AUTONOMOUS ROVER PATH PLANNING AND RECONFIGURATION

J. Graciano, E. Chester

AEVO GmbH, Friedrichshafenerstraße 1, 82205 Gilching, Germany
joao.graciano@aevo-technologies.de

ABSTRACT

The autonomous ability to plan collision-free trajectories is a critical enabler for planetary exploration rovers, potentially saving time and resources in a mission while enhancing its value to the science community undertaking the project. In this paper, we concentrate on autonomous guidance and navigation, and apply it specifically to the path planning function.

A method based upon artificial potential fields is presented for autonomous path planning and for path reconfiguration of planetary rovers. The artificial potential functions are re-interpreted as the payoff function in an iterated, asymmetric, and non-cooperative, non-zero-sum game, wherein each terrain obstacle and scientific target is represented by a game player. Since the players are implicitly independent, we show that the position at the target is the only Nash equilibrium of the game, thus guaranteeing a feasible path solution in every case, and thus eliminating the local-minima problem associated with prior artificial potential methods. The algorithm developed allows for path optimisation and the incorporation of constraints and furthermore is shown to offer a high degree of flexibility, enabling integrated combination with higher-level navigation planning functions.

We present simulations with several obstacles and targets, and show how the algorithm can cope with the selective activation of targets as a method to build extended trajectories through any terrain. We conclude by introducing ongoing improvements to the method and its application in adaptive trajectory planning.

1 INTRODUCTION

One of the aspects that gains more and more relevance as the autonomy of planetary robots increases, is the planning of a collision-free trajectory on a planetary surface. Until a few years ago, the planning of a trajectory was done by operators on Earth. However, this uses precious bandwidth and communication resources, while also reducing useful activity time of the rover. Several trajectory planning techniques have been applied to the problem of rover navigation, with the most well-known being Grid-based Estimation of Surface Traversability Applied to Local Terrain (GESTALT) [Maimone *et al.* 2006] as used by the Mars Exploration Rovers (MER). While directly handling obstacle avoidance, this method is augmented by stereo-vision navigation for compensating odometry estimation errors in regions of poor traction. Wide-angle Hazard Avoidance cameras at the front and rear of each MER are used autonomously to provide local range maps in support of the planning function.

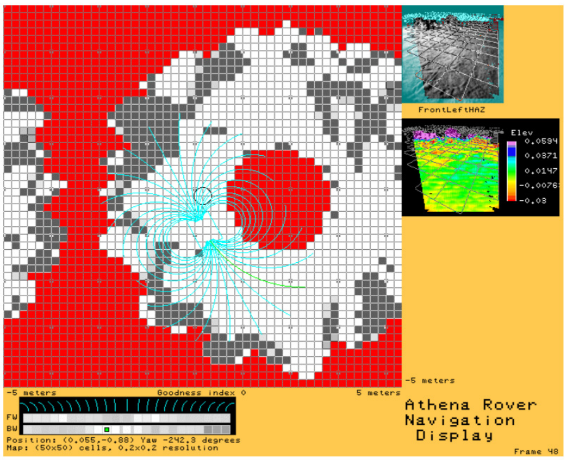

Figure 1: MER's GESTALT Navigation Planning [from Eisenman et al. 2005]

The drivers for most of the development in the area of path planning include increasing the degree of autonomy to cope with only poorly characterised terrains; risk reduction; enabling decision making executive functions in support of autonomous scientific investigation of targets of interest; managing onboard resources – especially power for locomotion; reducing operational costs of ground support; and mitigating the inconvenience of significantly time-lagged teleoperation. Taken together, these all in some sense contribute to an overall goal of enhancing mission utility or scientific return.

Beyond rover navigation, these considerations are increasingly important in other remote robotics or spacecraft operation scenarios, such as on-orbit servicing, requiring robust rendezvous strategies. A combination of robustness and flexibility is taken as a top-level requirement for path planning in such scenarios, together with the potential to be implemented in a variety of platforms for different operational contexts.

In this work, we concentrate on autonomous guidance and navigation, and specifically on the path planning function and its role for autonomous robotic planetary platforms. The motivation for developing the new method was threefold: to develop a method capable of dealing with a multitude of terrains, thus increasing a rover's autonomy; to increase the path planning reliability and testability by keeping the method as simple as possible; and to make it modular as to be incorporated into GNC subsystems across a wide range of autonomous platforms.

The remainder of this paper is organized as follows. In section 2 we describe the artificial potential method and in section 3 we briefly present the aspects of game theory that are essential for the understanding of the proposed method. In section 4 we define the proposed game algorithm used in this work, with simulations presented in section 5. Discussion and conclusions follow in section 6.

2 ARTIFICIAL POTENTIALS

In the group of methods for path planning, the artificial potential method is one of the most efficient. While providing a general field without the need for recalculation at every step, it also avoids the "stochastic" aspect of other methods (*e.g.* RRTs).

The artificial potential method is based on artificial potential (AP) functions or fields. In this scheme, obstacles are represented by repulsive potentials and goals by attractive potentials. A total potential field is built through the weighted sum of all the individual AP-fields, and the robot uses the gradient of that field to navigate, from high potential states to low potential states. However, unless very specific functions are chosen, the potential field can (and usually does) suffer from local minima, causing the robot to stop at undesired locations, and requiring an executive action (generally from human intervention) to be taken.

Some authors have identified functions that do not lead to local minima. Examples include harmonic functions [Keymeulen and Decuyper 1994], [Feder and Slotine 1997], superquadric potentials [Khosla and Volpe 1988], and functions that are solutions to Maxwell's equations [Hussein and Elnagar 2002]. Others have concentrated on strategies to avoid or escape the local minima through the use of random walks [Chang 1996], by iteratively modifying paths [Warren 1989], by positioning virtual obstacles near the local minima [Park and Lee 2003] or by using simulated annealing [Janabi-Sharifi and Vinke 2000], [Kirkpatrick, Gelatt and Vecchi 1983], [Park and Lee 2002].

Despite the multitude of solutions proposed, the artificial potential method remains very sensitive to the potential functions used, and when it is not, suffers from the local minima problem. Because of this, unexpected terrains and errors in the characterisation of the problem can lead to local minima in the total field, and gradient methods can suffer or fail completely.

3 APPLIED GAME THEORY

A game, in the mathematical sense, refers to the mathematical modelling of a situation of interaction involving two or more players or agents. We can describe a game by defining the number of players, the set of possible moves, the payoffs awarded to each one of them in function of all the possible combinations of moves (*i.e.* in function of the state of the game), and the game's rules. The objective of the game is usually for each of the players to obtain the highest payoff possible.

In the simplest case, the payoffs to each player are defined by matrices, and each element of the matrix is the payoff for the corresponding combination of moves. Assuming the players are rational in the sense that they try to have maximum payoff, and taken that the players (although knowing which moves are available to all players) do not know what a player's next move will be, the strategy to find the best move is: for every possible adversary move, find the counter moves that gives the best payoff. When all players behave in this "rational" way, the concept of Nash equilibrium applies. A given combination of players' moves is said to be a Nash equilibrium if no player has a garantee of increasing his reward by *unilateraly* departing from that given combination [Masterton-Gibbons 2001]. Formally, a combination of moves $(s_1^*, \ldots, s_n^*)$ is a Nash equilibrium if for each player *i* if the following holds:

$$\forall s_i \in S_i, s_i \neq s_i^* : \quad f_i\left(s_1^*, \ldots, s_i^*, \ldots, s_n^*\right) \geq f_i\left(s_1^*, \ldots, s_i, \ldots, s_n^*\right) \tag{1}$$

where f_i is the payoff matrix for player *i* and S_i is the space of possible plays.

The possibility for a player to unilateraly change his play is thus an essential aspect of the Nash equilibrium concept.

In general, to find the Nash equilibrium of a game, one proceeds by first finding the rational set of solutions for each player,i.e., the set of moves which are best answers to all possible moves by the adversaries. Performing the intersection of these sets gives the equilibrium combination of moves, such that if one players chooses to deviate from it, it will have no garantee that his payoff will increase.

4 A MULTI PLAYER GAME FOR NAVIGATION

The concept of Nash equilibria thus provides a very attractive concept from the point of view of path planning. It would be highly desirable to have, in general, an equilibrium solution, which in our case would be the target point at which the robot stops moving. To aproximate the formulation of the path planning problem to that of a game, we reinterpret the obstacles and the target of the terrain as players, where the payoff functions are now akin to artificial potential functions centered at the obstacles' and target's position. According to the game theoretic framework, each player tries to maximize its payoff by manoeuvering the robot (whose position now defines the state of the game) in a certain direction. Because each player tries to increase his payoff, the direction is given by the gradient of the artificial potential function. But given that now a combination of moves is a combination of coordinates, a change of move of a player is actually a change of robot coordinates, *which are the same for all players*. That is, no player can *unilateraly* change his strategy. To enforce this possibility, we implement the independence between players in the form of vector orthogonality.

We define a frame (attached to the robot moving in $\Re^2$) determined by the unit vectors $\mathbf{u_s}$ (in the direction connecting the target and the robot's position) and $\mathbf{u_t}$, respectively (see Figure 2). Variables are also defined for the players and their respective payoffs, as given in Table 1.

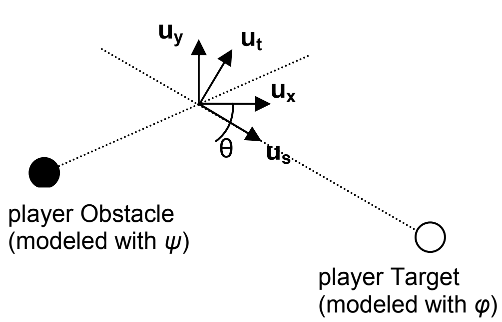

Figure 2: Graphical representation of the method

Table 1: Game variables

Players	*Player's Position*	*Player's payoff function*
Target P	(x_P, y_P)	$\varphi\ (x- x_P\ ,\ y- y_P)$
Obstacle Q	(x_Q, y_Q)	$\psi\ (x-x_Q\ ,\ y- y_Q)$

The gradient of each payoff function is then given by equations (2a) and (2b).

$$\nabla\varphi(x,y) = \frac{d\varphi}{dx}\mathbf{u_x} + \frac{d\varphi}{dy}\mathbf{u_y} = \varphi_x\mathbf{u_x} + \varphi_y\mathbf{u_y} \quad (2a)$$

$$\nabla\psi(x,y) = \frac{d\psi}{dx}\mathbf{u_x} + \frac{d\psi}{dy}\mathbf{u_y} = \psi_x\mathbf{u_x} + \psi_y\mathbf{u_y} \quad (2b)$$

$$\begin{bmatrix}\mathbf{u_x}\\ \mathbf{u_y}\end{bmatrix} = \begin{bmatrix}\cos(\theta) & -\sin(\theta)\\ \sin(\theta) & \cos(\theta)\end{bmatrix}\begin{bmatrix}\mathbf{u}_s\\ \mathbf{u}_t\end{bmatrix} \quad (3)$$

The vectors $\mathbf{u_x}$ and $\mathbf{u_y}$ can be defined in function of the $\mathbf{u_s}$ and $\mathbf{u_t}$ by a rotation (3), with the angles defined by (4):

$$\sin(\theta) = \frac{\varphi_y}{\sqrt{\varphi_x^2+\varphi_y^2}} \qquad \cos(\theta) = \frac{\varphi_x}{\sqrt{\varphi_x^2+\varphi_y^2}} \quad (4)$$

Substituting (4) and (3) into (2a) and (2b) gives:

$$\nabla\varphi = \nabla\varphi\,\mathbf{u_s} + 0\,\mathbf{u}_t \quad (5a)$$

$$\nabla\boldsymbol{\psi} = \left(\psi_x\frac{\varphi_x}{\sqrt{\varphi_x^2+\varphi_y^2}} + \psi_y\frac{\varphi_y}{\sqrt{\varphi_x^2+\varphi_y^2}}\right)\mathbf{u_s} + \ldots \left(\psi_y\frac{\varphi_x}{\sqrt{\varphi_x^2+\varphi_y^2}} - \psi_x\frac{\varphi_y}{\sqrt{\varphi_x^2+\varphi_y^2}}\right)\mathbf{u}_t \quad (5b)$$

To enforce the independence between players we require (5a) be orthogonal to (5b) and reduce the $\mathbf{u_s}$ component of (5b) to zero. We therefore build a vector field $\mathbf{L}$, given by

$$\mathbf{L} = \nabla\varphi\,\mathbf{u_s} + \left(\frac{\psi_y\varphi_x}{\sqrt{\varphi_x^2+\varphi_y^2}} - \frac{\psi_x\varphi_y}{\sqrt{\varphi_x^2+\varphi_y^2}}\right)\mathbf{u}_t \quad (6)$$

Because of the way it is built, (6) has the property that the set of points where both components of **L** are zero, in the ($\mathbf{u}_s$, $\mathbf{u}_t$) frame, has only one element – the point at the center of the function φ, i.e., at the target's position. This creates a global maximum in a field that is free from local minima. From the game theoretical point-of-view, each player's move is independent of all other player's moves, and the only Nash equilibrium point is the target's position.

Consider a simple example of a field with a target at position (2,1) and an obstacle at position (1,1). Figure 3 shows the gradient of the several functions involved. The following is immediate: the rational set of the player target has only the point at the target's position (upper figure); the rational set of the player obstacle is the line defined by $(x \geq 2, y=1)$ (middle figure); the intersection of these two sets has only one point, the only Nash equilibrium, at the target's position (lower figure).

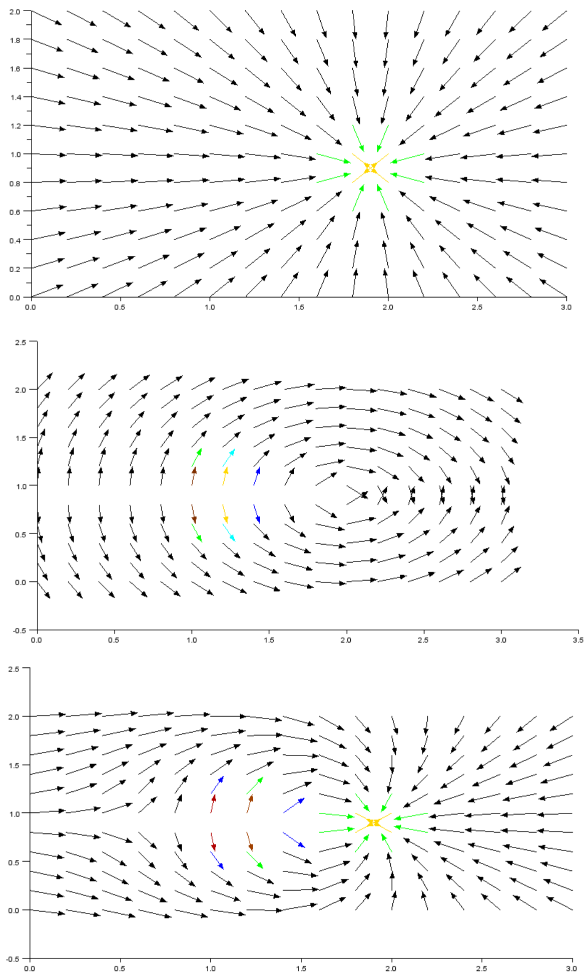

Figure 3: The gradient fields of the target function (upper picture), the modified obstacle function (middle picture) and the total vector field, given by (7) (lower picture).

The only potential breakdown of this method is found in the artificial case wherein: a) the total modified AP-field is exactly symmetric relative to the line connecting the robot and the target; b) the robot's velocity vector is parallel to that same line; and c) there is an obstacle exactly between the robot and the target. In this unlikely case, the robot will not see the obstacle. However, any perturbation from the exact path will cause the obstacle to become evident in the asymmetric field outside the line connecting the robot and its target.

To summarise, with the proposed algorithm, the robot constantly feels an attraction towards the target, while the obstacles infliuesnce its motion by pushing it *sideways*. The addition of more obstacles (even including other robots) is equivalent to the addition of more players, each manifested as an additional term in the second component of **L**.

This approach can be applied to planetary rover navigation through considering the objectives for safe path planning, including:

- avoidance of regions with a gradient higher than a safe threshold;
- avoidance of obstacles known to present a locomotion threat (rocks over threshold size, regions of suspect stability, soft sand, *etc.*);
- approach to, or passing through, regions expected to offer higher mission return (*e.g.* pass closer to a certain set of rocks);
- overall minimisation of energy usage (or traverse distance).

Several situations have been performed using this method in order to assess how well it responds to the above requirements. Quantitative assessment compared to other algorithms is not completed, but inspection of the trajectories has proven satisfactory for demonstrating that the rover trajectories found for a given scenario are as expected. For this paper, a random set of eight obstacles is considered, with a single target. Figure 6 through Figure 15 show the same obstacles (red circles) and central target (green circle). The black trajectories represent each of 200 paths selected by the algorithm in response to random starting locations for the rover.

A classic image of Martian terrain from the Viking mission (Figure) shows the types of obstacles rover operators try to avoid, and features extraction using a simple shadow-based method. The algorithms for generating local topology meshes, such as shape from shade, outwith the scope of this paper. In general, it is sufficient to assume terrain feature extraction (both targets and obstacles) is possible.

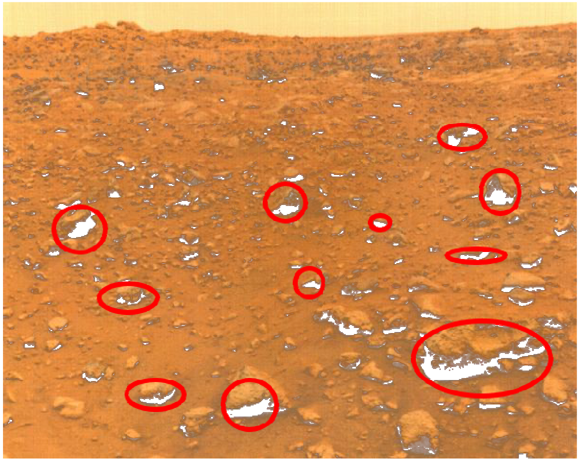

Figure 4: Viking terrain with shadow-based method for obstacle extraction and mapping

5 SIMULATIONS

The rover dynamics are described by constant-speed equations. The functions given by equations (7) through (12) and illustrated in Figure 5 were used for modelling different types of obstacles (impassable rocks, hills, sand regions), and thus used for testing the method's robustness

$$z = K/(x^2 + y^2) \tag{7}$$

$$z = e^{-(x^2+y^2)}/(x^2 + y^2) \tag{8}$$

$$z = -\log(x^2 + y^2) \tag{9}$$

$$z = -\sqrt{(x^2 + y^2)} \tag{10}$$

$$z = (\sqrt{(x^2 + y^2)} - 8)^2 \tag{11}$$

$$z = K/\sin(x^2 + y^2) \tag{12}$$

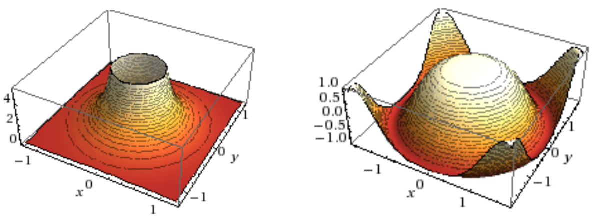

Figure 5: Representation of some types of obstacles

We first analysed the effect of the variation of the field strength (Figure 6 and Figure 7), and the possibility of modelling whole regions (Figure 8). Using (7), sets of 200 runs each were performed, first by varying the field strength, and then incorporating both forbidden and preferred regions.

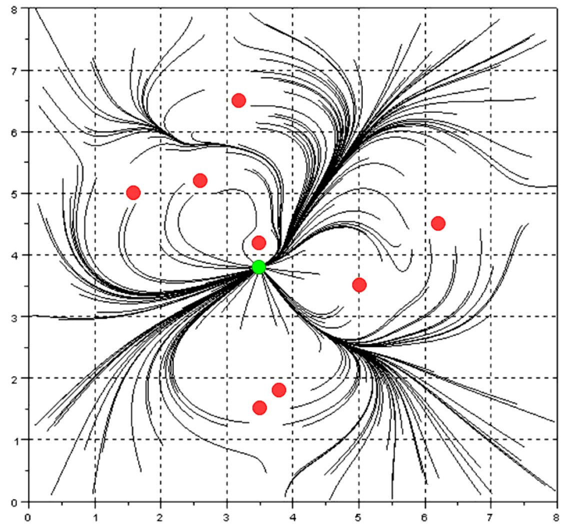

Figure 6: Using (7), with K=0.2

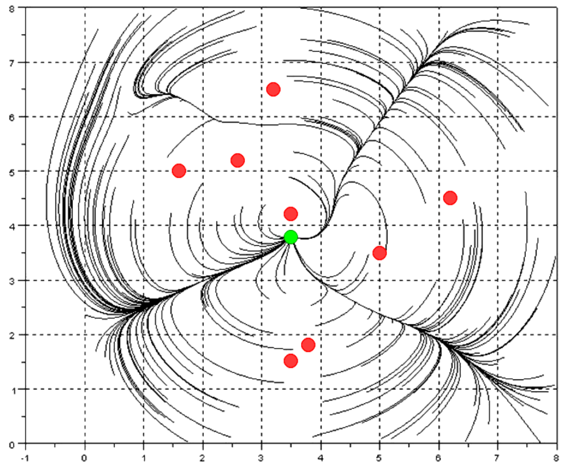

Figure 7: Using (7), with K=2

A region (square) to be avoided, and one to be preferred (in way of a tunnel) were also modelled.

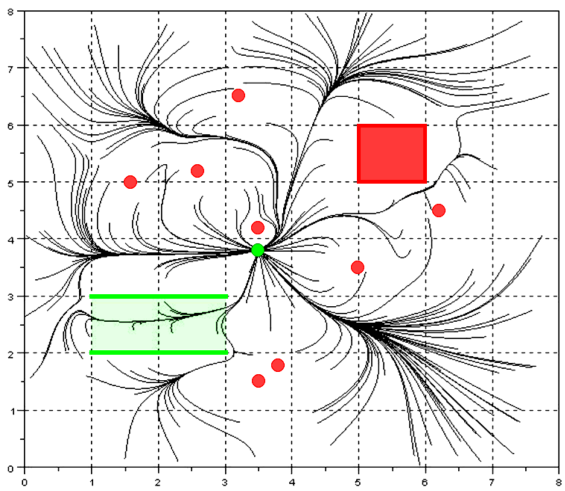

Figure 8: Regions to avoid (red) and prefer (green) (using (7))

The same set of obstacles was then simulated using functions 7-12.

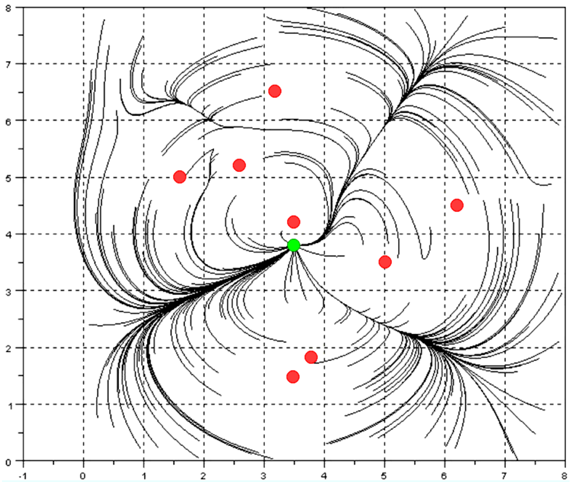

Figure 9: Using (7)

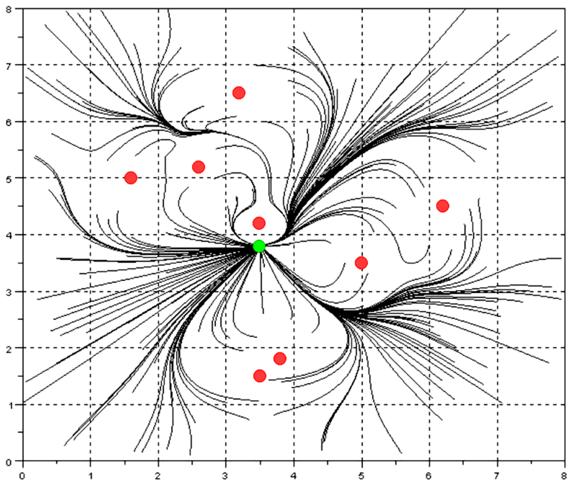

Figure 10: Using (8)

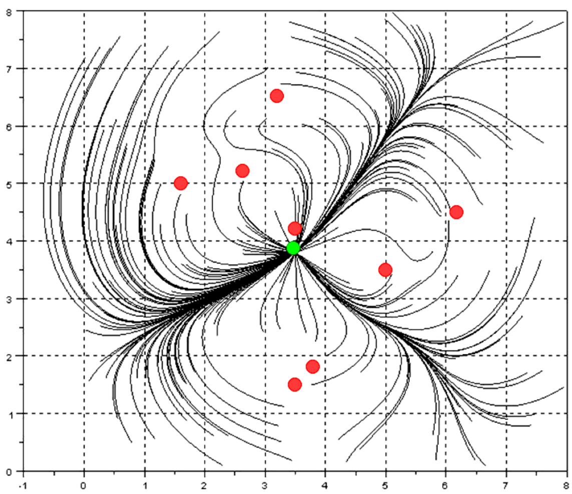

Figure 11: Using (9)

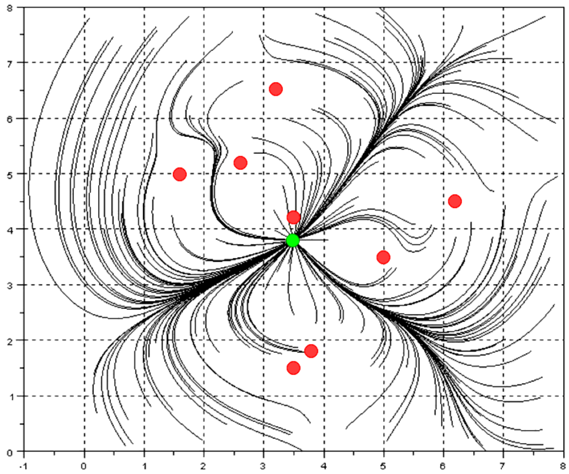

Figure 12: Using (10)

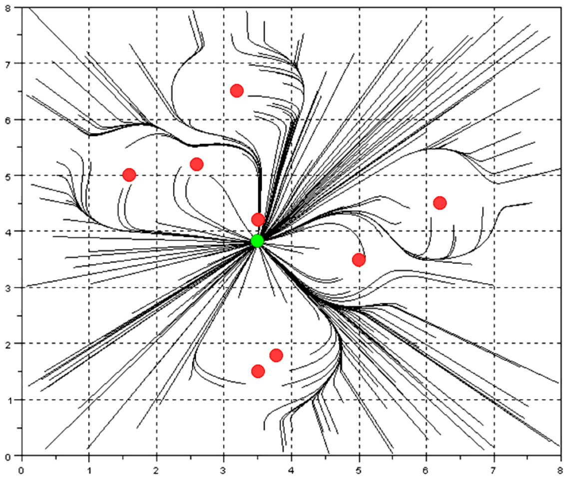

Figure 13: Using (11)

In Figure 13, a radius was set, outside of which the obstacle is not felt. This risk-minimisation strategy was again implemented using (12), and the action radius influence on the field geometry analysed.

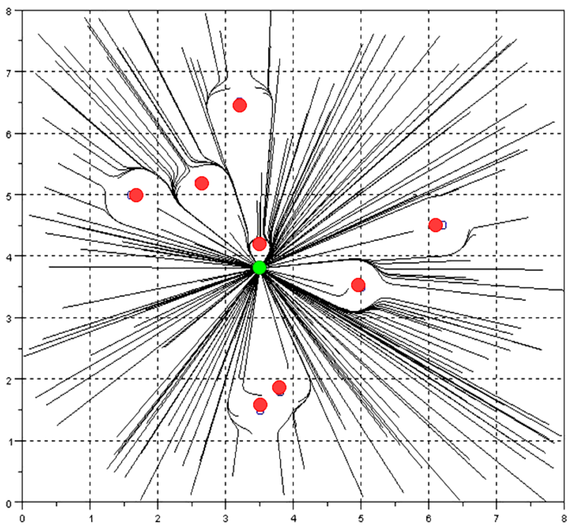

Figure 14: Using (12), R=0.5

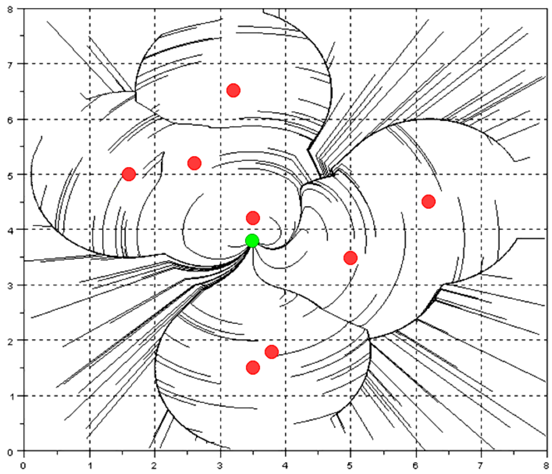

Figure 15: Using (12), R=1.5

Using a picture of Mars to achieve a slightly more realistic setup, from which specific obstacles and target were extracted, we added a procedure for target selection, in that a second target is activated once the first one is reached (Figure 16).

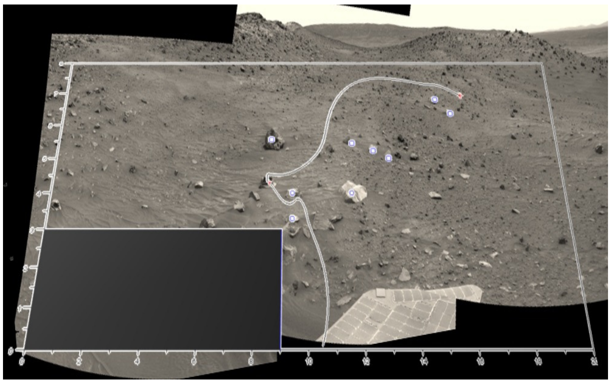

Figure 16: Superimposed on the terrain, are: a high-risk area (lower left), two targets (activated sequentially) and obstacles. Starting position is below, centre (Credits NASA/JPL).

Finally, to verify the method's capability in serving optimisation objectives, we coupled the method with an optimisation engine, and optimised the field's parameters for shortest path (Figure 16).

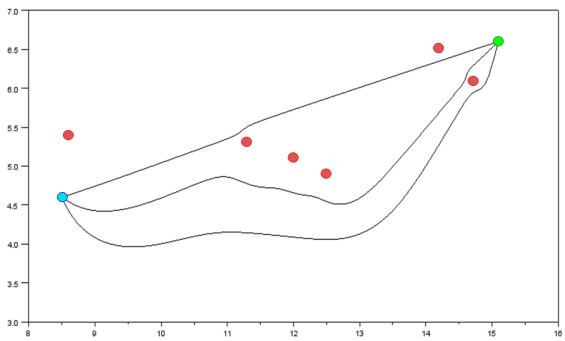

Figure 17: Two paths found during the optimisation process, and the shortest path found (using (7).

6 DISCUSSION AND CONCLUSIONS

In this paper the problem of navigating a robot using AP-functions was interpreted in the framework of game theory. That led to the implementation of player independence in the form of vector orthogonality between players' actions (*i.e.* independence between the directions along which the players make their moves). The main result is that the generated AP-field presents no local minima, i.e., in all test cases, all obstacles were avoided and all paths reached the target, bringing a significant improvement over the prior art.

In Figure 6 and Figure 7 we studied the impact of the field strength. As the field is intensified, the 'valleys' in the field geometry grow deeper, and nearby paths are quickly aggregated. In Figure 8, obstacles regions were modelled as a square and as a pair of lines to show the ability of the method to deal with obstacles of different dimensions. Figure through Figure 13 show the field modelled through the various test functions. Different functions influence differently the field's structure: some functions favour large 'valleys'; others allow much more freedom around obstacles while approaching the target more rapidly. In Figure 14 and Figure 15 an avoidance radius is modelled such that the rover maintains a certain minimum distance to the obstacle. This method also allows limitations on rover steering capabilities to be incorporated into obstacle avoidance. As expected, robots inside the radius first navigate away from the obstacle until they are clear of the avoidance radius, and then follow a path towards the target. We also demonstrated the ability of the method to cope with optimisation procedures. The optimisation, in our case, focused on the choice of the optimal field parameters in order to find the shortest path. Finally, we used a typical Mars surface to demonstrate how selective target activation can be used to create more complicated mission scenarios.

As mentioned in section 4, the method can lead a robot to ignore an obstacle positioned between it and its target. These particular cases are rare (the set of possible cases has null measure) and unstable, making relatively simple to define mitigation strategies. At the outset one of the main objectives was to create a flexible method, eventually to be implemented across platforms. We therefore also simulated cases of target pursuit, motion in formation and chase/avoid behaviours, aspects which could be relevant for on-orbit servicing applications and unmanned vehicle flight planning, to mention a few (Figure 18).

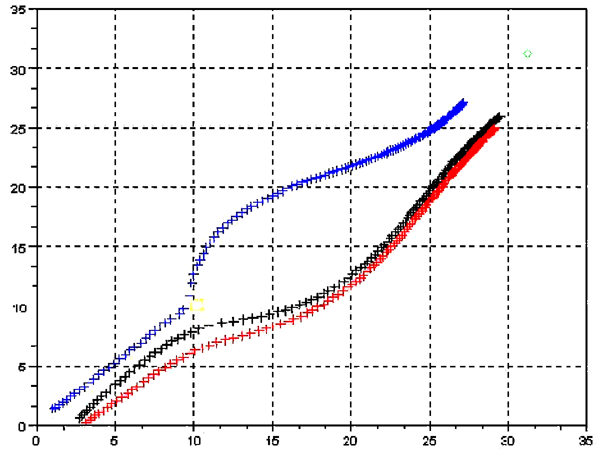

Figure 18: Obstacle avoidance by an agent group in formation

Several improvements are planned for the further development of the method, among which are the inclusion of a more realistic dynamical model and the generalisation to three dimensional spaces. Overall, the method developed shows considerable improvements over other methods:

- it guarantees reaching a target (does not suffer from local minima);
- it maintains robustness independently of terrain geometry;
- it allows the robot to instantly react to changes in the field (no need for complete path computation)
- is flexible enough to be implemented in several applications, over several platforms (rover navigation, on-orbit servicing, motion in formation);
- is simple, thus reducing costs associated with development and test;
- is easy to change to incorporate new functionality.

7 REFERENCES

H. Chang, “A New Technique to handle Local Minimum for Imperfect Potential Field based Motion Planning”, IEEE Int. Conf. on Robotics & Automation, 1996

A. Eisenman, C.C. Liebe et al., Mars exploration rover engineering cameras, Jet Propulsion Laboratory, 2005, Online article: http://www-robotics.jpl.nasa.gov/publications/Reg_Willson/MER_Cameras.pdf, Accessed: Jan 2011.

H. Feder and J.-J. Slotine, “Real-time Path Planning using Harmonic Potentials in Dynamics Environments”, IEEE Int. Conf. on Robotics & Automation, 1997

J. Graciano, "A simple navigation algorithm with no local minima", 12th International Conference on Advanced Robotics, 2005.

A. Hussein and A. Elnagar, “Motion Planning using Maxwell’s Equations”, IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems, EPFL, Lausanne, Switzerland, October 2002

R. Isaacs, Differential Games, SIAM Series in Applied Mathematics, John Wiley & Sons, Inc., 1965

R. Isaacs, Rand Reports RM-1391 (30 November 1954), RM-1411 (21 December 1954), all entitled in part Differential Games

F. Janabi-Sharifi and D. Vinke, “Integration of the Artificial Potential Field Approach with Simulated Annealing for Robot Path Planning”, IEEE Int. Conf. on Robotics & Automation, 2000

D. Keymeulen and J. Decuyper, “The Fluid Dynamics applied to Mobile Robot Motion: the Stream Field Method”, IEEE Int. Conf. on Robotics and Automation, 1994, pp. 378-386

S. Kirkpatrick, C. Gelatt and M. Vecchi, “Optimization by Simulated Annealing“, Science, Vol.220, No.4598, 1983

P. Khosla and R. Volpe, “Superquadric Artificial Potentials for Obstacle Avoidance and Approach”, IEEE Int. Conf. on Robotics & Automation, 1988

M. Maimone, A. Johnson et al., Autonomous Navigation Results from the Mars Exploration Rover (MER) Mission, in Experiment Robotics IX, Springer Tracts in Advanced Robotics, 2006, Volume 21/2006, 3-13, DOI: 10.1007/11552246_1

A. Massoud and M. Bayoumi, “Robot Navigation using the Vector Potential Approach”, IEEE Int. Conf. on Robotics & Automation, 1993

M. Masterton-Gibbons, An Introduction to Game-Theoretic Modelling, Student Mathematical Library Vol.II, American Mathematical Society, 2001.

M. Park and M. Lee, “Artificial Potential Field based Path Planning for Mobile Robots using a Virtual Obstacle Concept”, IEEE/ASME Intl. Conf. on Advanced Intelligence Mechatronics, 2003

M. Park and M. Lee, “Experimental Evaluation of Robot Path Planning by Artificial Potential Field approach with Simulated Annealing”, SICE 2002, Aug. 5-7, 2002, Osaka

R. Tilove, “Local Obstacle Avoidance for Mobile Robots based on the Method of Artificial Potentials”, IEEE Int. Conf. on Robotics & Automation, 1990

P. Vadakkepat, K. Cjen Tan and W. Ming-Liang, ”Evolutionary Artificial Potential Fields and their Application in Real Time Robot Path Planning”, IEEE Int. Conf. on Robotics & Automation, 2000

C. Warren, “Global Path Planning using Artificial Potential Fields”, IEEE Int. Conf. on Robotics & Automation, 1989

Y. Wang and G. Chirikjian, “A New Potential Field Method for Robot Path Planning”, IEEE Int. Conf. on Robotics & Automation, 2000