

Real-time Control of Soccer-playing Robots Using Off-board Vision: the Dynamite Testbed

Michael K. Sahota, Alan K. Mackworth, Stewart J. Kingdon and Rod A. Barman

Laboratory for Computational Intelligence
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada, V6T 1Z4
dynamo@cs.ubc.ca

ABSTRACT

The Dynamite testbed has been constructed for experiments with autonomous mobile robots. It consists of a fleet of radio-controlled vehicles that receive commands from a remote computer. All the robots share an off-board vision system that tracks the position and orientation of each robot at 60Hz. A novel aspect of this system is that computer vision is used not only for sensing the environment, but also for sensing the state of the robot itself. The former information is used for planning and the latter for control. The testbed has been used successfully for experiments with robot soccer games to compare the effectiveness of alternate control architectures. Testbed use is facilitated through a graphical simulator that allows control programs to be tested off-line at a workstation. The testbed has been developed to the point where we now have a portable version that has been demonstrated at a number of conferences.

1. INTRODUCTION

The Dynamite testbed provides a practical platform for testing theories in the soccer domain using multiple mobile robots. The testbed consists of a fleet of radio-controlled vehicles with perceptual, planning and control systems. The robots perceive their world through a shared off-board perceptual system. In an integrated off-board environment with dataflow and MIMD computers, vision programs monitor the position and orientation of each robot and the ball while planning and control programs generate and transmit motor commands. This approach allows umbilical-free behaviour for lightweight fully autonomous robots. As far as we know,

it is a unique and successful approach to the trade-offs involved in mobile robot design although a related scheme was independently proposed in [5].

2. ROBOTS

The mobile robot bases are commercially available radio-controlled vehicles. We have six 1/24 scale racing-cars, each 22cm long, 8cm wide, and 4cm high excluding the antenna. Although we plan to use all the cars, the most that have been used, so far, at any time is two. The testbed (244cm by 122cm in size) with two cars and a ball is shown in Figure 1. The cars have each been fitted with two circular colour markers allowing the vision system to identify their position and orientation. The ball is the small object between the cars. The robots are neither as flexible nor as competent as human soccer players, so we have modified the environment in two ways. First, there is a wall around the soccer field which prevents the ball, and the players, from going out of bounds. Second, there are barriers to prevent the ball from getting trapped in the corners. Since these are Canadian robots, it is not unreasonable for the soccer field to be shaped like an ice hockey rink.




Figure 1. Robot Players on the Soccer Field




Figure 2. Hardware Configuration

3. SYSTEM ARCHITECTURE

The original architecture of the testbed [1] is shown in Figure 2. There is a single colour camera mounted in a fixed position above the soccer field. The output of the vision system is common to all controllers. The different visual tracking hardware systems and algorithms are described in the next section. The reasoning and control components of a vehicle can be implemented on any number of transputers out of the available pool. Currently, each vehicle is controlled by a two-process user program running on a single transputer node. An arbitrary number of nodes, labeled 1 to n in Figure 2, can be used in parallel to control independent vehicles. The movement of all vehicles is controlled through radio transmitters attached to a single shared transputer node. Commands are transmitted to the vehicles at a rate of 60Hz.

A feature of the Dynamite testbed is that it is based on the "remote brain" approach to robotics [6]. The testbed avoids the technical complexity of configuring and updating on-board hardware and makes fundamental problems in robotics and artificial intelligence more accessible.

4. VISION

The basis of the vision system is the tracking of coloured discs that are placed on the robots. Even for this simple visual task, specialized hardware is required. Three vision systems have been constructed that use the same algorithms but run on different hardware.

The original system uses DataCube video processing hardware. The DataCube is a dataflow processor which has been programmed to classify image pixels into different colour classes at 60Hz (interlace video rate). This information is transmitted to a network of transputers which form a

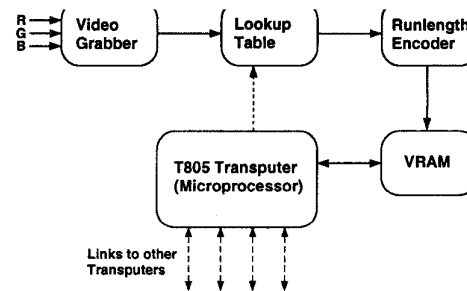


Figure 3. Architecture of the Colour Tracking System

MIMD computer. Additional vision processing is performed on the transputers to find the position, in screen coordinates, of the centroid of each coloured blob and to transform these positions from screen to world coordinates. The vision subsystem is called the Vision Engine [7]. The Vision Engine produces the absolute position of all the objects on the soccer field; the orientation of each car is also reported. This is done at 60Hz with an accuracy in position of a few millimeters.

The second vision system accomplishes the colour classification through the use of an analog circuit and performs the remaining computation on transputers as before. The analog circuit thresholds the video signal to provide a coarse division of colour space so that only red and green targets can be tracked [10]. This system is portable since all of the required computer hardware fits in a chassis that is the size of desktop computer.

The third system uses a digital colour video processing module, called the Colour Tracking System (CTS), which was developed to perform the colour classification. The main functional groups of the CTS hardware are shown in Figure 3.

The CTS is able to grab RGB video input in a variety of modes. Each pixel is classified, using the lookup table, into one of a number of classes based on its colour attributes. The runlength encoder compresses the description of the classified pixels so that they can be processed by the microprocessor in real-time. The hardware that comprises the runlength encoder is reprogrammable and allows the functionality of the CTS to be altered without hardware modifications. The INMOS T805 transputer is a general purpose microprocessor that segments the compressed classified pixels into connected regions and communicates this information to other




Figure 4. The Dynamite Simulator with Two Cars and a Ball

transputers in the system. The CTS conforms to the Transputer Module (TRAM) daughter card standard which allows it to be used in a variety of computing environments.

The speed of the system depends on the number and size of coloured objects present in the camera's field of view since the amount of computation performed by the transputer increases with the number of classified pixels to be segmented. For the Dynamite Testbed, the CTS is easily capable of 60Hz operation with a latency of several milliseconds following the reception of the video field.

5. SIMULATION

A physics-based graphics simulator for the Dynamite world is available for testing and developing reasoning and control programs. It has proven to be very useful in the development of control programs. The simulator supports an arbitrary number of cars and balls, as well as a simple contact model for collisions. A visual display of the simulator is shown in Figure 4.

6. PLAYING SOCCER

Soccer has been proposed as a task for the development and unification of divergent theories in Artificial Intelligence [9]. Soccer captures many important aspects of the real world such as competition and cooperation with other agents, a dynamic environment and real-time interaction.

7. REACTIVE DELIBERATION

A control architecture called *Reactive Deliberation* has been proposed for the task of soccer playing [8]. Under reactive deliberation, the robot controller is

partitioned into a deliberator and an executor; the distinction is primarily based on the different time scales of interaction. Informally, the deliberator decides what to do and how to do it, while the executor interacts with the environment in real-time. These components run asynchronously to allow the executor to interact continuously with the world and the deliberator to perform time consuming computations. This partition is inspired by recent architectures that attempt to integrate planners with more reactive components [3; 4] as well as architectures such as subsumption [2].

Several controllers based on reactive deliberation have been implemented to allow robots to compete in complete one-on-one games of soccer. A series of experiments called the LCI (Laboratory for Computational Intelligence) Cup has been conducted to compare alternate controller designs. The soccer tournament provided evidence that the features proposed for our robot architecture *Reactive Deliberation* are useful. There are three main results. First, intelligent action selection cannot proceed independent of planning in non-trivial domains. Second, goal-oriented behaviours are useful abstractions that allow sharing of scarce computational resources and effective goal-arbitration through inter-behaviour bidding. Third, experimental evidence has been provided that the goals and actions of a robot need to be evaluated at a rate commensurate with changes in the environment.

8. CONTROL

The deliberator is responsible for selecting an appropriate action or goal in the current situation. This includes the problem of non-holonomic path planning since the most appropriate action may depend on how long it takes to move from the current position to another location. The executor for the soccer-playing robot consists of five action low-level schemas: follow path, servo at ball, defend, stop, and idle. Only one action schema is enabled at a time and the enabled schema sets the robot's control outputs (throttle and steering angle) and sends messages to the active behaviour in the deliberator when it is having problems.

The current functionality of the controller includes various simple offensive and defensive strategies, motion planning, ball shooting and play-

ing goal. The robots can drive under accurate control at speeds up to 1m/s, while simultaneously considering alternate actions. We have produced a 10 minute video that documents these features.

9. CONCLUSIONS

The features and capabilities of the Dynamite testbed have been described in detail. The testbed has been used extensively for experiments in the soccer domain. The experiments allowed the comparison of alternative controller designs that has resulted in the Reactive Deliberation architecture. A more practical result is that we have demonstrated the viability of the remote brain approach for this class of experiments; the robots we have developed can drive under control at speeds of 1 m/s. Off-board computer vision, monitoring both the robot plant and its environment, can be used for real-time robot control.

9. ACKNOWLEDGEMENTS

We are grateful to Dinesh Pai, Jim Little, Heath Wilkinson, Steve Monai, Ying Zhang and David Weih for help with this. This work is supported, in part, by the Canadian Institute for Advanced Research, the Natural Sciences and Engineering Research Council of Canada and the Institute for Robotics and Intelligent Systems Network of Centres of Excellence.

10. REFERENCES

[1] Barman, R., Kingdon, S., Little, J., Mackworth, A. K., Pai, D., Sahota, M., Wilkinson, H., and Zhang, Y. Dynamo: real-time experiments with

multiple mobile robots. In *Proceedings of Intelligent Vehicles Symposium* (1993), pp. 261–266.

[2] Brooks, R. A. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation RA-2* (1986), 14–23.

[3] Firby, R. J. Building symbolic primitives with continuous control routines. In *First International Conference on Artificial Intelligence Planning Systems* (1992), pp. 62–69.

[4] Gat, E. Integrating planning and reacting in a heterogeneous asynchronous architecture for controlling real-world mobile robots. In *AAAI-92* (1992), pp. 809–815.

[5] Hallam, J. Playing with toy cars: An experiment in real-time control. Tech. Rep. DAI Research Paper No. 527, Edinburgh University, 1991.

[6] Inaba, M., Kamada, T., and Inoue, H. Rope handling by mobile hand-eye robots. In *Proc. ICAR-93* (1993).

[7] Little, J., Barman, R., Kingdon, S., and Lu, J. Computational architectures for responsive vision: the vision engine. In *Proceedings of Computer Architectures for Machine Perception* (1991), pp. 233–240. Paris.

[8] Sahota, M. K. Reactive deliberation: An architecture for real-time intelligent control in dynamic environments. In *Proceedings of AAAI-94* (1994), pp. 1303–1308.

[9] Sahota, M. K., and Mackworth, A. K. Can situated robots play soccer? In *Proceedings of Canadian AI-94* (1994), pp. 249–254.

[10] Wright, A. A high-speed low-latency portable visual sensing system. MIT course project, 1994.