


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Acta Astronautica

journal homepage: www.elsevier.com/locate/actaastro

Future robotic exploration using honeybee search strategy: Example search for caves on Mars

Áron Kisdi^{a,*}, Adrian R.L. Tatnall^b

^a University of Southampton, 36 Priory Road, Wantage, OX12 9DD, UK

^b SES, University of Southampton, SO17 1BJ, UK

ARTICLE INFO

Article history:

Received 19 February 2010

Received in revised form

4 January 2011

Accepted 24 January 2011

Available online 3 March 2011

Keywords:

Biomimicry

Swarm

Behaviour

Bees

Mars

Search

AI

Autonomous

Caves

ABSTRACT

Autonomous control has an increasing role in Earth and Space based applications. High level autonomy can greatly improve planetary exploration and is, in many cases, essential. It has been suggested during the Mars cave exploration programme, that an effective way to explore a larger surface area would be the use of many, small and fully autonomous robots. However, there are many challenges to overcome if such a swarm exploration programme is to be implemented. This paper summarises these challenges and focuses on one of the most crucial one: strategy. Many effective group exploration behaviours can be observed in nature, most of which are optimised to work with agents that have limited capabilities as individuals. For this paper a computer program has been written to simulate the way bees search for new hives and investigate whenever it is an optimal method to search for cave entrances on Mars. It has been found that this method, using simple autonomous robots which can be constructed using available technologies, could greatly improve the speed and range of a planetary exploration mission. The simulation results show that 50 swarm robots can cover an area of over 300 meters square completely in 5 sols while they are searching for cave entrances and returning results to the Lander which is a major performance improvement on any previous mission. Furthermore areas of interests found by the explorers are sorted in order of importance automatically and without the need of computational analysis, hence larger quantities of data were collected from the more important areas. Therefore the system – just like a hive of bees – can make a complex decision easily and quickly to find the place which matches the required criteria best. Using a high performance search strategy such as the one described in this paper is crucial if we plan to search for important resources or even life on Mars and other bodies in the solar system.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

1.1. Aim of research

Autonomous systems are increasingly important in all engineering applications. In the case of space systems

Artificial Intelligence (AI) is highly desirable as agents operate in a remote environment and there is often a delay in communications and sometimes a fast response is required. However AI requires major improvements in the space sector for future missions and there is much to be learned from nature in this field. Advanced autonomy can make systems more robust, speed up planetary exploration and would also reduce the cost of operation. Even the simplest biological systems exhibit very effective decision making and sometimes cooperation. The individual agents of the swarm have limited capability but the

* Corresponding author. Tel. +07901871711.

E-mail addresses: aron.kisdi@stfc.ac.uk (Á. Kisdi), art4@soton.ac.uk (A.R. Tatnall).

decisions they make collectively are complex, while the process is efficient. Inspired by the fascinating way bees search for new hives, this paper considers the use of a similar artificially simulated methods in order to find areas of interest in planetary exploration. A secondary aim is to use biomimetics in such a way that the complexity of AI is minimal, allowing hardware designs with minimal processing power, meaning that compact and robust hardware can be built.

1.2. Honeybees in nature

There are more species of bees than mammals and birds combined [1]. The species are very different both in appearance and in behaviour. Contrary to common belief most bees are loners and only a few species live in hives. In this report the subject of interest is the *apis mellifera* or European honeybee.

1.2.1. When bees search for a place for a new hive

If the bee hive reaches a certain size, explorer bees set off to locate a place for a new hive. If the explorers find a place which they believe is suitable they return and 'tell' some of the other bees the location. These bees then set off to inspect the location. If they like the place they return to the hive and they recruit more bees to inspect the location. In this way if a place is not so good for a hive fewer and fewer bees will visit it. However if it is a good location more and more bees will go there until the number of bees reaches a certain size and the original hive splits and starts the construction of the new hive [2].

This search is interesting because of many aspects. The hive as a collective will receive multiple opinions on each possible location. These opinions will come from different bees who visited the location at different times. At the end of the search the best location is selected however any other location can be easily compared with one another. Like all biological agents, bees are not identical, their physical and mental state is defined by their growth and learning. On the other hand the simulated bees are identical just like the mass produced swarm robots would be.

1.2.2. Advantages of bee-like search

Now we have to study the advantages of a bee-like search. The hive as a collective receives different reviews of each place and compares them before the decision is made. But the review of a bee is not subjective; it searches for the characteristics that are best for the new hive. Each bee has the same senses but a place can receive different reviews because of a number of different factors. Partially it is down to individual opinion which is impossible to simulate but there are other aspects that need to be considered; most importantly time and perspective. Bees will visit the same place at different times therefore collectively gather much more information about the changes at the place over time, such as during a day-night period. Differentiating perspective can be achieved in multiple ways, one of which is equipping the robots with different sensors. Having task specific agents could be interesting for a future study however the aim of this

study is to show that improvement can be made on past missions even with the simplest agents if the bee-like search is used. Even if the agents are identical it is possible to increase the information gathered from revisits by making the robots visit the same place from different directions allowing them to see additional detail which might be hidden from other angles. This can be achieved by randomizing obstacle avoidance which will lead to different paths to the same point. Alternatively, the AI of the robots can be improved so that they approach the previously discovered point from a different direction by choosing another path when they are close to the target. Both of these improvements can be added to the simulation however to change the results a higher resolution model of the terrain and caves should be used as differences will depend on the temperature profile of the caves.

The individual bees make the decision whenever they think the location is good or not based their senses. It is debatable if bees have individual preferences or they search for characteristics required for the hive. Therefore these biological agents might or might not introduce bias in the selection however the robots based on classical AI only compare a limited number of digital inputs therefore in the simulation the agents are identical and without bias which does not affect the overall strategy. However it is interesting to note that such a system can potentially manage bias well. Inaccuracies from sensors would only affect individual agents on the other hand from the perspective of the hive it is all down to numbers. A location is better if more explorers returned from it with a positive opinion.

In summary the advantages without running the simulation to compare the search performance to past missions:

- Identical agents improve swarm robustness and reduce production costs
- Possible areas of interests are revisited by other agents producing data over time
- Revisits could potentially show the area of interest from different perspectives
- Revisit efforts are concentrated on areas which are cumulatively found relevant
- The areas likely to be relevant can be identified autonomously with very simple reactive AI

The simulation discussed in this study was set up to find out if such a search strategy can be artificially replicated and also to quantify the performance of swarms of different sizes and compare its speed with past missions, and to identify any disadvantages.

1.3. The importance of Mars exploration; caves and methane

The surface of Mars is hostile to any organisms mainly because of its very low temperature and high radiation levels. Caves would not only provide shelter from radiation but could also have a warmer and more even

temperature. A model presented in a recently published article [3] implies that liquid water is present under the surface of Olympus Mons. Deep reservoirs are warmed up by geothermal gradients and magmatic heat meaning that even in the present day could contain liquid water. Atmospheric observations from both Earth and Mars orbiting spacecraft confirm the presence of methane in the Martian atmosphere [4]. Methane is quickly destroyed due to radiation. This means that it is continuously generated under the surface. The source could be biological or geological, however, either way; it indicates the presence of water underground.

Images from Mars Odyssey spacecraft revealed deep holes on the surface which are likely to be cave entrances [5]. The reason for this is that these holes are colder during the day and warmer during the night than the average surface temperature. The temperature measurements were taken by Mars Odyssey's Thermal Emission Imaging System camera in the infrared spectrum. However the resolution of these observations from orbit is limited. The smallest possible cave entrance of the seven discovered at the time this report was written is around 50 meters diameter while the largest is 200 meters. These holes have deep vertical walls which make them very difficult to explore with rovers; however, there could be much smaller and much more accessible cave entrances on Mars.

1.4. How can caves be found?

Large holes on the Martian surface were first spotted by the Mars Odyssey spacecraft. It took infrared images during the day and night of the same area. These images can be seen in Fig. 1 clearly show that the temperature is warmer inside the hole during the night and slightly colder during the day than the surface temperature. This characteristic is very similar to what can be observed at cave entrances on Earth, however due to the low atmospheric pressure and wide temperature fluctuation the signal strength is expected to be stronger than on Earth [6].

Therefore, the best method to search for potential caves is to measure the temperature and compare it to the expected average surface temperature or a natural point on the surface. Measurements of this temperature difference of the same region at different times over the day–night cycle can be used to confirm whenever something is a cave entrance or just a crater. Recent studies led

by SETI help us to understand the thermal behaviour of Martian caves [6].

The average surface pressure is around 0.6 per cent of the surface pressure on Earth. The average at zero altitude is around 0.6 kPa. [7] Therefore, the atmosphere is much steadier than on Earth and weather features are less intense and more widely spread. This enables altitude to be calculated by measuring the pressure and using atmospheric models available for Mars. This can help in identifying caves.

The change in pressure with altitude is given by Eq. (1)

$$P_{lan} = e^{-\frac{z-z_{ref}}{H}} P \quad (1)$$

where P_{lan} is the mean surface pressure at the landing site while P is the mean surface temperature at Viking 1 landing site at the same time in a Martian year. The equation $z-z_{ref}$ gives the altitude relative to the Mars reference surface. H is a constant which can be calculated for the lower few kilometres of the Martian atmosphere using Eq. (2)

$$H = R \frac{T}{mg} \quad (2)$$

where R is the universal gas constant, T is the mean temperature at the altitude, g is the local acceleration due to gravity and m is the mean atmospheric molecular weight.

Therefore, to measure altitude to one meter accuracy a pressure sensor capable of measuring one twentieth of a Pascal would be required coupled with a temperature sensor because pressure varies with temperature. The atmosphere of Mars is stable enough to use this method for local altitude approximation so small robots can better track their path and correct odometry.

As caves shield sunlight and radiation they could be identified by measuring light and radiation intensity during the day. However using equipment to measure these characteristics would require the robot to enter the cave to some extent which might be impossible in cases where the wall of the cave pit is vertical.

All these parameters can be measured using low power, mass and compact integrated circuit or MEMS based sensors. The search strategy described in the following section is aimed at agents with similarly simple digital or analogue sensors operating in a known range.




Fig. 1. From left to right: visual image, infrared image in the afternoon and during night—Credit: NASA/JPL.

2. Simulation programme

2.1. Objective

The simulation was set up using the Python programming language. One of the objectives was to keep the simulation flexible to modification and development. The program of the graphical interface was kept separate from the description of the Martian environment and the code which controls the bees. Therefore the different aspects can be developed separately and different codes can be tested without the need to rewrite anything. The software is open source and available for download free of charge [8].

2.2. Overview of the simulator

In the following sections all parts of the program are discussed in detail: the scientific background and the included features and theories behind the working mechanism.

2.2.1. Environment description

The simulation is two dimensional however at a later development stage this was extended to pseudo three dimensional. Altitude is included but only the surface is simulated. The area is limited to 1 kilometre square. Multiple simulations can be run to test for larger areas however, the 1 kilometre square is sufficient to show the concepts and to compare different scenarios in order to optimise the number of robots and the robot code. The simulated area is shown in two separate screens. The main screen shows the 100 meter square with the Lander in the centre. This has a resolution of 0.1 meter square. This means there is a different value for each environmental aspect for every 0.1 meter square. Objects which are smaller than this are not considered and it is assumed that a robot can cover 0.1 meter square with its on board sensors. The second screen shows more of the surrounding up to one kilometre square and it has a resolution of 2 meter square.

The main tool used to look for caves is the observation of temperature. As on Earth, caves on Mars have a more even temperature therefore the cave entrances are colder during the day and warmer during the night [6]. As the goal is to find the temperature difference between the cave entrance and a neutral point on the surface regardless of the time of the day, the software is simplified and does not include day–night cycles.

Temperature values were approximated to give similar results to those given in the Mars Environment Specification document by NASA's Jet Propulsion Laboratory [9]. The temperature is included for every point of the surface – within the resolution – in a two dimensional matrix.

Pressure, light and radiation intensity and other characteristics are not fully implemented in the simulation at the time of the publication of this paper.

Possible caves can be generated at random or can be user defined. Caves have a size and notable difference in temperature and some difference in pressure compare to other points of the surface. The larger the difference the

more likely it is to be a cave. Robots will locate all possible entrances however the aim of the programme is that the swarm compares the locations and decides which one is most likely to be a cave. Instruments on orbit are limited by their resolution and viewing angle. Robots can be used to detect smaller and more accessible caves. In order to make it possible to compare different runs of the programme the caves were defined manually instead of being generated randomly, a capability which is also included in the program.

The environment file can be developed separately from the rest of the program. A major improvement in the future would be to include changes due to the day–night cycle and weather conditions hence making the environment change dynamically with time.

2.2.2. Robot movements simulation

The way the robots explore their surrounding, avoid obstacles or return to the Lander was aimed to mimic the behaviour of the bees in the simplest possible way. A random walk pattern was tailored specifically for the simulator. Other basic decisions – for example which way the robot should turn if it hits an obstacle – are made at random.

2.2.3. Time, outputs and user interface gives the mainframe of the software

The main design goal of the interface was to be graphical and user friendly so simulations are easy to set up and the behaviour of the robots can be observed before the final data is returned. This way the user can concentrate on the development of the robot code and can get graphical feedback identifying errors immediately.

When the program starts it asks for some user input. The type of robot control needs to be specified first. The options are controlled or behaviour based robots. If behaviour based is selected the program asks for the number of robots. In theory the software can run with any number of robots however it is optimised for less than one thousand. The last parameter the user needs to select is the background. The user can choose between temperature and pressure display in the current version.

The programme displays three windows. The first window shows the close surrounding of the Lander, a 100 meter square zone while the second window shows a 1 kilometre square area. Both screens have various mouse and keyboard commands. The third window displays the time, the Lander's memory and certain events such as when a robot finds a cave. Time is displayed in two different ways. On top the simulation time is displayed in seconds and under it the Martian time. Martian time is displayed in minutes, hours and sols and one sol is 24 h and 39 min.

The average speed of the MER rovers is about 3 meters per minute. The maximum speed of the robots is also included in this code and it is assumed to be 0.1 meter per minute. The reason this value is so low is because the small size of the robots and the assumption that they avoid or climb over small obstacles, stop to take measurements or turn around. It assumes the worst case, on which a real life scenario should only improve.

3. Robot control

The most important attribute of an agent working in a remote environment is reliability. Generally a simple system can be made more robust than a complex system. This applies to both hardware and software design. Therefore the main aim of the robot control was to achieve the required level of autonomy with the simplest possible code. This can be achieved using subsumption architecture [10], meaning decomposing intelligent behaviour into many simple modules.

At the start of the program the robots are in the Lander and set off to explore their surrounding in a random manner. While one might argue a more structured search could prevent repetition and ensure there are no areas missed, such a strategy would need advanced AI or collaboration between agents, hence more complex robots. The results in this report show that a simple random walk can still improve greatly on previous planetary exploration missions while it gives the option to use very simple and therefore robust hardware. Reducing the hardware requirement also means agents can be made smaller or lighter allowing more advanced robot designs such as jumping or flying locomotion. They continuously read the on board sensors and record the distance they have travelled from the Lander. They avoid obstacles and have a certain chance of returning to the Lander if they have not found anything of interest. This chance is time dependent and described by the Poisson distribution. This model mimics how bees return to the hive after some time if the search was unsuccessful. If a robot finds an area which might be a cave entrance it returns to the Lander immediately. Whenever a robot identifies something as a cave or not depends on certain threshold values which can be changed. In the simulation the temperature difference threshold was set to 5 degrees. It is assumed that the robots continuously receive temperature measurements from the Lander wirelessly and compare them to its own reading. The results show that a relatively fast decision can be made without the need of advanced sensors such as cameras. Again this design aims to show that the strategy can be used with very simple agents. If the Lander is mobile the swarm of “bees” can identify the most interesting destination rapidly in order to maximise the data of interest returned by the Lander’s sensor array which could include cameras and other high bandwidth devices.

Each robot stores the values of the last sensor measurements and does not store any new values on the way back, unless it finds a ‘better’ – i.e. has a larger difference than a neutral point on the surface – cave on the way back. The robots take the shortest route, the ‘bee line’ back to the Lander. As the robot keeps track of the distance it travelled in a Cartesian frame it can calculate the direction of the Lander and move directly back instead of going back the way it came. For a real mission another possibility to explore would be a continuous homing signal sent from the Lander.

If a robot returns to the Lander it checks the memory of the Lander where all the discoveries are stored. Again the lack of wireless transmitter on board of every agent

minimalizes the hardware requirement. If it found a cave entrance which is better than any in the Lander’s memory, it uploads the data to the Lander which then sorts the list of stored data. This way the Lander acts as the hive of bees which can communicate past findings to any explorer bees which returns. If some possible cave sites are already discovered, the robot decides whenever to start a new search or go to one of the possible caves and take measurements again. A better cave has a higher chance to be revisited.

This way the robots interact with each other by the use of stigmergy [11]. Stigmergy means that there is no direct interaction between agents, yet they are able to communicate through a third party. They modify their environment, in this case the hive (Lander), and other bees (robots) are then able to detect this change therefore receiving information from the other agents.

Another advantage of such a simple control is that in theory the robots could be built purely on analogue technology without microprocessors. An analogue system would be much more robust than digital technology as it would not be vulnerable to radiation and could tolerate a much wider temperature range. Also analogue systems are much faster and are therefore able to respond to changes such as the sudden appearance of a cliff or cave entrance. On the other hand digital control would enable the robots to be improved further and could significantly increase the amount of data returned. Therefore both options should be considered in future investigation of the mission.

4. Results and discussion

To be able to quantify the findings from the simulator, some numerical results have to be recorded. The simulation can be set up in many different ways. For this section several set ups were used in order to collect comparable data and try to answer the main questions regarding the bee search algorithm:

- Were the ‘better’ cave entrances visited by more robots than others?
- Is there a relationship between the size of the swarm and its effectiveness?
- Is there an optimal number of robots?
- Can the program be used to find the optimal number of robots?
- How long does it take to find all caves in a certain area?
- How large an area can the robots explore in a certain period of time?
- How do the findings compare with existing missions?

The variables which can be changed without modifying the software:

- Number of robots.
- Number of caves, with the choice of random cave generation or manual set up.
- The number of times the program should run subsequently.

- Time of simulation with the option not to have a time limit i.e. the simulation runs forever. (Time in terms of simulated Martian time.)
- Whenever the simulation should be stopped after a number of cave entrances have been found.

4.1. Comparing results with bee behaviour in nature

The program was designed to mimic the behaviour of bees in nature in order to distinguish important and less important results and concentrate on retrieving more data from the important locations. All the caves which were found by the robots are included in the results as well as the number of times a robot visited that cave (number of 'hits on cave' in results). The cave with the lowest temperature, hence highest temperature difference from the Lander should be the one visited most times; while other cave entrances should have considerably less visits. More visits in nature removes potential bias, and similarly in case of artificial agents if a robot's sensor malfunctions or has high level of noise the result from the system is not effected. Due to the random walk of the agents multiple simulation results have to be averaged to draw an acceptable conclusion.

4.1.1. Results from 100 simulation runs

Data of over 100 simulation runs is summarised in Table 1. The average of the simulation data is shown for a certain number of robots to one decimal place.

The results show an increase in both the total number of caves found and the total number of times robots visited a possible cave site with increasing number of robots, while the time it took the robots to find the first cave decreased. An interesting exception is the average of the simulations with 100 robots where the results show a decrease in performance. This can be regarded as an error due to the randomness of the simulation, however further investigation should be carried out for large swarm sizes. The standard deviation for number of caves found is approximately 1.

The most important part of the results regarding the bee-like behaviour is the last three columns, which show

the number of times a robot investigated one of the best three potential cave sites and returned with a result. Fig. 2 shows data from another set of experiments plotted on a scatter graph. It can be seen that while the number of visits increases for all caves, the rate of increase is larger for the better (higher temperature difference for this simulation) caves. Therefore, the software successfully simulates a swarm of bees as it utilises more resources to investigate better caves.

From the trend lines in Fig. 2 it can be seen that the cave which fitted the criteria best always got the most visits from robot agents while the second best option got considerably less visits on average. The linear regression is shown as R^2 . Because of the variation due to the randomness of the search, the second best site sometimes got less hits than the third best, but it is also because of the time limit of 10 sols. If the search is run for an extended period of time the possible cave entrances which better satisfy the criteria conditions will always receive more visits.

Unlike the simulation, in the real world, the sites might change over time and their fit to the criteria could become better or worse. In theory the search is capable of dealing with such a situation and the cave which is best most of the time will receive most attention from the explorer robots. This could be tested by introducing environmental changes over time in the simulator.

While the number of times robots visited a cave show a linearly increasing relationship with the number of robots, if we consider the total number of caves found the relationship is not so clear. In this simulation scenario 10 caves were hidden on the 1 kilometre square area and the allowed search time was 10 sols. The results are shown in Fig. 3. The number of caves found increases with the number of robots however the results better suit a logarithmic relationship. Therefore increasing the size of swarm and hence the mass of payload will only increase the effectiveness by a small amount after a certain size. This means there is an optimal number of robots for a certain scenario and the program can be used to find this optimum number or at least a range where the trade-offs should be considered.

The logarithmic relationship suits the results best with a coefficient of determination (square of correlation coefficient) value over 0.9. Further increase in the swarm

Table 1
10 caves, 10 sol limit, variable number of robots. Average of 10 run each (overall 100 runs).

Number of robots	Number of hits	Time it took to find first cave				Avg. no cave	Hits on best caves		
		Sol	Hour	Minute	Sum in Sol		First	Second	Third
10	14	5.4	7.7	23	5.7	2.2	9.8	3.6	0.4
20	61	2.1	7.8	34.6	2.4	3.1	30.1	23.2	7
30	80	2.3	11.3	37.1	2.8	3.6	44.2	33.3	1.7
40	79	1	11	24.4	1.5	4	42	27.3	7.9
50	140	0.5	11.8	33.6	1.0	3.4	115	15.8	8
60	190	0.4	13.5	33.8	1.0	4.3	95.7	60.3	28.8
70	191	0.2	11.2	21.1	0.7	4.4	88.4	70.7	26.6
80	240	0	10	55	0.4	5	18	179	37
90	524	0	14	33	0.6	5	461	22	2
100	228	0.2	12.4	32.1	0.7	4.6	98.7	96.6	28




Fig. 2. Number of times a robot found (visited) the best, second best and third best cave against number of robots. For this simulation higher temperature difference is better.




Fig. 3. Number of caves found against number of robots; 10 caves 10 sols. The logarithmic curve is a good approximation for relatively small (less than 500 agents) swarms.

size would still results in increase in number of caves found however large numbers of additional robots are required to increase it with a small amount. This suggests that if higher accuracy is required the search time should be increased rather than number of robots. If there are a

finite number of caves (like in these simulations) to find all caves within a relatively short time such as 10 sols a very large number of agents would be required. While the simulation can be used to find from the available data the prediction is that this number it would be more than 200

which would result in an impractically large swarm. Hence a search with a reasonable number of agents is in most cases incomplete.

4.1.2. Effects of changes in the size of the swarm

To determine the relationship between the number of robots in the swarm and its effectiveness, a single scenario was set up and results from multiple simulation runs observed. Five cave entrances were added in the 100 meter square area around the Lander and five more further than 50 meters but closer than 500 meters away. The simulation was run for 5 sols, ten times with a certain number of robots. The minimum size of the swarm was ten robots and it was increased by five robots each time up to 70.

Table 2 summarises the results by averaging the outcomes from the 10 separate runs for each number of robots. The simulations ran for 5 sols each time and the position and size of the caves were fixed. The first column shows the number of robots, the second column the total number of times a robot discovered or visited a cave, the third column the number of caves that were discovered and the last column the 'effective radius' of the search. This means that over 80 per cent of the caves had been found within the radius.

As discussed in Section 4.1.1 the amount of data returned from potential cave sites increases with the size of the swarm. This suggests there is always a benefit in increasing the swarm size. However for a given mass budget, an increase in swarm size would reduce the mass assigned to a single robot and therefore the capability of the explorers. The relationship between the number of caves found and the number of robots also show a non-linear trend. For the data obtained for five sols runs the trend line which fits best is a logarithmic curve. As discussed previously it is the same for ten sols run the logarithmic curve was a best fit. Whilst determination of a general equation would require many more simulation runs some conclusions can be drawn from these results:

- There is a strong correlation between the number of robots and the science returned.
- The relationship between the number of caves found and the number of robots is nonlinear.

Table 2

Results from 130 simulation runs with a 5 sol time limit.

Number of robots	Total number of hits	Number of caves found	Average distance (m)
10	5.5	1.7	267.9
15	7.7	2.3	269.85
20	8.9	1.9	263.15
25	15.6	2.5	279.4
30	15.7	2.8	322.7
35	15	2.7	288.85
40	21	2.9	304.85
45	23.1	3	321.45
50	28.6	3	314.05
55	25.1	3	346.9
60	30.9	3.2	315.85
65	31.6	3	312.25
70	32.5	3.2	352.2

- After a specific number of robots any further increase of swarm will only improve the performance by a small amount—this number can be found by multiple simulation runs.
- The simulator can be used to determine the optimal number of robots for a certain mission.

4.1.3. Time to find caves and search a certain area

Sometimes finding results fast is the key point of a mission. If the mission lifetime is limited in a hostile environment it is desirable that the surface exploration returns scientific results as soon as possible. A swarm of robots could be used to locate interesting areas rapidly. Fig. 4 show the time it takes for a certain swarm size to find the first cave. Increasing the number of robots decreases the time it takes to find the first area of interest.

The results depend on the position of the caves and it was only calculated for a specific case. If the approximate proximity of area of interest is known and time is a major factor, the simulator program can be used to calculate the optimal number of robots. In this case the trade-off region is between 60 and 80 robots.

For some missions it might be required to search an area and find all or most of the places of interest. Fig. 5 shows how well the swarm performs in 5 sols. The distance on the y-axis is the average distance the robots returned results from. Because of the nature of the swarm there might be important areas closer than the given distance however it is unlikely. On the other hand there will be robots returning results or exploring at greater distance.

The data shows a liner trend, as expected in case of maximum distance travelled. However, the speed of the robots is a limiting factor and increasing allowed search time would also increase the area searched.

If a time limit is set the swarm could explore the surface and based on the results autonomously evaluate the best solution for a movement of a larger and more capable rover. The accuracy of the solution will depend on the swarm size, the overall search time and the number of the factors the swarm robots needs to look for.

4.2. Comparing results with other Mars surface exploration missions

The most successful surface exploration missions to Mars were the two MER rovers. Opportunity, which travelled the longer distance, only moved 3 kilometres a year on average. This limit in speed is mainly due to the limited autonomous control even though 100 meters a day was reached later in the mission. On the other hand a swarm of small robots would navigate on the surface and move around the obstacles for which simple behavioural AI can be used. Because the payload and control required for the bee-like search strategy can be accomplished using very simple hardware it opens up the possibility for small, innovative robot designs and locomotion methods which are not possible on heavy robots. Small swarm robots can combine rolling with jumping, gliding or powered flight to reach better obstacle avoidance capabilities. The swarm




Fig. 4. Time to find first cave against number of robots based on Table.

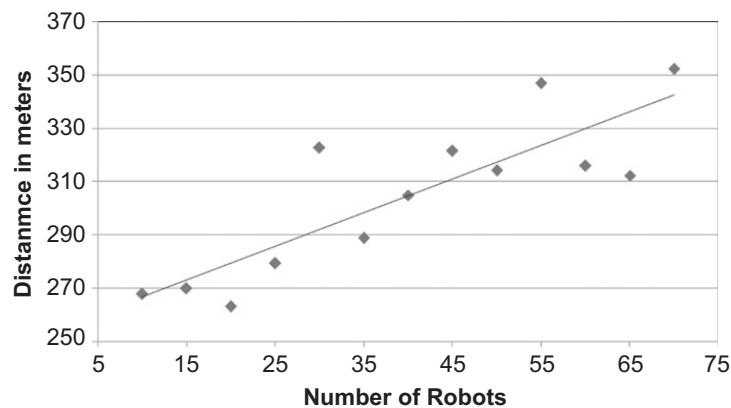


Fig. 5. Radius of effective search is against number of robots based on Table.

could gather important results even with the loss of several robots. Even though the simulation assumes simple and – considering current technology – slow swarm robots, the results show, with a sufficiently large swarm size they are capable of searching a 1 kilometre square area in 10 sols. From work on robots with similar capability at the University of Bath [12], if we assume 3 kg per robot over 50 robots could be delivered to Martian surface for the same mass budget as one of the MER rovers. However this strategy could allow the use of even smaller designs because of the minimal required processing power and small sensor mass. In the future sub-kilogram robots could be built based on behaviour AI.

Future missions will make more use of artificial intelligence and autonomous control. The ExoMars mission aims to achieve a speed of 125 meter per sols. Assuming it roves without a stop it could move 625 meter

in a straight line in 5 days. The simulation results show that 50 swarm robots can cover an area of over 300 meters square in 5 sols while they are searching for cave entrances and returning results to the Lander. Considering real life robots would likely to be able to move faster than the maximum velocity assumed for the simulator this could even be improved further.

In conclusion using behaviour based artificial intelligence and a swarm of robots to search the surface has the potential to considerably increase the searchable area and return results faster than previous missions or near future single rover missions.

5. Conclusion

The simulation was set up with multiple runs to collect data. As discussed, the analysis of the data shows that bee-like behaviour can be observed in the simple robot

control algorithm. The areas of interest found by the explorers are sorted in order of importance by the number of robot visits and hence larger quantities of data were collected from the more important areas.

The simulation also shows that surface exploration mission with multiple fully autonomous explorers are potentially much more efficient than previous and near future rover missions. The main advantages are the speed and the increased size of the searchable area, but at a cost of some error.

The simulator software can be extended in several ways. The Martian environment can be described in more detail introducing time element and true three dimensions and the resolution can be improved to include details of obstacles. Different environmental files can be added to allow simulations of the Earth, Moon and other possible places for swarm missions.

The next step would be to design and build swarm robots and test the search strategy in the real world instead of a simulated or controlled environment. This would help to understand the feasibility, efficiency and the minimal possible hardware design better.

Furthermore other use of the bee-like search can be investigated for applications such as communication networks, scientific missions or robotic search and rescue.

References

- [1] vanEngelsdorp, Dennis. Where have the bees gone? July 2008. <http://www.ted.com/index.php/talks/dennis_vanengelsdorp_a_plea_for_bees.html> (accessed 26 March 2009).
- [2] P. Miller, The Genius of Swarms, National Geographic Society, 2007.
- [3] Patrick J. McGovern, K. Julia, Morgan. *Volcanic spreading and lateral variations in the structure of Olympus Mons*, Mars., Geology (The Geological Society of America (2009) 139–142.
- [4] Steigerwald, Bill. Martian Methane Reveals the Red Planet is not a Dead Planet. 15 January 2009. <http://www.nasa.gov/mission_pages/mars/news/marsmethane.html> (accessed 08 March 2009).
- [5] Webster, Guy, and Dwayne Brown. NASA. 21 September 2007. <http://www.nasa.gov/mission_pages/odyssey/odyssey-20070921.html> (accessed 6 December 2008).
- [6] Judson J. Wynne, et al., Earth-Mars Cave Detection Program, Phase 2, The Explorers Club, New York, 2008.
- [7] Glen Research Center. AtmosModeler Simulator. 17 July 2008. <<http://www.grc.nasa.gov/WWW/K-12/airplane/atmosi.html>> (accessed 31 March 2009).
- [8] Áron Kisdi, Bee search simulator. 30 April 2009. <<http://cid-440e9ad70737c878.skydrive.live.com/self.aspx/Biomimetics/simulator.zip>> (accessed 28 September 2009).
- [9] M. Desjean, F. Forget, C. Farguson. Mars Sample Return Mars Premier. Report, CNES and JPL, 2001.
- [10] Rodney A. Brooks, Anita M. Flynn, Fast, cheap and out of control: a robot invasion of the Solar System, Journal of the British Interplanetary Society 42 (1989) 478–485.
- [11] Mark Ayre, Ian Nicholas, Biomimetic Engineering for Space Application. Vols. SP-1297, ESA Publications Division, Noordwijk, 2006.
- [12] Rhodri Huw Armour, A biologically inspired jumping and rolling robot. PhD Thesis, Bath: University of Bath, 2008.