN). Continuous time recurrent neural networks can be shown to be a class of dynamical systems capable *in principle* of replicating to an arbitrary degree of accuracy the behaviour of any other dynamical system with a finite number of components.

2.2 Evolutionary Robotics results

Initial experiments done at Sussex [3] used simulations of a round, two-wheeled, mobile robot with touch sensors and just two visual inputs — simulated photoreceptors, with (genetically specified) angles of acceptance, and of eccentricity relative to the frontal direction of the robot. Simple visually guided navigation tasks were evolved.

Moving from simulations to reality, we have developed a specialised piece of visuo-robotic equipment — the gantry-robot. This allows evolutionary development and evaluation of robot control systems using real noisy vision, in a controlled environment. Successful experiments in incremental evolution have used a sequence of tests of increasing difficulty, with a minimal vision system of effectively only 3 photoreceptors. The following sequence of tasks of increasing difficulty were used:

1. Forward movement.
2. Movement towards a large target.
3. Movement towards a small target.
4. Distinguishing a triangle from a square.

The successful result was a minimal control architecture (an ANN of less than 20 nodes) which could distinguish between a triangle and square with effectively just 3 photoreceptors and navigate the robot to the target.

2.3 Evolvable Hardware

The work discussed so far has generally used a genetically specified dynamical system as the control system for a simulated or real robot. But this dynamical system, for instance when conceptualised as an ANN, has in practice been implemented on a computer. There is a related approach of evolving control systems directly onto hardware, which has been taken within our group by Thompson [11, 12]. Using silicon chips such as FPGAs (Field Programmable Gate Arrays) it is possible for a designer (or in this case a GA) to reconfigure a real physical circuit embedded in silicon.

This has been used at Sussex to produce the first evolved hardware controller for an autonomous mobile robot. Also the first evolved FPGA hardware design for pattern recognition. Further experiments demonstrate that evolution may be an effective method of producing hardware designs tolerant to hardware faults, and offers a promising approach to solving some of the problems associated with achieving robustness and reliability for planetary exploration vehicles.

2.4 Comparison of simulation and reality

When using simulations it is an important question to decide just how realistic the model should be, and how noise should be handled. With these questions in mind Jakobi [9] built a simulator, *Khepsim*, for the *Khepera* robot, from EPFL, Lausanne in Switzerland [5]. The simulation is based on a spatially continuous, two dimensional model of the underlying real world physics and not on a look-up table approach as in [10]; parameters were set using empirical information from a *Khepera* robot.

Temperature Variation	60-140 K @ poles 215-295 K @ equator
Wind Velocity	20 km/hr typ, 400 km/hr max
Atmospheric Density	7mb
Particulate Environment	Dust problem, meteorites less so
Gravity	0.38 x Earth gravity

Table 1: Significant features of the Martian environment

Life (Mars day = 24.62 Earth hours)	200 Mars days
Total traverse dist.	200m
Range from lander	20m
Traverse time	50 hours
Target vehicle mass	1.55kg
Payload mass	1.45kg
Speed	0.5 to 2.0 m/h
Hibernation/Standby	20x8 Mars days
Terrainability	15 degrees up slope 20 degrees down slope 15 degrees lateral slope 100mm obstacle climbing 100mm hole crossing
Power	5-8 W
Transmission time	3 mins @ Solar opposition 22.3 mins @ Solar conjunction

Table 2: Significant mission requirements/constraints

Comparisons were made with the results from a number of evolutionary runs performed in simulation and then tested on the real robot. One conclusion is that simulations can *in some circumstances* be good enough to be used for artificial evolution, with the resulting designs successfully downloaded onto a real *Khepera*. It seems likely that there are strong limitations on how far it is realistic to extend this approach.

A second conclusion is that the noise used in such simulations should be at a level similar to that observed experimentally. If there is a significant difference in noise levels, then whole different classes of behaviours become available which, while acquiring high fitness scores in simulation, necessarily fail to work in reality. This is true both for too little noise, and for too much noise.

3. THE EFFECT OF MISSION CONSTRAINTS AND REQUIREMENTS

3.1 Environment

The robustness and reliability characteristics required of the evolved vehicle will, naturally, be shaped by the target planetary environment and by the mission constraints and requirements. Whilst complex and extensive in full, the more dominant design drivers arising from these aspects of the ROSA/M mission are summarised here in tabular form.

Of the various characteristics of the Martian environment, the thermal and particulate environments are major design drivers. Various thermal control regimes are available to counter the temperature problem, although adequate temperature control for any appendages would be difficult, particularly during periods of inactivity. Here there is the option of retracting the vehicle into its stowed shape, which would be helpful. Whether fully passive or active solutions are appropriate is uncertain at this time. Novel techniques for thermal control, such as active disconnection of critical thermal paths remain a possibility. Regarding particulate environments, micrometeorite impact can probably be discounted as a major hazard. However, ingress of dust into mechanisms remains of major concern. This is one area where the evolutionary approach advocated can bring significant benefits. Where partial loss of articulation due to dust clogging poses a significant threat, control strategies that are able to cope with modified architectures offer significant advantages.

All the various mission requirements tabulated pose their own specific difficulties. In addition to those mentioned, achieving the very tight mass budget, and the design of an adequate, environmentally friendly power supply are also particularly intractable.

Figure 1: The vehicle concept illustrated is built around the current representative ROSA/M payload of Close-Up Imager (CUI), Alpha-Proton X-Ray Spectrometer (APX), Neutron Detector (NED), APX and NED Electronics, Moessbauer Spectrometer (MOS), and Evolved Gas Analyser (EGA)

.

3.2 Trade-Offs

The dilemma faced in chosing appropriate approaches in order to achieve the required characteristics is epitomised by the trade-off between walking type robots, and wheeled robots. Many options for the mobility systems of planetary exploration devices exist, which are well documented, and not explored in detail here. Traditionally, walking robots have been unpopular because of their increased mechanism complement (hence mass), their greater demand on computing to achieve gait control, and a variety of other, valid arguments. However, there is a danger that traditional wisdom can obscure the potential of developments capable of overturning such perceptions. In the case of evolutionary robotics, this potential can be seen quite clearly, and demonstrated by established trade-off techniques.

The trade-offs which could be performed in establishing the optimum configuration / design for an instrument deployment device such as that discussed here are many and varied, but a treatment at subsystem level is considered representative at this stage. Although these devices can be considered to comprise the seven basic subsystems: Locomotion, Structure, Thermal, Navigation and Control, Telecommunications and Data Handling, Vision, and Power, it is arguable whether this is the correct approach for such a small, closely integrated device as that under discussion, since it is considered that in reality, many of the functions will be almost inseparable. However, the divisions make for clarity of treatment, and are retained for convenience. These subsystems, together with the the additional aspects of payload interfaces, launcher/lander interfaces, mass, reliability, cost, programmatics, and miscellaneous functionalities, may be considered to comprise the main aspects for consideration as trade-off criteria.

Such an assessment (wheeled vs. legged) was carried out using a “paired comparisons” based approach for the current situation, and for that expected to prevail in the near future. The trade-off criteria used were those identified above. It was found that whilst traditional wheeled approaches carry the day at the present time, the potential exists for a comprehensive reversal of the situation as the effects of new techniques come into play.

3.3 Target Configuration

A number of constraints dictate that the breadboarding phases of the programme have to be conducted with intermediate standards of hardware, which is the normal approach. Nonetheless, this breadboarding has to be set within the context of a target mission configuration in order for it to be focussed towards an identifiable set of end objectives. That this target mission configuration may itself be subject to development as the programme progresses has to be accepted as the norm. Figures 1 and 2 show this target mission baseline design, as currently defined.

3.4 Mechanisms Architecture

A major objective of the programme of work is to develop a mechanisms architecture suitable for exploiting the inherent advantages of work by the University of Sussex. This should demonstrate the required features of graceful degradation and articulation redundancy, and exhibit a good match with the silicon based evolved control system such that adequate scope exists for learned behaviour to produce the required degree of flexibility in hardware response.

A typical top-level baseline activity profile will always be substantially series: the robot moves to a site, takes

Figure 2: The vehicle stows within a volume equivalent to the declared ROSA/M stowage dimensions of 0.3m x 0.2m x 0.2m. This is achieved by providing the legs with sufficient articulation range for them to be folded tightly within the available envelope. The configuration illustrated may also have advantages as a thermal control hibernation approach.

measurements, moves to the next site, and so on, over a period of about 180 - 200 days. However, when specific activity groups, for example the deployment of the robot from its lander, are examined in detail, different results are obtained.

Traditional design philosophies are fundamentally based on an expectation of success. Even designs incorporating redundancy, employ that redundancy as a means of guaranteeing the primary functionality of the device. Thus, the behaviour of such an object, is aimed at protecting the critical (primary) path, and achieving a substantially series sequence of events - see Figure 3.

By contrast, a design can be envisaged which is intended to operate in its primary state, in a scenario filled with failure conditions. The advantage of evolutionary robotics is that it offers a way of approaching hardware design such that the output is a design intended for operation in fault conditions, not one which has a design aimed at protecting it from faults. This is a fundamental difference in the philosophical approach to design. Characteristically, this type of device will not have an identifiable primary / critical path, but a number of equally valid parallel tracks, as illustrated in Figure 3.

The flow chart (right) shows a simple logic flow for a deployment system intended to have three alternative modes of operation, reinforced with traditional redundancy at product level.

Such an approach allows flexibility for alternative deployment strategies to be evolved, and to be reinforced with additional advantages such as automatic identification of optimum power dissipation strategies, and incorporation of recharge periods as required by replanning.

In the past, such approaches have been limited by their impact on aspects such as system mass. Evolutionary robotics, however, provides the means of implementing this fundamentally different design philosophy into hardware controls capable of driving a multi-functional mechanical system such that a very large number of behaviour strategies can be employed in achieving mobility goals.

4. DEMONSTRATION STAGES

The demonstration stages of our proposed programme of work are as follows.

1. *Wheeled vehicle.* Simple wheeled vehicle on which a full control system, equipped with low resolution vision and infra-red sensing, will be demonstrated. Robustness to hardware faults and environmental uncertainty will be demonstrated.
2. *Legged Robot.* An existing octopod robot will be used to demonstrate a version of the control system able to locomote a legged robot. As well as demonstrating the criteria from the previous stage, robustness with respect to mechanical faults/failures associated with the legs will be demonstrated.
3. *Leg Breadboard.* The final stage will involve breadboarding a leg of the final target vehicle to demonstrate that the overall concept is feasible.

Figure 3: Vehicle deployment sequence: traditional design approach on the left, revised design approach on the right.

5. FUTURE WORK

The development programme described will proceed naturally into the preparation of a full Engineering Model (EM). This EM is required to fulfill two distinct purposes. Fistly, it will provide the first chance to exercise the evolved software with a fully representative locomotion subsystem. Secondly, it will provide the workhorse for investigation of the advanced techniques and new development products which will be required in order to achieve vehicle mass and functionality compliant with the mission objectives, since this still remains some way away.

The type of developments required are typified by the recent ESA ITT for Miniaturised Drive Development, which recognises the critical power/mass ratio and size problems faced by devices such as those described here. As discussed earlier, in order for the full advantages of the evolutionary approach to be capitalised upon, a mechanisms architecture must be established which is appropriate for the control strategies being developed. The EM stage of the development will be the point at which the decision to proceed along the legged vehicle line of development can be fully tested.

It also needs to be borne in mind that the vehicle in question is part of an overall mission, and it is, in effect, the reliability and robustness of this overall mission that is receiving attention. Consequently, at the EM stage, it will most likely be necessary to start representation of additional systems in order to give adequate insight into overall system behaviour. Typical of such areas would be the various lander-based deployment systems, the behaviour of which exert a strong influence on mission reliability. Alternative logic approaches for such systems were illustrated earlier.

6. CONCLUSIONS

This paper has argued for an alternative to traditional design approaches as a way of tackling the problems of ensuring robustness and reliability in the control of planetary exploration vehicles. It has outlined a programme of work aimed at exploring the use of artificial evolution for developing highly robust, yet minimal, control systems for a legged vehicle. We believe that a legged vehicle, appropriately controlled, offers advantages over wheeled vehicles of flexibility and the potential of more complex autonomous behaviours involving gripping objects and exploring surfaces locally with its legs.

Acknowledgements

This work was in part funded by EPSRC grant GR/J18125. Many thanks to Adrian Porter of MMS for generating the Euclid visualisations of the vehicle concept.

References

[1] R.D. Beer and J.C. Gallagher. Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122, 1992.

[2] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE J. Rob. Autom.*, 2:14–23, 1986.

[3] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):73–110, 1993.

[4] M. Colombetti and M. Dorigo. Learning to control an autonomouys robot by distributed genetic algorithms. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2, Proc. of 2nd Intl. Conf. on Simulation of Adaptive Behavior, SAB'92*, pages 305–312. MIT Press/Bradford Books, 1992.

[5] D. Floreano and F. Mondada. Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *From Animals to Animats 3, Proc. of 3rd Intl. Conf. on Simulation of Adaptive Behavior, SAB'94*. MIT Press/Bradford Books, 1994.

[6] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, Reading, Massachusetts, USA, 1989.

[7] Inman Harvey. Species adaptation genetic algorithms: The basis for a continuing SAGA. In *Proceedings of the First European Conference on Artificial Life*, pages 346–354. MIT Press/Bradford Books, Cambridge, MA, 1992.

[8] Inman Harvey. Evolutionary robotics and SAGA: the case for hill crawling and tournament selection. In C. Langton, editor, *Artificial Life III*, pages 299–326. Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Vol. XVI, Addison-Wesley, Redwood City CA, 1994.

[9] N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life*, pages 704–720. Springer-Verlag, Lecture Notes in Artificial Intelligence 929, 1995.

[10] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 190–197. MIT Press/Bradford Books, 1994.

[11] A. Thompson. Evolving electronic robot controllers that exploit hardware resources. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life*, pages 640–656. Springer-Verlag, Lecture Notes in Artificial Intelligence 929, 1995.

[12] A. Thompson, I. Harvey, and P. Husbands. Unconstrained evolution and hard consequences. In E. Sanchez and M. Tomassini, editors, *Towards Evolvable Hardware.* Springer-Verlag, 1996.