MADbot: A *M*otivated *A*nd Goal *D*irected Ro*bot*

A. Coddington and **M. Fox** and **J. Gough** and **D. Long** and **I. Serina**
Computer and Information Sciences, University of Strathclyde, Glasgow, UK
Email: {intial.secondname}@cis.strath.ac.uk

In most work in plan generation and execution the assumption has been made that the goals being addressed by the planning system (and executive) are imposed externally and that once a plan has been constructed to achieve these goals the activity of the planner can cease. Similarly, once the plan has been successfully executed and a state satisfying the externally imposed goals has been reached, it has been assumed that the planning and execution behaviours will suspend until a new goal set and consequent plan has been imposed. These assumptions do not hold for fully autonomous systems, which are capable of directing their own behaviour and prioritising their own goals. The problem we are most concerned with is determining how goals arise during the autonomous behaviour of a system.

Recent work by Knight *et al* (Knight *et al.* 2001) and Chien *et al* (Chien *et al.* 2001) relaxed the above assumptions by introducing the idea of continuous planning in the Casper system. Casper is an architecture for an autonomous planning and control system intended for application in space missions involving onboard autonomy. In (Chien *et al.* 2001), the authors indicate that the system can respond to failures and opportunities during execution of a plan within a given horizon and, in principle, update its goal set in response to environmental factors. The RAX system (Jonsson *et al.* 2000) demonstrated the integration of planning, execution, failure diagnosis and subsequent plan repair. Our approach to autonomous planning and execution is to address the issues of goal generation and management according to the changing *motivations* of an autonomous system.

Although in general an autonomous system will be created to carry out externally-imposed tasks, any extended period of autonomous behaviour will require a system to react to its environment and to be able to create its own goals. The MADbot project concerns the development of a *motivated* autonomous system, capable of generating its own goals in accordance with a system of drives and impulses (Coddington & Luck 2004) and monitoring its own behaviour in the execution of plans. Execution monitoring is done with respect to stochastic models of the actions that the robot can execute. At any point during the execution of its plan the robot can estimate its most likely state based on the appropriate model and the current observations of the system. If its most likely state is a failure state the system can replan to achieve its outstanding goals.

The execution behaviour of the system differs from that exhibited by most plan executives because it might replan behaviours during the execution trajectory in order to manage goals generated in response to newly observed situations. These occur when:

1. The plan being executed leaves resources unused, making them available for previously unplanned activities.
2. Changes in the execution environment cause the motivations of the system to change in value, so that some exceed their "comfort" thresholds.
3. The plan experiences a failure, requiring the system to reconsider the priorities of its outstanding goals.

In these situations the system will incorporate additional goal-achieving activities into its behaviour. A plan adaptation system (Gerevini & Serina 2000), is used to plan how to incorporate the new goals into the remainder of the plan. In deciding which goals to address, and how to incorporate these into the plan, the top level goals of the original plan under execution remain paramount. However, it might be that the pressing need to respond to a motivated goal temporarily suppresses the plan and allows an originally unplanned-for task to take priority. This would only happen if failure to prioritise this task could lead to the failure of the overall plan.

To create a concrete context in which to explore this system behaviour we have constructed a simple problem domain in which a roving robot is gathering scientific data at various sites. This scenario is motivated by the current activities of the NASA MER rovers and the planned ESA ExoMars mission. The system we have developed is currently equipped with five motivations: the motivation to avoid long periods without gathering science-data (called *opportunism*), the motivation to free memory by communicating stored data (called *data communication*), the motivation to self-localise (called *relocalisation*), the motivation to identify interesting sources of data (*science target identification*) and the motivation to remain charged (we term this *replenish-energy*). The replenish-energy motivation is modelled as a function of time and the difference between the current and expected power levels of the system. Relocalisation is modelled as a simple function of time spent navigating between locations. The longer a navigation task takes

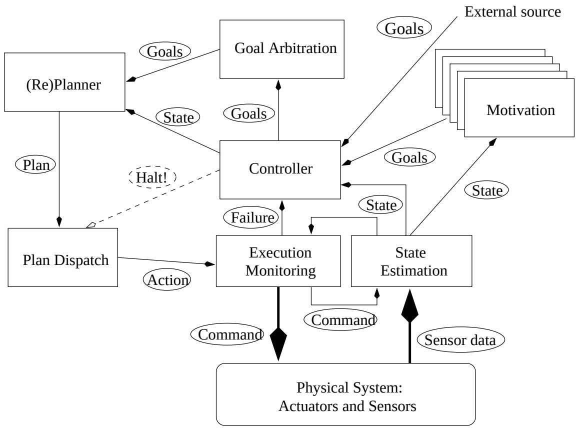

Figure 1: The architecture of MADbot. Rectangles represent components of the architecture. Ovals represent the data types communicated between the components in the directions indicated by the arrows. The dotted line and oval indicate a Halt instruction being communicated from the controller to the dispatcher.

the more pressing the need to relocalise, until eventually the motivation generates a goal to take a panoramic photograph, which is intended to help to relocalise the robot. The level of the opportunism motivation is affected by the acquisition and communication of scientific data and the time taken over these tasks. When the robot begins executing its plan these functions have low values, but as time passes they increase according to simple polynomial functions. The science target identification motivation causes the system to look for features in its wider environment that indicate the potential to gather interesting science. The opportunism motivation might subsequently prompt the system to explore such features. The data communication motivation ensures the system downlinks any stored data whenever opportunities to do so arise between planned actions in the appropriate temporal windows.

The architecture of our system is illustrated in Figure 1. Goals can be injected into the system from an external source whilst others are created internally according to the motivations. Motivations can react to the passage of time, monitoring the perceived state over time and generating goals if the system does not make appropriate observations in motivation-specific intervals. Motivations can also react to state changes when they occur. The goal arbitration sub-system decides which subset of a collection of goals should be given priority. This decision has been explored by other authors as the over-subscription problem (Smith 2004; Sanchez & Kambhampati 2005), although here it is also required to account for the distinction between top level goals (from the external source) and internal goals triggered by motivations. The planner sub-system must face the possibility of goals changing before a plan is completely executed. It is therefore important that it be capable of replanning to meet a combination of new and old goals. We have adopted a strategy based on planning by plan adaptation (Gerevini & Serina 2000) which exploits the existing (unexecuted) plan fragment as a starting point to seed a search for a new plan that achieves the newly formed goal set each time the planner is invoked.

The granularity of operations is an important parameter in the design of the behaviour of each element of the system. It is essential that the system should not thrash, continually replanning in response to the smallest changes in its environment but, equally, it is necessary to ensure that motivations are capable of reacting to changes in the (perceived) state in order to generate goals in a timely and appropriate way. Chien *et al* (Chien *et al.* 2001) also emphasise this need for responsiveness of the system at different levels of granularity. This partly motivates their introduction of a continuous planning system with a shifting plan horizon.

It is important that the motivations of the system do not interfere with its planned activities whilst these are progressing well. Therefore, the system maintains a log of its progress, keeping track of the time and resource consumption of its actions and how these relate to the expectations imposed by the plan. If resources become available, because the plan was too conservative in its estimates of resource usage, the system can use the accumulated resource to satisfy its opportunism motivation. Additional activities are only incorporated into the plan if the system can be adequately confident (to some probabilistic threshold determined a priori) that these additional activities will not prevent the rest of the plan from being executable. If resources become scarce the replenish-energy motivation of the system will increase and the system will eventually insert a refuelling sub-plan into the plan. The objective is to ensure that at least the expected science gains of the plan are achieved, and that they are often exceeded because of the goals generated by the motivation system.

References

Chien, S.; Engelhardt, B.; Knight, R.; Rabideau, G.; Sherwood, R.; Hansen, E.; Ortiviz, A.; Wilklow, C.; and Wichman, S. 2001. Onboard autonomy on the three corner sat mission. In *International Symposium on AI, Robotics and Automation in Space*.

Coddington, A., and Luck, M. 2004. A motivation-based planning and execution framework. *International Journal on Artificial Intelligence Tools* 13(1).

Gerevini, A., and Serina, I. 2000. Fast plan adaptation through planning graphs: local and systematic search techniques. In *Proceedings of AIPS 2000*.

Jonsson, A. K.; Morris, P. H.; Muscettola, N.; Rajan, K.; and Smith, B. D. 2000. Planning in interplanetary space: Theory and practice. In *Proceedings of AIPS 2000*.

Knight, R.; Rabideau, G.; Chien, S.; Englehardt, B.; and Sherwood, R. 2001. Casper: Space exploration through continuous planning. *IEEE Intelligent Systems* 16(5).

Sanchez, R., and Kambhampati, S. 2005. Planning graph heuristics for selecting objectives in over-subscription planning problems. In *Proceedings of International Conference on AI Planning and Scheduling*.

Smith, D. 2004. Choosing objectives in over-subscription planning. In *Proceedings of International Conference on AI Planning and Scheduling*.