

Quantitative Evaluation of Human-Robot Options for Maintenance Tasks during Analogue Surface Operations

Graham A. Mann

School of Information Technology,
Murdoch University
South Street, Murdoch WA 6150
g.mann@murdoch.edu.au

Abstract. Due to the scarcity of human labour plus the harsh conditions at any human Mars base of the foreseeable future, robots are likely to be employed in to assist with at least some assembly, deployment, transportation, inspection, servicing or repair tasks. By the first human landing, robotic technology is expected to have made possible the use of robot teams already on the surface to prepare the landing site, ensure the functioning of ISRU equipment and survey the local area for the arriving astronauts. Robots are also likely to assist them during their stay and after their departure. Today's researchers are increasingly interested in the question of how to systematically choose the best combination of robots and/or humans for particular tasks, and how to actually demonstrate and measure teams performing these tasks in realistic simulations. This paper critically examines a quantitative method developed by Roderiguez and Weisbin of JPL for computing performance/resource scores for a range of human-machine systems on a variety of tasks. It then proposes a practical experiment, to be conducted at a future Mars Society surface operations simulation, that will apply the method to quantitatively compare human maintenance task scores with those of a hexapodal service robot that the author is currently building.

Keywords: Mars analogue studies, field robotics, evaluation, sliding automation.

1 Introduction

It is now commonly envisaged that due to the scarcity of human labour and the harsh conditions at any human Mars base in the foreseeable future, teams of robots will assist with assembly, deployment, transportation, inspection, monitoring, maintenance, mapping, science and safety tasks. By the time humans are ready to land, perhaps around 2020, the technology is expected to have advanced to the point where robot teams already working on the surface will prepare the way for the arriving astronauts - checking equipment, surveying the site, moving boulders, etc. Later, robots deployed on the surface will join a cooperative network, communicating with and working alongside the humans in flexible ways to get the best out of both agencies.

Mission planners and engineers are increasingly interested in the question of how to choose the best teams of robots and/or humans for particular tasks, and how to actually demonstrate and measure teams performing these tasks in realistic simulations. Human teams are constantly being evaluated in ever more realistic surface simulations, involving sophisticated electronic communications, planning, recording and monitoring. For example, a team lead by William Clancey demonstrated the value of their Mobile Agent Architecture at the Mars Desert Research Station (MDRS) in Utah in April, 2003 [1]. Once the necessary physical communications infrastructure had been made reliable, this voice-driven software proved capable of acting as an intelligent 'CapCom', automatically monitoring, route planning and generally assisting its human agents, transferring large volumes of logged data such as maps, models, photographs, voice logs and other science data around the local area, and communicating this remotely to distant "back office" teams for later analysis.

In the past few years, several remarkable demonstrations of the capabilities of robot teams have also been made by the combined efforts of NASA's Ames Research Centre, Carnegie-Mellon University, the Jet Propulsion Laboratory [2, 3]. For instance, the Collaborative Decision Systems (CDS) Demonstration that took place at Ames in September, 2005 showcased an integrated network of cooperating human and robot agents. The scenario included 'K9' a highly autonomous, six-wheeled MER-class wheeled rover; 'Grommit', a smaller, four-wheeled high speed personal assistant robot, one space suited individual in the test area, a remote coordinating 'habcom' and a crewmember remotely commanding each robot, as if from a habitat. The test field was a sandpit scattered with rocks, and pieces of equipment serving as landmarks. The task was a science sampling EVA, in which the robots, in the process of performing their own exploratory tasks, could be requested by an astronaut to interrupt their work, and assist with another task, which request would be granted conditionally according to a policy that prioritised tasks.

Such demonstrations depend on the accumulated efforts of dozens of paid government researchers, costly equipment and the expenditure of substantial sums of

money and time. The key consideration for the purposes of this paper, are then: without such resources, what answers to the question of robot usage can the Mars Society Australia hope to answer? It will be argued here that a niche opportunity exists for rigorous experimentation in this domain using what resources are now available to MSA. This work learns from the prior experiments. It would be a mistake to deny the value of JPL's cooperative agent network (team of communicating robots and humans) so this basic concept will be accepted in what follows. The approach taken here is to develop a specific offering for such a network: a machine oriented toward maintenance tasks instead of field science (the justification for this is given in Section 3).

Once a prototype robot has been built, it will be field-tested during a Mars simulation in order to answer questions such as: How does a maintenance robot compare with humans performing maintenance? Is a robot-human combination preferable? What are the requirements for a suitable maintenance robot? What are the best mode(s) of control (assuming "adjustable autonomy" [4]) for a eminence robot - teleoperation, high-level commanding or full automation? How simple and reliable could a robot be made that still served a maintenance role? What other tasks could a maintenance robot be expected to perform?

Answering these questions will require:

- A good evaluation method for quantifying the contribution or "value-added" expected from a human, robot or human-and-robot system for a given category of task
- A suitable example task(s) that can be modeled in a realistic surface simulation
- A human work team capable of performing the example task(s) in simulation
- A robot capable of performing the example task in simulation

The remainder of this preliminary paper will attempt to provide these four requirements. Section 2 critically examines an interesting quantitative method developed by Roderiguez and Weisbin [5] for evaluating a range of human-machine systems on a variety of tasks. Section 3 justifies the choice of maintenance as a category tasks suitable for these experiments, and analyses these tasks into independent task primitives as a step toward applying the Roderiguez and Weisbin method. Section 4 introduces the Mascot experimental field robot, currently being developed by the author, as a robot system potentially capable of inspection, servicing and maintenance tasks. A sliding automation control system is planned for the Mascot, i.e. it will eventually be capable of being teleoperated by a remote human, commanded at a high-level by a remote human or operating fully autonomously. Section 5 then proposes a practical field test, to be conducted at a future Mars Society surface operations simulation, that will apply the remaining steps of the method to

compare the value added by a human performance of the task with that of the Mascot robot system operated in one or more of its modes.

2 Roderiguez and Weisbin's method

I choose to focus on a method developed by engineers Guillermo Roderiguez and Chuck Weisbein of JPL for evaluating the performance of different agentive systems on particular tasks [5]. The method is interesting in that it allows measurements taken on very different systems, using almost any suitable criteria and metrics, with different units, to be directly, quantitatively, compared. It can be also be applied at any scale. Briefly, the method consists of the following steps:

1. A scenario involving the tasks of interest is analysed into a complete set of *functional primitives*, i.e. physically independent operations an actor might perform in carry out the task, such as Plan Path, Traverse, Find Rocks, Carry Rocks or Sense Atmosphere.
2. For each functional primitive, define one or more *performance metrics* to be used in evaluating each candidate system. E.g. for Traverse, one would include distance to be travelled, as well as degree of difficulty of the terrain to be negotiated.
3. Specify a set of *agent systems* to be evaluated: these can be particular robots, humans or a combination of both.
4. For each agent system, specify the *resources* needed to deploy it on each functional primitive. E.g. for Traverse, the mass and power of individual agent systems might be measured. This keeps the comparison fair, by compensating for the differences in performance which might be due to different classes of machine or human tackling the operation.
5. Either by analysis, simulation or experiment, the performance of each agent system is then evaluated on each of the functional primitives and a composite score $s(m)$ is computed that estimates the aptitude of each agent system for each operation. This is combined with a composite score $r(m)$ estimating the resource consumption to form a comparative ratio called *value-added* $v(m)$.

Values of $v(m)$ are the output of the process and may be interpreted as “the ratio of additional performance due to system m to the additional resources needed to implement this system when compared against the performance and resources of the reference system [5, p.173]. They can thus compare any of a number of competing systems.

From a practical perspective, it could be difficult to specify the input parameters for the calculations. In particular, Step 1 could require some effort to create new

functional primitives and ensure that they are independent, although the total possibilities for these should be limited, and a common pool of “standard” primitives would soon become available if these were always well-described in publications. In Step 4, appropriately characterising resource requirements would be difficult for some tasks. Would the dollar cost of a system be an appropriate parameter? Do these also have to be independent? Once chosen, it would generally not be as conceptually difficult to decide on appropriate metrics, but this could still present some difficulty: how would one account for the resources expended in a human or robot system that opportunistically used an existing measuring device to gather extra data? The difficulty is not one of differences between the units of measures because the method is specifically designed to use a multiplicity of measurement units and reduce them to standard units of the bit by the final step. It is rather, about understanding the abstract relationship between inputs and outputs well enough to make good choices.

Is the method theoretically sound? Most of calculations involved in the method are straightforward and uncontroversial. Mathematically speaking, a potential problem arises from the choice of an information-theoretic measure. In [5], Equations 4 and 7 describe task growth and resource growth, respectively, as base 2 logarithmic functions, drawing inspiration from the human performance work of Fitts [6]. That work proposed an Index of Task Difficulty (ID) for experiments involving human placement of limbs at a target position:

$$ID = \log_2(2A / W) \quad (1)$$

where A is the size of the motion required to place the limb at the target and W is the width of the target.

The unit is bits, because Fitts was interested in quantifying the information processing capacity of the motor nervous system and wanted to apply Shannon’s information theory [7], where the bit is the fundamental unit of complexity. However, according to MacKenzie [8], Fitts may have erred by adopting a simplified variation of Shannon’s work. MacKenzie argues on theoretical and experimental grounds that, unless $A:W \gg 1$, the behaviour of Equation 1 will depart from Shannon’s well established model of the information capacity of a channel as limited by its signal-to-noise ratio, and that Shannon’s original formulation ([7], Equation 17, p.100-103) should have been used instead. According to MacKenzie, in some applications of the Index of Task Difficulty, the ratio has been observed at unity or less, which condition would have invalidated the measure.

Our concern is that Roderiguez & Wiesbin, in adopting Fitt’s idea, may have inadvertently made the same error. Now combining Equations (1) through (3) from their paper and replacing a product of ratios with a ratio of products, one of the two affected corresponding measures for our purposes is the performance s of system m

$$s(m) = \log_2 (|P(m)| / |P(1)|) \quad (2)$$

where each P(m) is the product over all performance measures of system m

and system $m=1$ is arbitrarily chosen as a standard reference.

A second measure, $r(m)$ is similarly defined, but for resource consumption. The question becomes: are these ratios likely to approach unity? The answer is clearly yes, since any system m might return very similar performance measurements to those of the reference system 1. The same would be true of resource consumption. Fortunately, the solution is at hand; as MacKenzie points out, there is no reason why Shannon's original equation may not be used instead. In the case at hand, that would amount to replacing Equation 1 in the Roderiguez & Wiesbin paper with

$$p(k,m) = p(k,m) + p(k,1)/p(k,1) \quad (3)$$

as well as the corresponding alteration to the $c(k,m)$ resource ratio. If the proposed field trials can be realised, calculations using both variations can be compared to gauge the actual magnitude of this problem.

Another problem is one common to econometric analysis of this kind: that it could be focused too narrowly on achieving readily-measurable outcomes at the expense of less tangible, but still real, outcomes. Suppose an analysis based on science productivity measures such as number of sites visited, hypotheses generated, etc. [e.g. 9] returned a finding that the optimal science could be done by leaving human astronauts in Mars orbit and conducting the exploration by controlling robots on the surface (as is actually proposed by Landis [10]). Choosing this option might well be a cheaper, safer and more efficient way of doing science but from a broader, cultural perspective such a mission is clearly deficient, both for the human crew and for the taxpayers vicariously experiencing it. They would be "spared" the experience of landing, ascending, living and working on another planet - and nothing of these important matters would be learned. The remedy to this drawback is to find a way of properly valuing the less obvious benefits of human presence so that it can be input to and accounted for by the Roderiguez and Weisbin method. This would be the equivalent of efforts by environmentalists to revolutionise business accounting so that it does not undervalue the contribution of natural resources or a clean environment as inputs. But although theoretically feasible, deciding how to include intangible benefits in the Roderiguez and Weisbin method would complicate the already difficult matter of how to choose and weight the component primitives¹.

Despite these problems, the potential usefulness of the Roderiguez & Weisbin method can scarcely be overstated and so it should be refined and applied by all means.

¹ One of the example primitives offered by Roderiguez & Weisbin, called "Be There", accumulates risk to the human astronauts over the time taken in an EVA. It is zero for robots and apparently all negative for humans.

3 Choice of task

The chosen demonstration scenarios for most of the robot development at NASA centres over the past decade still reflects the prevailing funding environment prior to the Bush administration's commitment to a return to human spaceflight in 2002: a culture of Earth-controlled robots doing exploration and science. As the emphasis moves to human spaceflight, the trend now is toward cooperative human-robot teams,



Fig. 1. Maintenance tasks will represent a considerable, ongoing burden for future Mars explorers, unless the workload can be reduced by robots. Here engineer Matt Bamsey repairs a collapsed water pipe support outside the Mars Desert Research Station, Utah during a 2003 simulation. .

but the focus is still on glamorous field science. This category of task is therefore quite well studied, and probably not worth revisiting for our purposes.

On the other hand, taking care of the base has not received so much attention. Monitoring and maintaining all the equipment required to support human exploration in optimal condition over a many months will represent a lot of work. Examination of actual crew workloads on the International Space Station (ISS) reveals that a substantial proportion of even the science crew's time is spent on planned and unplanned maintenance tasks [11,12]. From the author's experience at the MDRS [13], maintenance work on for small crew at the first Mars base is likely to be even more demanding (Figure 1).

An Ames Research Centre study of human versus robot rover science returns in a Mars simulation suggested that human beings are 1-2 orders of magnitude more productive than robots at field science^[9]². But even if it were shown that human astronauts were inferior to robots at this category of task, it is difficult to imagine a realistic scenario in which they took the trouble to fly to Mars but did not actually take a lead role in exploratory science (see Section 2). Once humans arrive, it is far more likely that robots will be cast into supporting roles, not the least of which would be relieving the astronauts of the burden of servicing and maintaining the other equipment, and themselves. For many outside tasks that did not deserve the direct attention of humans, the time, effort and risk reduction of robot work would be highly desirable.

For convenience, I shall categorise tasks into three levels of increasing difficulty for a robot, depending on the nature and predictability of the task.

Level 1. Location-based non-manipulation tasks (e.g. still and video imaging; transport of tools and consumables; instrument positioning) It is only necessary for the robot to navigate accurately to a location such as a possible trouble spot and take high-resolution photographs of the equipment concerned for transfer to an engineer's station. In teleoperation, the machine is guided by the human operator; in high-level commanding and full automation, the robot must plan a path between waypoints that avoids obstacles. Such a robot also could fetch and carry tools, equipment and samples on command.

Level 2. Use of manipulators for planned, structured tasks (e.g. repair-by-replacement; spraying of paint, lubricant or sealant; changeout of a dust filter or replacement of a gas cylinder; staking or pegging structures such as solar panels or antennae; connecting and tightening electrical or stay cables; loosening or tightening bolts and nuts; sweeping or blowing dust off solar panels, instruments or cameras). This task requires in addition to accurate navigation the provision of one or more manipulators and/or specialised tools and the skill to bring those tools to bear on a particular work item. In teleoperation, the skill is that of the remote human operator; in high-level commanding it requires sophisticated sensors and intelligent control software. Scheduling would come from human-supervised, overrideable, automated scheduling software working to a routine maintenance schedule (both off-board the robot).

Level 3. Use of movement and manipulators for unplanned, unstructured tasks (e.g. repair on demand, given a diagnosis; disassembly and assembly of machines according to manufacture's procedure; repositioning of fallen or displaced equipment; opening or closing stuck valves, doors and panels; unfreezing pipes; simple testing of electronic and mechanical components). These tasks require everything required in Level 2, but also presuppose a certain degree of problem-solving, and error recovery. This would come from human intervention, planning and

² This is actually a claim of the kind that should be better quantified using the method described in Section 2.

reasoning overriding automated routine maintenance schedules. A larger selection of tools, probably more sophisticated sensors and probably a greater amount of applied force from the manipulators would be required. Such skill is difficult, but not impossible, to demonstrate [14]

From an evaluation point of view, what functional primitives would be involved in the performance of such tasks? Because it is advantageous to have a small, standardised set of these available to all, the first step in any such specification should be to examine the existing primitives and try to use what is there. New primitives should be created reluctantly, and only if there is nothing suitable on the shelf. From the list of examples in [5] we see that Traverse (moving from one specified location to another, characterised by distance, speed, and terrain difficulty), Recover From Mishaps (overcoming relatively simple operational mishaps such as a fall and verify that no damage has occurred) and Carry Rock, renamed as Carry Equipment, (characterised by mass, volume and distance carried) could be used. Similarly, Find Rocks should be renamed as Find Jobsite (speed and accuracy with which vision systems could locate, recognise and project a working calibration onto a specified object of interest). To this we should add Grasp (ability to apply force to turn or lift an object, characterised by force/torque applied and mass and dimensions of object). Skilled Tool Use should also be added, to capture the need to apply human and machine skill to the use of a specific tool (correct selection of tool, speed and accuracy of placement, time to completion).

4 Mascot field robot

A team of two (simulated) human astronauts working on specific maintenance tasks would form one agent system to evaluate (and would probably be chosen as the reference system). The Mascot field robot (Figure 2), currently being developed by the author, is another. It is designed as a service robot and could be adapted to serve as a simulated maintenance machine for Mars explorers. This machine is designed to demonstrate that six-legged locomotion can provide good speed, traction and stability in uneven or broken terrain that cannot be matched by wheeled machines. However, in order to avoid the well-known problem of unreliability in complex, jointed leg systems, the mechanism has been greatly simplified. Inspired by a similar machine called RHex [15], the Mascot has six, simple passive spring legs, each mounted on an independent revolute axis (6 DoF in total) and driven by an 18V Metabo 100W DC motor fitted with a 150:1 planetary gearbox. The six motors are driven by Jeffrey Kerr LLC PIC-SERVO control boards connected to a 32-bit RS485 multidrop network controlled by an onboard Sony Viao laptop running Windows XP.

This configuration is simple, reliable and robust, yet provides remarkable agility and control. The machine is 590mm long, 570mm wide at the middle legs and 700mm from the ground to the top of the current camera mast. It weighs approximately 14kg. As with insects, the machine moves by "tripod walking": at any instant three legs are on the ground, and these alternate between the sides of the robot.

The machine is steered by altering the phase relationship between the tripods on either side. Although not yet measured, the machine is expected to be able to achieve a speed of at least 0.5 m/sec. on uneven ground. The main power supply for the motors consists of two 18 volt, 13Ah Lithium-Ion battery packs with built in voltage regulators and thermal shutdown circuitry. A 12v 2Ah Lithium Ion battery supplies logic power. The camera and control receiver are both independently powered by small NiCd battery packs. The camera platform is designed around a pair of EO5-380



Fig. 2. Prototype of the Mascot field robot being developed by the author.

CCD cameras mounted on a tilt-pan head. Each camera is capable of transmitting 380-line PAL colour video over at 2.4GHz wireless link. At this stage the cameras are not used by the robot as a vision system, but only as part of a low-cost teleoperation control system. This also depends on a commercial 6-channel 36MHz FM wireless remote control system, designed for model aircraft. Two channels of this control the effective left and right steering, and two channels control the tilt and pan motors of the camera head. When completed, the system will enable a remote operator to control high-speed motion of the robot while viewing real-time video from the cameras on a small LCD monitor. Depending on the performance of the machine, the project may progress to a high-level commanding mode or even to full automation.

What tasks could a robot like the Mascot take on while setting up and operating a surface base, and in exploratory work? The answer depends on the task, the mode of control required (teleoperation, high-level commanding or full automation) and the provision of hardware and software for the total robot system (see Table 1). We will restrict our attention to maintenance tasks for the reasons discussed in Section 3. This table should be interpreted as showing increasing, cumulative demands on the equipment and behavioural competence of the robot as we move from the lowest demands in the top left of the table to the greatest demands in the bottom right. Thus simplest operational form of Mascot would be capable of Level 1 tasks if teleoperated, because those tasks only need accurate navigation and photography. At the other extreme, a long, well funded research effort would be required to provide the all requirement specified in the table, including real-time planning and error-recovery software in order to automatically cope with unplanned, unstructured repair jobs.

Table 1. Cumulative requirements of the Mascot field robot relative to developed control mode and level of task.

	<i>Teleoperation</i>	<i>High-level Commanding</i>	<i>Full Automation</i>
<i>Level 1 Tasks</i>	functioning basic system allows inspection, photography, fetch and carry	add compass, GPS and obstacle- detecting sensors; add software planning layer on behaviour based reactive layer	connect cameras to vision system; add vision software to frame photographs, recognise humans
<i>Level 2 Tasks</i>	add at least one 6 DoF manipulator arm; routine maintenance manuals for operator	add specialised, detachable tool ends; vision software and touch sensors to guide tool use algorithms	add human-interruptible maintenance scheduling software (off-board); algorithms for selection of tools
<i>Level 3 Tasks</i>	Provide more force at manipulator, tool end; more manipulators; detailed troubleshooting manuals/software for operator	add more and better sensors depending on task; voice command software	add best available real-time planning software (off-board); error recovery software

Table 1 also suggests a research direction for future work on the Mascot robot: left to right and top to bottom. It makes sense to try to add high-level commanding only once teleoperation is perfected, and progressing to full automation will be easier once high-level commanding is perfected. For example, the necessary skilled motions for the robot to open an access panel and remove the circuit board inside on its own might be able to be acquired by a learning algorithm in the robot while it is being guided through these actions in teleoperation. If more resources can be made available to the robot, such as better sensors and lightweight, multi-jointed arms with

manipulators for the front of the machine, it will be possible to progress down the table to the higher levels of functional skill.

5 A Possible Experiment in the Field

How can the use of the Mascot robot be tested in an analogue surface simulation, such as those conducted by the Mars Society and how can its performance be evaluated quantitatively in comparison with human astronauts on maintenance tasks like those described in Section 3? Conceptually at least, if we choose a simple Level 1 task - maintenance photography - we now have the four requirements of Section 1: a good evaluation method, a suitable task, a human team that can do the task in simulation and a robot that can do the task in simulation. At a minimum, the Mascot robot will be able to carry out this task at the next Mars Society simulation, at which it is also extremely likely that a volunteer human astronaut team of two could be found for the comparison.

Physically, the task would require each agent system - the pair of astronauts, and the teleoperated Mascot, and a combination of the two - to take a series of high-resolution photographs at a number of key equipment sites at various distance from the habitat. A taxing list of real or dummy equipment panels, bolts, connectors etc. would be nominated or set up in the vicinity of the habitat. The amount of time, resources, risk taken as well as the quality of the resulting photographs would be assessed. These data would then be processed using (both variations of) the Roderiguez & Weisbin method to decide how they compared.

The Mascot robot can probably be modified to carry out Level 2 tasks, but Level 3 tasks are expected to require a more massive, better engineered machine. It is, however, neither necessary nor wise to tackle all levels of tasks at once. In developing robots to tackle ever more complex tasks, it might also be possible to assess an important matter about maintenance by robot: to what degree does extra complexity need to be added to the system in order to carry out the higher level tasks, and at what point do the maintenance needs of the robot itself begin to impose more of a burden on the mission than they are worth?

References

1. Clancey, W.J., Sierhuis, M., Alena, R., Crawford, S., Dowding, J., Graham, J., Kaskiris, C., Tyree, K.S. and Hoof, R.V., The Mobile Agents Integrated Field Test: Mars Dessert Research Station, 2003. Proc. of FLAIRS 2004, Miami Beach, Florida (2004).
2. Culbert, C., et.al. Activities of the NASA Exploration Team Human Robotics Working Group. Proceedings of the American Institute of Aeronautics and Astronautics Space 2003 Conference, September, Long Beach, Calif. (2003).

3. Pedersen, L., Clancey, W., Sierhuis, M., Muscettola, N., Smith, D., Lees, D., Rajan, K., Ramakrishnan, S., Tompkins, P., Vera, A., Dayton, T. Field Demonstration of Surface Human-Robotic Exploration Activity, AAAI Spring Symposium 2006 (2006)
4. Dorais, G. A., Bonasso, R. P., Kortenkamp, D., Pell, B. and Schreckenghost, D. Adjustable Autonomy for Human-centered Autonomous Systems on Mars. Proceedings of the First International Conference of the Mars Society (1998)
5. Roderiguez, G. and Weisbin, C.R.: A New Method to Evaluate Human-Robot System Performance. *Autonomous Robots* 14, 165–178 (2003)
6. Fitts, P. M. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47, 381-391 (1954).
7. Shannon, C. E. and Weaver, W. *The Mathematical Theory of Communications*. Urbana, IL: University of Illinois Press. (1949)
8. MacKenzie, I. S.. A Note on the Information-Theoretic Basis for Fitts' law. *Journal of Motor Behavior*, 21, 323-330 (1989)
9. Glass, B., Briggs, G., Jasper, J. and Snook, K. Evaluation of Human vs. Teleoperated Robotic Performance in Field Geology Tasks at a Mars Analog Site. Proceedings of iSAIRAS 2003, Nara, Japan, May (2003)
10. Landis, G.A. Robots and Humans: Synergy in Planetary Exploration. *Acta Astronautica* 55, p.985–990 (2004)
11. Acquisti, A., Sierhuis, M., Clancey, W.J., Bradshaw, J.M. Agent Based Modelling of Collaboration and Work Practices Onboard the International Space Station. Proceedings of the Eleventh Conference on Computer-Generated Forces and Behaviour Representation. Orlando, Florida (2002)
12. Russell, J. F., Klaus, D.M. and Masher, T. J. Applying Analysis of International Space Station Crew-time Utilization to Mission Design. *Journal of Spacecraft and Rockets*, 43, 1, p.130-136 (2006)
13. Persaud, R. Rupert-Robles, S. Clarke, J.D.A., Dawson, S., Mann, G.A. Waldie, J., Piechocinski, S. and Roesch, J. Expedition One: A Mars Analog Research Station 30-Day Mission (AAS 03-304). In Cockell, C. (Ed.) *Martian Expedition Planning*. AAS Science and Technology Series, Vol 107, Univelt Publishing, San Diego, California, p53-89 (2004)
14. Nickels, K., Kennedy, B., Aghazarian, H., Collins, C., Garrett, M., Magnone, L., Okon, A. and Townsend, J. Vision-Guided Self-Alignment and Manipulation in a Walking Robot. Proc. of the 2006 IEEE International Conference on Systems Engineering, April 24-26, 2006, Los Angeles CA (2006)

15. Saranli, U., Buehler, M. and Koditschek, D.E. RHex: A Simple and Highly Mobile Hexapod Robot, *Int. Journal of Robotics Research*, 20, 7, p616-631 (2001)