

Experiments in Autonomous Earth Moving

Adrian Bonchis* Nicholas Hillier* Julian Ryde** Elliot Duff*
Cédric Pradalier***

* *Autonomous Systems Laboratory, CSIRO, Brisbane, Australia*
(e-mail: name.surname@csiro.au).

** *University of Buffalo, Buffalo, NY 14203 USA*
(e-mail: jryde@buffalo.edu)

*** *Autonomous Systems Laboratory, ETH Zürich, Zürich, Switzerland*
(e-mail: cedric.pradalier@mavt.ethz.ch).

Abstract: This paper presents a technology demonstrator currently under development and describes experiments carried out to date in autonomous bulk material handling using mobile equipment. Our primary platform is a Bobcat S185 skid-steer loader instrumented with an onboard computer, a sensor suite, and a communication link that support various levels of automation, from remote control to supervised autonomy. We present the main system components and discuss the autonomous cleaning of spillage and carryback, common bulk handling task in mining, currently executed exclusively using manually and/or remotely operated loaders. The system architecture is based on Spring, a Robotics Software Framework developed by CSIRO to support rapid development of new robotic systems, distributed as an Open Source package.

Keywords: Autonomous mobile robots, control applications, software tools, scene segmentation.

1. INTRODUCTION

Since the early 1970's the Australian mining industry has grown at an enormous rate in the area of production and technology, but the extraction of minerals remains vastly a mechanised material handling exercise. Nevertheless, the last few years have seen an expansion of robotic and autonomous systems in various surface and underground operations, driven mainly by three key factors.

- Safety: although safety performance over the last decade has improved, mining is still an industry which has a comparatively high death rate (Minerals Council of Australia, 2006).
- Productivity gains: mining is a capital intensive proposition, and improving machine utilization, reducing machine shutdowns due to faults caused by human operators, lead to significant economic benefits.
- Workforce shortage: this is a serious concern in Australia. The National Industry Skills Council predicts that despite slowdown in activity caused by the global financial crisis, an ageing workforce will have a significant impact on the trained workforce in the resources industry over the next five years (SkillsDMC, 2010).

In this context CSIRO is actively pursuing research focused on automating typical mining tasks, such as excavation, loading, haulage and dumping. In particular, the Autonomous Systems Laboratory in CSIRO is focusing on common bulk handling problems such as cleaning spillage and carryback. Traditionally, spillage and carryback material are removed using manually and/or remotely operated loaders. Operating a loader in conditions that are

dynamic, with hazards that can vary and can be subtle and difficult to recognize, adds to the complexity of the tasks the workers perform on the machine. Many of these tasks require high levels of skills in ground engaging tool control, machine positioning, judgement and decision making. In addition, operators are exposed to a variety of risk factors that may lead to health problems, some of the most common being: whole-body vibration, awkward postural requirements, dust, noise, temperature extremes and shift work.

This paper presents a technology demonstrator currently under development and describes experiments carried out to date in autonomous bulk material handling at CSIRO. Firstly, in Section 2, we review some state-of-the-art research results in autonomous earth moving. In Section 3 we introduce the robotic system architecture implemented on our test vehicle, and discuss specific hardware and software aspects. Subsequently, we discuss the the main system modules implemented to date: the vehicle and ground engaging tool controllers in Section 5 and the situational awareness module in Section 6, where we present results obtained in segmenting the dig and dump scenes. Finally, in Section 7 we future development work.

2. RELATED RESEARCH

Research in the area of autonomous earth moving appears to be gravitating around systems (including systems of systems), traditional robotics topics (sensing, planning, control), and tool-ground interaction. A framework for an intelligent earthwork system is suggested by Kim and Russell (2003). The authors discuss factors that can affect

earthwork operation performance, identify key emerging technologies to support the implementation, the system architecture, and the system control strategy. A comprehensive review by Singh (2002) examined various aspects of sensing, planning and control as they apply to earth-moving automation.

The work described in our paper can be classified as scooping. The wheels and/or tracks of excavators and backhoes are stationary during one excavation cycle (dig-dump), whereas in scooping, the vehicle needs to move forward during the digging phase. The vehicles most commonly used in mining to perform scooping include load-haul-dump (LHD) vehicles, wheel loaders, and dozers. The requirements for robotic scooping using an LHD have been analysed in detail by Hemami (1994). Experimental results can be found in the work of Hemami (1993), Petty (1997), Sarata et al. (2008) and Marshall et al. (2008). Invariably, all researchers agree that the biggest challenge to robotic excavation is the interaction between the tool and the terrain (ground, pile). This interaction is shaped by the properties of the media (e.g., density and hardness), the rock pile geometry, and the distribution of particle sizes and shapes. The problem is compounded by the fact that it is difficult to predetermine the exact nature of interactions prior to the execution of any particular excavation operation.

Significant effort was dedicated to modelling the interaction between the ground engaging tool and the terrain in excavation. The classical fundamental equation of earth-moving (Reece, 1964) used to determine the forces required for digging will fail to predict these forces if the soil parameters are not estimated correctly. Luengo et al. (1998) for example developed a numerical scheme that predicts resistive forces experienced at the tool during digging based on a reformulated version of the fundamental equation of earthmoving. Soil parameters are estimated from a small number of digs and used to predict forces for candidate digs in the future. Tan et al. (2005) estimate soil parameters using a modified Newton-Raphson algorithm that minimizes the error between the measured failure forces and the failure forces predicted using a Mohr-Coulomb soil model.

To avoid the explicit use of the fundamental equation of earthmoving in autonomous excavation, some researchers have turned to fuzzy logic and neural networks for control. Variations of this approach are presented for example in the work of Singh (1995), Shi et al. (1996), and Wang (2004). Other approaches treat the reactive forces in digging as external disturbances. Ha and Rye (2002) for example implement a variable structure systems methodology to the control of a robotic excavator and assume that the disturbance force is slowly time-varying.

3. SYSTEM ARCHITECTURE

Our primary platform is a Bobcat S185 skid-steer loader equipped with a Bob-Tach mounting system that enables us to switch between a combination bucket and a blade dozer. The vehicle, shown in Figure 1, is instrumented with an onboard computer, a sensor suite, and a communication link that support various levels of automation, from remote

control to supervised autonomy. The automation system architecture is described next.



Fig. 1. The robotic Bobcat developed by CSIRO.

3.1 Hardware System

The following hardware systems were retrofitted:

- Machine Control PC
- Computer Vision PC
- Pilz Safety System
- Advantys Distributed I/O
- Hetronic radio remote control
- Cisco Cannonball Wifi
- Throttle actuator.
- Vehicle and tool position and attitude sensor suite

The signals controlling the vehicle and ground engaging tool can originate from three main systems: the factory-default machine control unit, the remote radio control unit, and the machine control computer. These signals are fed into the system through two main interfaces, a Pilz PLC and an Advantys Distributed I/O device which is configured with digital I/O, analogue I/O and counters. Safety is handled by the Pilz programmable safety system, which is configured to monitor and control a number of E-STOP pushbuttons, the operating mode selector switches, the safety-bar safety switch, the park brake and the forward-reverse maximum speeds. The Advantys I/O system also incorporates a safety function, and turns the commanded values to 0 after 10ms if no signals are received from the machine control PC. This will force the machine to come to a stop if the PC hangs while controlling the machine. The Computer Vision PC is currently not used, but was installed to support further research in enhanced situational awareness and visual servoing for autonomous digging.

The Bobcat automation system operates in three main modes: Manual, Remote, and Auto. In Manual mode, we use the factory-fitted in-cab controls for operating the vehicle and tool. These control signals are available to the retrofitted machine control PC, such that control sequences performed by a human operator can be recorded if desired. In Remote mode, we operate the vehicle and tool using the Hetronics radio remote control transmitter-receiver unit. These signals are also fed to the machine

control PC. The Auto mode refers to operation in closed-loop under computer control; the machine is driven from the signals generated by the CSIRO machine control PC through the Advantys Distributed IO module.

The vehicle and tool position and attitude sensor suite consists of:

- Novatel SPAN-CPT GPS/INS receiver
- Spinning Laser
- Boom tilt sensor
- Bucket tilt sensor

We use a high-end Novatel SPAN-CPT GPS/INS operating in RTK mode to provide vehicle position and attitude estimates. The receiver incorporates in the same enclosure a OEMV-3 GPS receiver with a 6DoF IMU manufactured by KVH Industries. The SPAN engine developed by Novatel provides a tightly coupled GPS/INS solution which can propagate a vehicle pose estimate during short periods of GPS outage while also providing a faster satellite reacquisition and faster RTK initialization after outages. Workspace mapping and situational awareness information is provided by a SICK LMS 291 laser range finder mounted on a spinning platform. In contrast with a camera based system, the laser is not affected by lighting conditions, while the narrow beam-width, high accuracy and angular resolution makes it preferable over other ranging sensors such as radar and sonar. Finally, the tool pose is retrieved using information provided by two string potentiometers.

3.2 Low Level Software Loops

The low level loops in the system interface with the various sensors and actuators using Spring, the robotics software framework developed by the Autonomous Systems Laboratory at CSIRO, available as an Open Source (CSIRO, 2007) package. At the heart of Spring is DDX (Dynamic Data eXchange) which is our core platform for building distributed robot controllers (Corke et al., 2004).

DDX uses a publish-subscribe communication system to transfer data between processes running on networked computers. The system consists of two main programs, the catalog and the store. A typical system consists of a collection of stores running on several hosts (one per host) all communicating with a catalog. The store is used as a data repository (via shared memory) and the catalog keeps track of the contents of each store (via UDP/multicast). Each client interacts with the local store only which stores all data in shared memory. Data between hosts is synchronized by the stores using UDP/multicast, thus minimizing the number of copies of the data in the system. This relationship is shown in Figure 2.

Spring includes standard tools to interface with various GPS receivers, SICK laser scanners, Riegl laser scanners and a variety of I/O devices, including Modbus/TCP. Spring also supports various Firewire cameras as well as V4L. These tools allow the devices to be used in the system without requiring any extra code be written. The toolkit supports several languages: C, C++, Java and Python. Support for Matlab is also provided, but is optional. Standard tools are provided to query the store and the catalog for diagnostic purposes.

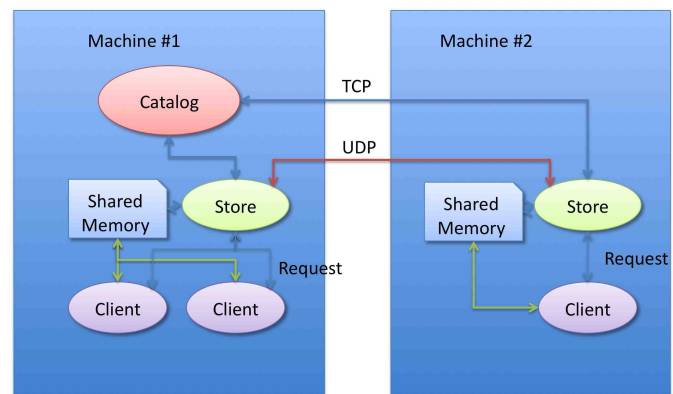


Fig. 2. The Dynamic Data eXchange mechanism at the core of the CSIRO Bobcat Automation System.

For details, please refer to the documentation in CSIRO (2007).

4. BOBCAT SIMULATOR

Field robotics is a difficult research domain that poses significant challenges in development and operation due to the unstructured environments (both geometrically, and in material properties), common use of multiple sensing modes and the high degrees of freedom inherent in most field robotics platforms. Additionally, there are pragmatic difficulties associated with the development of field robotics platforms stemming from long set-up and shut-down times (pre and post-start checks) and access to suitable test sites that do not occur with most indoor and small platform robotics. In particular, the use of the Bobcat as a research platform has significant operational safety risks induced by high turn rates, large actuation forces and high moving mass.

The above considerations have led us to investigate some of the tools available for simulating mobile robotic systems and adopt one suitable to our needs (Hillier and Ryde, 2010). Amongst the numerous offerings available, we identified the OpenRAVE platform as the simulation tool of choice, for the following reasons:

- it offers simple and flexible interfaces in a variety of languages (C++, Octave/Matlab and Python);
- it provides a dynamically loadable plug-in type architecture that allows for significant customisation of very low-level interfaces (such as the employed physics engine and collision checking) and ease of integration into middleware;
- platform independence (support for Windows, MacOS and Linux) with the transparency of open source code.
- it provides the user with the ability to dynamically create, destroy and change the geometric properties of items in the simulation environment. This is a particularly powerful option for operations that manipulate the environment.

Furthermore, the OpenRAVE environment allows for environment cloning, a powerful feature that could be exploited for more advanced simulation tasks such as parallel




Fig. 3. A simulator screenshot showing the simulated vehicle, the spinning sensing plane of the laser rangefinder, and the bulk material as discrete cubes.


execution for run-time planning tasks, although we are yet to exploit this feature.

The first significant test case scenario involved autonomously moving material from one location to another, and executing this task in both reality and the simulation environment. The simulation of earthen material is typically not an implemented feature of robotics simulators, and whilst work on accurately modelling earthen material in simulation for robotics has been conducted (Halbach, 2007), we chose to use a number of 0.1m sided cubes as an approximate representation, as shown in Figure 3.


5. VEHICLE AND GROUND ENGAGING TOOL CONTROL SYSTEMS

The Bobcat is a short base wheel skid-steered vehicle, steering is achieved by controlling the relative velocities of the left and right side wheels. Considerable research exists in the area of modelling and control of wheeled and tracked skid-steered vehicles. The main problems faced by automatic steering controllers for this type of steering stem from slippage that occurs during turning, and in the case of off-road vehicles, the complex interaction between wheels and terrain. Slippage occurs as a result of the rigid alignment between the wheels and the longitudinal axis of the vehicle.


Our approach in dealing with the vehicle control problem was to start from simple control schemes and evolve to more complex control scheme later. In doing so, we were able to concentrate on the higher level tasks on the actual operation that we are trying to automate, namely moving a pile of dirt from location A to location B. From a control perspective, the lateral slippage manifests itself in a coupling between the lateral and longitudinal vehicle dynamics. Our simplified vehicle controller ignores this coupling and the speed and steering loops are implemented as independent PID loops. We operate off-road, and the wheel-terrain interaction is affected by the soil compactness and humidity. We are relying in a first instance on the robustness of the PID controllers to achieve acceptable control errors. In Figure 4 we present the results of a ten-trip run between the dig and dump locations.



(a) Track on ground.



(b) Heading control actual vs. demand.



(c) Speed control actual vs. demand.

Fig. 4. Experimental results from ten dig-to-dump trips.

The kinematic structure of the bucket mechanism is shown in Figure 5. Relative to the vehicle chassis, this mechanism can be seen as an open-chain mechanism, consisting on two independent sub-mechanism, each implementing one degree of freedom, up-down (lift) and crowd-

dump. The lift mechanism consists of two independent closed chains. Three such combinations can be formed: 0,1,2,3,4,0, 0,1,2,5,0 and 0,4,3,2,5,0. The crowd-dump is delivered by the chain 2,6,7,8,2), with link 2 being seen as the “fix” element and link 8 being the bucket (end-effector). The mechanism is actuated by two single-ended hydraulic cylinders, which are in turn controlled by two independent PID loops. A simple point-to-point path generator is sufficient for prescribing the set points for the two loops. Experimental results are shown in Fig. 6





Fig. 5. The kinematic structure of the ground engaging tool mechanism on a Bobcat S185.

6. SITUATIONAL AWARENESS


From a situational awareness perspective, the main problems we are facing are that of segmenting the environment and identifying the dig and dump regions, as well as detecting and avoiding obstacles in the workspace. The more recent work of Marshall et al. (2008) and Schmidt et al. (2010) present a number of approaches to the dig and dump planning processes. Such approaches utilise a 2.5D (height-field) representation of the environment constructed from a registered point cloud and compare with a desired surface profile to obtain the material to be removed in the dig process. The approach we used for segmentation is analogous to background subtraction (BGS) discussed in Chalidabhongse et al. (2003). A “background” digital terrain map (the base scene) is stored in a *base* map and through an alignment process, the differences between this background surface and the current terrain map are highlighted. For digging, the background is subtracted from the current surface profile to identify the material to be removed. A similar approach is used for identifying the dump point at the dump location. Quite often, the dumping is constrained to a free-standing bin or the tray of a haul truck.

In our experiments we aim to move a pile of bulk material from one location to another. We broke the task into three main sub-tasks: digging, hauling and dumping into a 3 sided bin. As part of the digging process, the pile must be segmented from the surrounding environment in order to determine a desired digging pose for the Bobcat. Similarly, dumping in the bin requires that we segment any existing pile in the bin in order to determine a desired dump pose for the Bobcat.

We illustrate in the following the segmentation process we implemented at the dump location. The process consists of the following phases:



(a) Up-down control loop.



(b) Crowd-dump control loop.

Fig. 6. Bucket control (experimental results).

- (1) create and store a base map of the empty enclosure;
- (2) adopt a pose approximately in front of the enclosure;
- (3) acquire a 3D scan using the rotating laser;
- (4) align the scan with the base map;
- (5) remove scan points coincident with the base map;
- (6) analyse the remaining points (the segmented data of the pile) and determine the dig pose.

To generate the base map of the empty enclosure, the Bobcat[®] was tele-operated to various poses around and in front of the empty material enclosure. It was paused at each pose for a short period in order to get an accurate 3D point cloud from the rotating laser. Adoption of a number of variant poses, including some from behind the enclosure, although not entirely necessary, help in making the base map more complete. This in turn improves the efficacy of the segmentation subtraction. For enhanced accuracy the 3D scans used for base map generation were only extracted when the sensors report that the vehicle is stationary. These scans were iteratively aligned to the map and co-registered to produce an empty enclosure base map, which was then saved.

During the course of its missions, the Bobcat adopts a pose in front of the material enclosure and obtains 3D range

data of the enclosure and any contained material. This scan is aligned with the map as with the base map creation, using the INS pose as an initial guess, to initiate the segmentation process. The alignment and segmentation approach we employed is described in full in Ryde and Hillier (2011).

7. FURTHER WORK

The application described in the paper is work in progress, and experiments carried out to date have revealed a number of avenues to improve the system performance, the vast majority of them focusing on high level planning processes. At the dig location, we are investigating methods that optimise the volume of material loaded in the bucket, such that we avoid under- and over-loading the machine. At the dump location, we are working towards including a collision predictor in the dump planning phase, and minimising the amount of material dumped outside the bin. We also continue to refine our scene segmentation. Since we are focusing on cleaning spillage and carryback material, it is extremely important that we remove the dependency on GPS for localising the vehicle. To this extent, we are currently working towards incorporating SLAM-based and RF localisation technologies developed within CSIRO. In parallel, we are trying to further improve the control performance, and this this extent, we are investigating alternative robust control paradigms.

ACKNOWLEDGEMENTS

We would like to thank and acknowledge the contribution of our colleague Paul Flick, who designed and built the hardware retrofit system, of former colleagues Polly Alexander for her work in developing the low level controller software in Spring, Leon Stepan and Hendrik Erckens for their assistance with the implementation of the control systems and field trials for the work presented. We also acknowledge the project funding provided by the CSIRO's Minerals Down Under Flagship.

REFERENCES

- Chalidabhongse, T.H., Kim, K., Harwood, D., and Davis, L. (2003). A perturbation method for evaluating background subtraction algorithms. In *Joint IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance (VS-PETS 2003)*. Nice, France.
- Corke, P., Sikka, P., Roberts, J., and Duff, E. (2004). DDX: A distributed architecture for robot control. In *Proc. Australian Conference on Robotics and Automation*. Canberra, Australia.
- CSIRO (2007). Spring – Robotics Software Framework. <http://research.ict.csiro.au/research/labs/autonomous-systems/field-robotics/robotics-software-framework>. Accessed 22 September 2010.
- Ha, Q.P. and Rye, D.C. (2002). Robotic excavation in construction automation. *IEEE Robotics and Automation Magazine*, 9(1), 20–28.
- Halbach, E. (2007). *Development of a simulator for modeling robotic earth-moving tasks*. Master's thesis, Lule University of Technology.
- Hemami, A. (1993). Study of bucket trajectory in automatic scooping with load haul dump loaders. *Trans. Institution of Mining and Metallurgy Section A, Mining Industry*, 102, A37–A42.
- Hemami, A. (1994). Modeling, analysis and preliminary studies for automatic scooping. *Advanced Robotics*, 8, 511–529.
- Hillier, N. and Ryde, J. (2010). On the use of simulation for dependable field robotics. In *Proc. 19th Int. Symp on Mine Planning and Equipment Selection (MPES 2010)*. Freemantle, Australia.
- Kim, S. and Russell, J.S. (2003). Framework for an intelligent earthwork system. part i. system architecture. *Automation in Construction*, 12(1), 1–13.
- Luengo, O., Singh, S., and Cannon, H. (1998). Modeling and identification of soil-tool interaction in automated excavation. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 1900–1906. Victoria B.C., Canada.
- Marshall, J.A., Murphy, P.F., and Daneshmend, L.K. (2008). Toward autonomous excavation of fragmented rock: full-scale experiments. *IEEE Trans on Automation Science and Engineering*, 5(3), 562–567.
- Minerals Council of Australia (2006). The mobile equipment incident causation survey 2005-2006. Technical report, Minerals Council of Australia.
- Petty, M.K. (1997). Autonomous LHD loading. In *Proc. 4th Annual Conference on Mechatronics and Machine Vision in Practice*, 219–224. Toowoomba, Australia.
- Reece, A.R. (1964). The fundamental equation of earth-moving mechanics. *Proc. of the Institution of Mechanical Engineers*.
- Ryde, J. and Hillier, N. (2011). Alignment and 3d scene change detection for segmentation in autonomous earth moving. Submitted to the 2011 IEEE International Conference on Robotics and Automation (ICRA 2011).
- Sarata, S., Koyachi, N., and Sugawara, K. (2008). Field test of autonomous loading operation by wheel loader. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2661–2666. Nice, France.
- Schmidt, D., Proetzsch, M., and Berns, K. (2010). Simulation and control of an autonomous bucket excavator for landscaping tasks. In *Proc. IEEE Int. Conf. Robotics and Automation*, 5108–5113. Anchorage, AK.
- Shi, X.B., Wang, F.Y., and Lever, P.J.A. (1996). Experimental results of robotic excavation using fuzzy behavior control. *Control Engineering Practice*, 4(2), 145–152.
- Singh, S. (1995). Learning to predict resistive forces during robotic excavation. In *Proc. IEEE Int. Conf. Robotics and Automation*, 2102–2107. Nagoya, Japan.
- Singh, S. (2002). State of the art in automation of earthmoving, 2002. In *Proceedings of the Workshop on Advanced Geomechanics*.
- SkillsDMC (2010). Environmental scan 2010. Technical report, The Australian National Industry Skills Council.
- Tan, C., Zweiri, Y.H., Althoefer, K., and Seneviratne, L.D. (2005). Online soil-bucket interaction identification for autonomous excavation. In *Proc. IEEE Int. Conf. Robotics and Automation*, 3576–3581. Barcelona, Spain.
- Wang, F.Y. (2004). Agent-based control for fuzzy behavior programming in robotic excavation. *IEEE Trans. Fuzzy Systems*, 12(4), 540–548.