An r-strategy Architecture for the Robotic Exploration of Mars

Duncan Steel[1] and Harrison Steel[2]
[1]Australian Centre for Astrobiology, University of New South Wales
[2]Wesley College, University of Sydney
E-mail: TMA1@grapevine.com.au

Abstract

Biologists often divide the reproductive strategies of living organisms into two broad camps: K-selectors and r-selectors. K-selection leads to high efficiency (the quality and the probability of survival of individual offspring) while r-selection leads to high productivity (the number of offspring produced). Prime examples of the former include humans, whales and elephants, as characterized by a large investment in a very limited number of progeny, of whom great care is taken as they mature. Conversely, r-selectors produce huge numbers of offspring and may abandon them to their own devices; examples include weeds, insects, and most fish. There is a spectrum between the extremes of r- and K-selection (e.g. small rodents; various birds, amphibians and reptiles).

Generally, r-strategists are better suited for an unpredictable environment: they tend to be able to occupy a new habitat rapidly, but require nutrients to be readily available. With high motility and large numbers, there is a greater chance that at least some of the offspring will survive. In K-strategists the parents focus on passing down memes (behavioural traits or cultural information) that benefit survival and future proliferation, whereas in r-strategists the future behaviour of progeny is defined by their genetic inheritance and physical environment. The young of r-strategists are independent and self-sufficient (i.e. precocial), whereas the progeny of K-strategists are dependent on external support and are helpless in themselves (i.e. altrical). In an unstable or dangerous environment r-selection is more robust, leading to a higher probability of avoiding complete mortality of the species.

To date our attempts at the surface exploration of Mars have clearly been based on a K-strategy: one or perhaps two landers/rovers, in which a huge amount is invested and consequently the least-risky (and therefore least-interesting) landing sites are chosen. We argue here that if we want to take a closer look at the more interesting locations on Mars then we should adopt instead an r-strategy. That is, we should employ an architecture based on landing/scattering many small, self-sufficient robots, each of which has modest capabilities and is individually expendable.

Imagine a swarm of a hundred microbots − each perhaps the size of a tennis ball or a football – being dropped onto Mars, with individual probabilities of success of just 20 percent. The probability that <u>none</u> would be successful is just 2×10^{-10}. We'd probably settle for 10 being successful, but the expectation value would be 20 if we'd done our risk analysis correctly; and we might get 25−30 or even more successes!

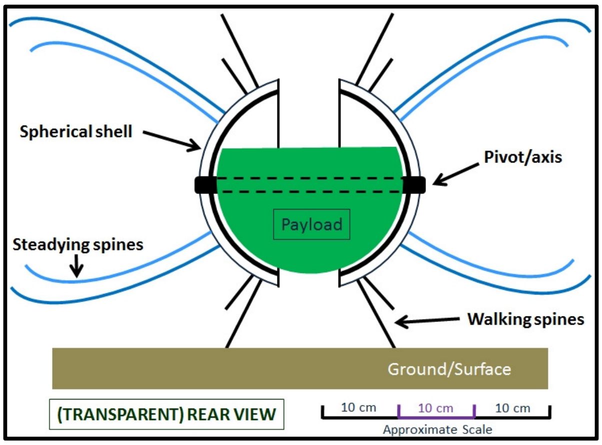

In this paper we will present a discussion of the above concepts, but our main aim is to stimulate discussion by others, and the generation of new ideas.