The Next Generation of Unmanned Ground Vehicles

Geoff Hester, Chris Smith, Pete Day and Antony Waldock
BAE Systems Advanced Technology Centre
Bristol, UK
antony.waldock@baesystems.com

Abstract

Over the last decade, the development of Unmanned Ground Vehicles (UGVs) has received significant attention through technology competitions, such as the DARPA Grand Challenges, where an unmanned vehicle autonomously navigated across a desert or through semi-urban roads. Although this marked a significant step forward in autonomous navigation, the current generation of UGVs are only capable of operating in controlled environments where the dynamics of a scene are well-understood. For example in both of the DARPA Grand Challenges, the participants were given detailed maps and the environment was carefully controlled during the competition. The next generation of UGVs will need to operate in uncontrolled environments where the terrain and infrastructure is uncertain and humans could be present. This paper discusses the challenges and current developments in the areas of sensing, localisation and planning to realise the next generation of UGVs.

1. Introduction

The last decade has seen a rapid increase in the deployment of Unmanned Ground Vehicles (UGVs) as a means of reducing the three "Ds"; the dull, dirty and dangerous tasks. For example, the mining company Rio-Tinto has recently announcement it will deploy one hundred and fifty autonomous trucks by the end of 2015[1]. To make today's UGVs viable in a wider set of applications and more complex or uncontrolled environments will require greater levels of autonomy. To date, the shift to higher levels of autonomy has been relatively cautious and heavily reliant on having a controlled and well-understood operating environment such as in a warehouse and a mine. However, UGVs capable of autonomously navigating over a wide range of uncontrolled environments, where the terrain and infrastructure are unknown, could offer significant benefits over today's generation of UGVs.

The ability to develop a UGV that is capable of operating autonomously, and, more specifically, navigating over a wide range of uncontrolled environments requires the development of technologies in a number of key areas. The remainder of this paper discusses three such areas; terrain sensing, localisation and motion planning, which are fundamental for any UGV. Firstly, to understand how these technologies fit together it is important to understand the overall architecture of any autonomous system.

2. Autonomous Systems Architecture

The use of a generic architecture (e.g. AURA [1], NIST 4D RCS [2], ESA Functional Reference Model) is beneficial to the development of an autonomous system by providing a framework that supports rapid experimentation and helps to ensure that best practice principles are adopted. The Six Box Model, shown in *Figure 1* is a conceptual architecture that is used to describe autonomous systems and place sub-system components and technologies in context within the overall system. The architecture identifies six functional areas into which autonomous systems architectural components can typically be placed.

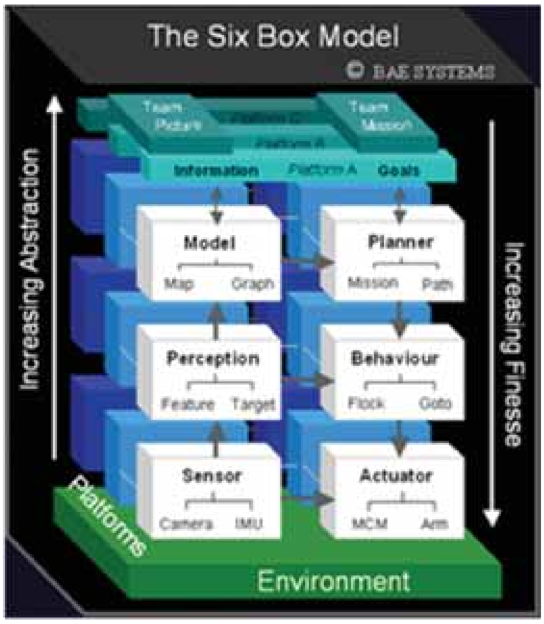

Figure 1: 6 Box Model, showing multiple vehicles

It should be noted that not all six boxes are required for an autonomous system; For example it is quite possible to create a system with no models or planners. However, components such as models and planners typically provide more cognitive ability, and hence higher levels of autonomy. The model is highly modular and encourages the use of different closed control loops through a system, for example steering control (sensor-actuator), waypoint following (sensor-perception-actuator) and exploration (sensor-perception-model-planner-behaviour-actuator). It can be thought of as an extension or recursion of the OODA loop developed by USAF Colonel John Boyd, where you first Observe (sense), then Orient (perceive/model), then Decide (plan/behave) and finally Act (actuate). The technology areas discussed within this paper are concerned with the modelling, perception and planning elements of the Six Box Model.

3. Modelling the World

Accurately modelling the environment around the vehicle using on-board sensors is a critical aspect for higher

[1] http://www.riotinto.com/documents/111102_Rio_Tinto_boosts_driverless_truck_fleet_to_150_under_Mine_of_the_Future_programme.pdf

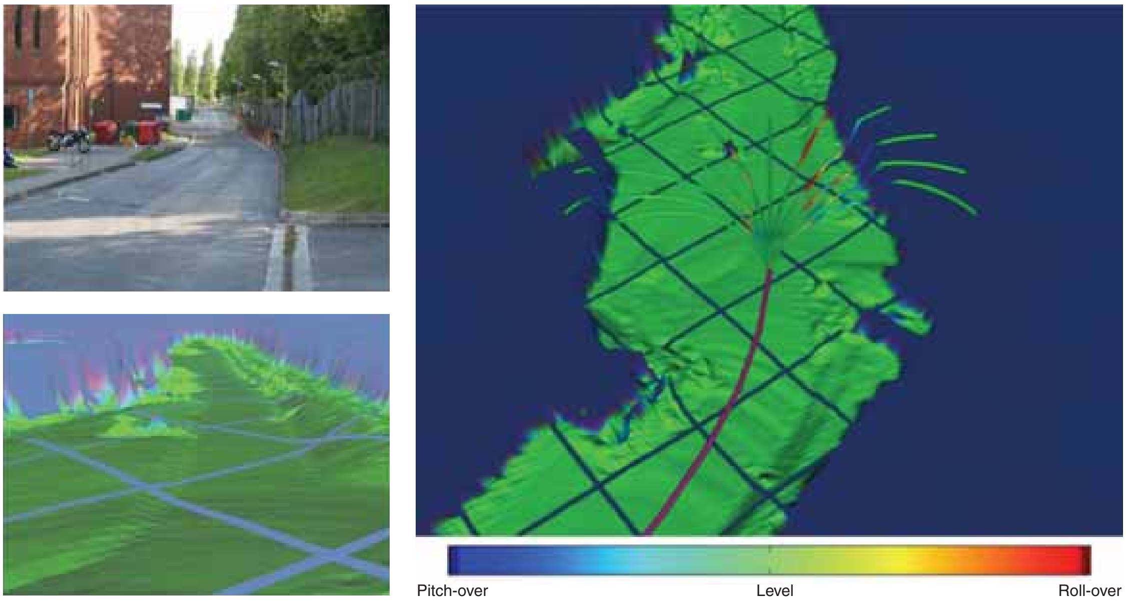

Figure 2: NURBS–KF modelling of surface geometry. Top left – Visual viewpoint of scene to be modelled. Bottom left – A NURBS–KF surface generated from Velodyne HDL-64 data captured from the same area. Right – Trajectories evaluated across the surface to determine potential pitch-over and roll-over conditions

levels of autonomy because it drives the on-board decision making processes and reduces the reliance on prior maps. In this section, we outline two key elements of modelling an unknown environment, estimating the terrain geometry and classifying the terrain type.

The estimation of the terrain geometry is the reconstruction of an estimate of the local surface topology. A key concern when operating in unknown environments is making sure that the vehicle is capable of safely navigating over different sections of the terrain. Therefore it is necessary to estimate the geometry of the surrounding terrain to be able to plan appropriately i.e. avoid excessive slopes or slow down for particular traverses. Most techniques are based on inferring geometry from point clouds derived from range measurement sensors. Traditional methods, such the well-known Delaunay Triangulation algorithm, are susceptible to error because they do not explicitly handle noise nor capture and represent how certain the estimate is. In particular, large (undetected) errors can be introduced in sparsely sampled regions due to a linear assumption about the surface between sample points. The handling of sensing errors is particularly important when the position of the platform between scans is uncertain due to positional estimates derived from an external infrastructure such as the Global Positioning System (GPS). To address this, we have developed a real-time system that automatically senses and constructs a probabilistic surface model of local terrain geometry from arbitrary range sensors (laser range finders, stereo cameras etc.) operating on-board a moving ground vehicle.

The current solution is based on Non-Uniform Rational B-Splines (NURBS) inspired by [9] who showed how NURBS could be combined with Kalman Filters (KF) such that model parameters can be estimated from direct surface observations. Constructing a NURBS surface is typically a manual, iterative, process whereby users manipulate control points, or some higher-level interface, through dedicated modelling tools. This approach however treats the z-position of each control point as a state variable. Under the KF formulation this uses a mean and covariance to represent each control point, where the covariance represents uncertainty. Projecting surface observations (i.e. range measurements) into the control space provides indirect observations of the underlying control points. These projected observations can then be used to update the control point estimate by standard KF methods. An example reconstruction from Velodyne HDL-64 data is shown in *Figure 2*.

In addition to understanding the terrain geometry, the terrain type is also a key characteristic that affects driveability. A series of methods have been investigated for automatic semantic labelling of terrain directly in front of a vehicle according to its type. The system automatically identifies the type of terrain ahead of a vehicle to provide inputs into the reasoning engine. The system is based on [10] which was deployed in the DARPA Grand Challenge. That method however aimed to detect similar terrain without attempting to provide a semantic labelling. By introducing labels and using a multi-class classification scheme we can successfully differentiate between different terrains. One advantage of using a multi-model approach is that the system can identify regions that have a mixed (and thus inconclusive) appearance. In these circumstances it is concluded that the region is mixed and it is labelled as such rather than forcing it to be a specific class as per *Figure 3*.

These two methods for estimating the terrain geometry and type have been demonstrated on diverse terrains from urban (as shown in *Figure 2*) to off-road to enable detection of possible roll-overs (as shown by the colour coding of the paths in *Figure 2* and discussed later in this paper).

4. Localisation

The second key technology area addressed in this paper to enable navigation over uncontrolled environments is localisation. Localisation is the problem of estimating the motion and position of the platform and, within the robotics community, there has been a great deal of effort expended on trying to construct a perfect (or near-perfect) localisation sub-system.

The availability of a perfect, drift free and jump free localisation system is a key assumption made in many autonomous systems. Given that localisation is a core input to the processes of information fusion (including

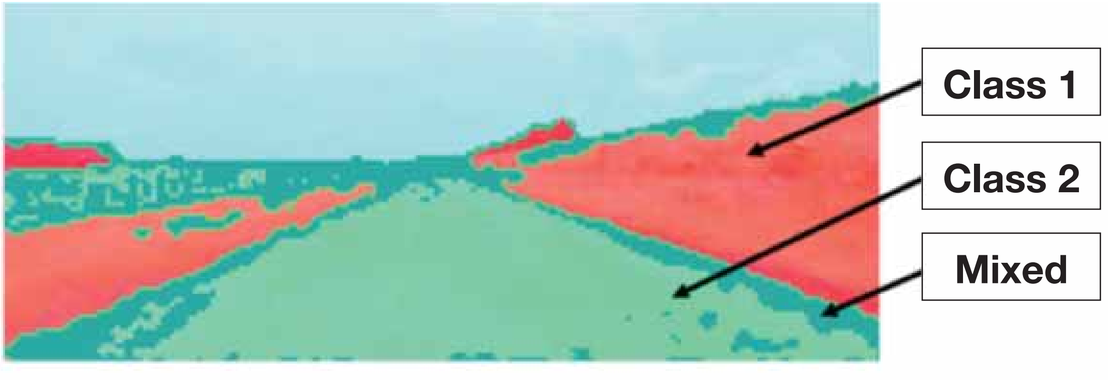

Figure 3: Mixed Class Labelling Example

the terrain estimation shown in the previous section), planning and control, the performance of a platform's localisation system is thus tightly coupled to the overall system performance. Within the Defence and Security sectors most full-scale autonomous ground vehicles utilise Inertial Navigation Systems (INS) from companies such as Applanix, Oxford Technical Solutions or Novatel. These systems typically cost tens of thousands of pounds per platform, are reliant on adequate satellite coverage, and are fragile to GPS dropouts and denial.

In contrast to other approaches that fuse additional sensors such as wheel odometry or build a new co-ordinate frame such as Simultaneous Localisation And Mapping (SLAM), a different approach is suggested. An autonomous navigation solution has been architected to use a system level approach that supports multiple distinct co-ordinate systems which allow constant drift, and collectively describe both position and motion. These co-ordinate systems are combined by dynamic, uncertain transforms which allow significantly more flexibility in introducing new localisation systems. The architecture has been evaluated using a series of localisation systems including traditional inertial navigation systems, visual place recognition [3], cellular transmitter tracking [4], visually stabilised inertial [5], visual odometry [6] and Inertial Measurement Unit/SLAM hybrids [7]. A selection of these systems allows for cost-effective, indefinite-range autonomous navigation in GPS-intermittent and denied environments.

The use of multiple independent frames allows the system designer to de-couple and simultaneously optimise different aspects of the control system to various operating points without requiring that localisation performance be sufficient for all uses. This means that different categories of system element (e.g. planning and control) can use the frame most suited to their own problem. For example, a collision avoidance behaviour rarely requires knowledge of the global geodetic position since its purpose is to avoid locally sensed obstacles. Similarly a mission planning element would benefit from accurate geodetic information at key decision points i.e. at a road junction or on approach to its final destination. The independent frames are optimised into two main categories. The first category represents the dynamics of the platform motion-optimised in a frame with an origin that drifts constantly compared to geodetic. The motion optimised frames can be driven by localisation technology such as wheel odometry, visual odometry or SLAM without loop closure. The second category represents the position and orientation of the platform position-optimised in a frame that is tied to geodetic but updates sporadically. There can be multiple position optimised frames including topological systems [3], each representing a common reference (e.g. zip code, or GPS/WGS84).

Within the motion optimised frames, functions such as sensor fusion, obstacle avoidance, and vehicle control are operated without any need to know the geodetic location or correctly determine if the vehicle has returned to a prior location. Within each position

Figure 4: Integration of Candidate 1 (Left) and Candidate 2 (Right) sub-system into Multi-frame Autonomous Navigation Architecture provides GPS denied operation of LiDAR registration system over 40 times longer, and at one tenth the cost, of baseline INS system

optimised frame, route and mission planning type tasks can be performed as and when data are available and accurate, and are then transformed into, and cached in, the motion optimised frames. *Figure 4* shows the output of the terrain sensing registration system operating within a multiple navigation frame based control system on the BAE Systems autonomous systems test bed at speeds of up to 40 MPH. At the time of capture the low cost localisation system is still operating terrain sensing, path planning and obstacle avoidance functions after 35 minutes of GPS denial.

Whilst utilising multiple navigation frames removes the reliance of the overall system on a single "all purpose" localisation source there is nonetheless the need for sub-system(s) capable of providing an estimate of platform position and platform motion. Hence, a key challenge to enable autonomous navigation for the next generation of UGVs is to extend the availability of positioning systems by reducing or removing their reliance on GPS. Our approach is not simply to replace GPS but to use a myriad of alternative positioning sensors that can either augment GPS or provide localisation information appropriate to the problem of each individual navigation frame.

For example, BAE Systems has been working closely with the University of Oxford, to develop and evaluate how visual features in a set of camera images can be used as a positional sensor to aid the availability and accuracy of localisation information in both a global and local frame of reference. For example, in the global reference frame – recognising a road junction in an image with a known location or in the local reference frame – calculating the motion of the vehicle between camera frames. These technologies have been demonstrated in the civil sector and appear to show significant promise with the recent acquisition of Plink Art Ltd by Google after only 6 months in existence.

Another method of extracting positional information to lessen the impact of GPS dropouts has exploited radio positioning information from both cooperative beacons and non-cooperative transmitter sources [4]. The low signal strength of GPS leaves it highly susceptible to jamming and also results in intermittent or denied service indoors and in canyons. The higher powers of terrestrial signals allow radio positioning in challenging GPS-denied environments, and when combined with the covert nature of using opportunistic signals also reduces the likelihood of jamming. A global positioning system has been developed based on opportunistic radio sources to rival GPS and other satellite-based positioning systems. At the core are a series of simultaneous localisation and mapping algorithms developed specifically for indoor and outdoor scenarios. These exploit the specific signal features available in these environments, and can incorporate measurements from digital television, digital radio, 3G cellular, WI-Fi, Wi-MAX, and the future 4G wireless protocols all of which provide a greater fundamental positioning accuracy than GPS when the infrastructure is available.

Localisation is a critical technology area as it typically drives the rest of the system performance. In the next-generation of UGVs, it will be essential to deliver multiple localisation feeds derived from a myriad of sources to be able to cope reliably and gracefully with unknown environments.

5. Motion Planning

The final technology area addressed in this paper is motion planning. Motion or path planning for a UGV is an exercise in the real-time balance of safety, efficiency and speed in a world that is constrained, changing, and uncertain. For operation in GPS denied environments the planning system must work in multiple frames, with the entities in the surrounding world and goal locations being described with either relative or absolute positions. Motion planning on roads demonstrates this dual localisation requirement. Whilst an UGV may have access to a map which details the geodetic locations of the road network, requiring that this map exists at sufficient accuracy to enable autonomous operation drastically reduces the usability of the system. With a relatively small amount of registration error for a road, a system which simply tracks the road centre line could be caused to drive on the pavement or into oncoming traffic. A much better solution is to locally sense the road and rely on a motion planning system to drive down the centre of the lane.

Current motion planning techniques are, in the general case, a formulation of a constrained optimisation problem. This can take the form of finding a trajectory which minimises some cost function, when evaluated against the model of the world, and will result in the vehicle transitioning between the start location and some specified goal location, whilst not violating the kinematic or dynamic constraints of the vehicle. This problem is then extended to include additional constraints, such that the generated trajectory should not result in the vehicle colliding with any perceived obstacle.

The motion planning problem described is non-convex with potentially an infinite number of feasible solutions and numerous local minima. A system was developed, based on the approach of [8], which uses a discretised state space to reduce this complexity. Representing each state in this discrete space as a vertex, a graph is then generated such that each edge represented a kinodynamically feasible trajectory. Using this graph representation a standard shortest path graph search algorithm can then be used to find the optimal set of trajectory-edges which reaches the goal state.

When determining the quality of a candidate trajectory a popular approach is to model the expected behaviour of the vehicle executing the trajectory and assigning some concept of incurred cost for performing that action. As part of the research into terrain classification and modelling, an algorithm was developed to detect static roll-over using the planned state of the vehicle and the surface normal at that state. From this, the predicted pitch and roll can be calculated and compared to operational limits. A cost function can then be developed which associates lower cost with trajectories that do not include high pitch or roll angles, causing the planning optimisation to prefer these trajectories. Figure 2 shows several candidate trajectories evaluated for static rollover against a terrain surface model in real-time. With the increasing maturity of autonomous systems, capable of perceiving their environment and deciding on a course of action to complete a specified goal, much attention has been focused on incorporating these systems into human environments. For land vehicles this means enabling the system to operate in the presence of other vehicles, many of which may be being driven by a human, as well as obeying the rules of the road.

Given the nature of optimisation based planning algorithms, with the desire to minimise the cost incurred by the planned trajectory, generated plans can often appear unintuitive. When avoiding a collision the planning system may choose to avoid the obstacle by

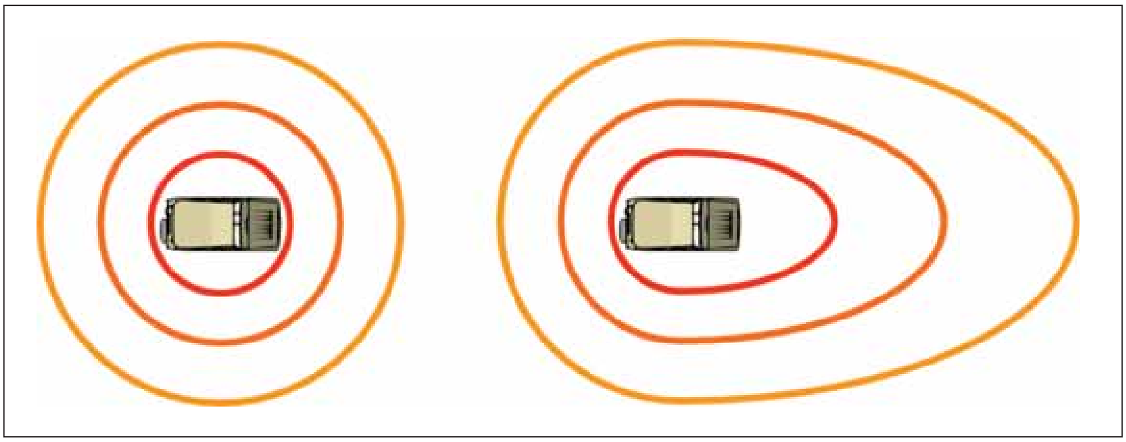

Figure 5: Modified cost field to dissuade AGV from planning to drive in front of other vehicles

the minimum required distance. In the case of a road junction these approaches may produce plans which result in the UGV pulling out in front of traffic as the system has found a collision free trajectory. Whilst, this plan would not require the other vehicles to take any avoiding action this aggressive driving style would be unacceptable to many human drivers. The planning system must now ensure that generated plans comply with the rules of the road and result in the vehicle being perceived as a considerate by other drivers.

Two approaches were developed to achieve this requirement. The first approach inserts additional intermediary goal locations as constraints, with the effect that the generated plan must pass through these points. By selecting these points appropriately, compliance with the rules of the road can be achieved. The second approach modifies the shape of the cost profile surrounding other vehicles sensed in the world. Instead of just a region resulting in a collision incurring high cost, the area in front of the vehicle was also assigned higher cost values. This caused any trajectory which passed through these regions to be less likely to be selected for the final plan. *Figure 5* shows the difference in the cost field surrounding another vehicle. In the left hand image a radial avoidance cost field is shown. The modified increases the cost of the region in front of the vehicle making it preferable for the motion planning system to pass behind.

6. Summary

This paper has introduced and discussed three key technology areas that are being addressed in order to realise the next generation of unmanned ground vehicles. The requirement to cope gracefully with uncontrolled environments has many potential benefits in a wide number of domains. The technology areas are not likely to deliver simultaneously but as these technologies mature we will start to see the greater deployment of UGVs, first in isolation but then alongside manned vehicles.

7. Acknowledgement

The work presented in this paper has been conducted by the Autonomous Systems team at the BAE Systems Advanced Technology Centre with support from a range of partners and funded, in part, by the Systems Engineering for Autonomous Systems (SEAS) Defence Technology Centre (DTC).

8. References

[1] R, Arkin, T. Balch, "AuRA: Principles and Practice in Review", Journal of Experimental and Theoretical Artificial Intelligence, Vol. 9, No. 2, 1997, pp. 175-189
[2] RCS: The Real-time Control Systems Architecture, http://www.nist.gov/el/isd/rcs.cfm, accessed 01/05/2011
[3] M. Cummins, P. Newman, "FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance", The International Journal of Robotics Research 27(6), 647-665
[4] R. Faragher, "The effects of multipath interference on radio positioning systems", Thesis, University of Cambridge, 2007
[5] A.I. Mourikis, S.I. Roumeliotis: "A Multi-state Constraint Kalman Filter for Vision-Aided Inertial Navigation", Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2007, 3565-3572.
[6] C. Mei, G. Sibley, M. Cummins, P. Newman, I. Reid, "RSLAM: A System for Large-Scale Mapping in Constant-Time using Stereo", International Journal of Computer Vision, 2010
[7] L. Connolly, S. Scheding, "Comparative Analysis of Sensors, Algorithms, and Models within a Navigation System", Proceedings of IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, 2008
[8] T.Howard, C.Green, A.Kelly and D.Ferguson "State Space Sampling of Feasible Motions for High Performance Mobile Robot Navigation in Complex Environments", Journal of Field Robotics 25(6-7), 325-345 (2008)
[9] Y. Huang and X. Qian, "Dynamic B-spline surface reconstruction: Closing the sensing-and-modeling loop in 3D digitization", Computer-Aided Design Volume 39, Issue 11, November 2007, Pages 987-1002
[10] S. Thrun et al, "Stanley: The Robot that Won the DARPA Grand Challenge", Journal of Field Robotics, 2006
[11] R. Simpson, J. Cullip and J. Revell, "The BAE Systems Wildcat Data Sets", 5th SEAS DTC Technical Conference, 2010