Reactive and Motivational Agents: Towards a Collective Minder

Darryl N. Davis⋆

School of Computing, Staffordshire University,
Beaconside, Stafford, ST18 0DG, UK
D.N.Davis@soc.staffs.ac.uk,

Abstract. This paper explores the design and implementation of a societal arrangement of reflexive and motivational agents which will act as the building blocks for a more abstract agent within which the current agents act as distributed dynamic processing nodes. We contest that reactive, deliberative and other behaviours are required in complete (intelligent) agents. We provide some architectural considerations on how these differing forms of behaviours can be cleanly integrated and relate that to a discussion on the nature of motivational states and the mechanisms used for making decisions.

1 Introduction

This paper reports on the *Architectures for Intelligent Agents* project within which computational complete agent architectures are being investigated (using a two dimensional simulated world). The work associated with this project is open-ended but primarily relates to the investigation and exploration of the possibilities associated offered by different agent architectures, for the modelling of motivational and other control states. The motivation for this work include:

- By producing plausible computational models of simulated agents we may further our understanding of biological, psychological and social agents.
- By designing and implementing agent architectures based on different theories of the mind, we may better understand the strengths and inadequacies of these theories.
- By developing working agent architectures in a dynamic and potentially hazardous (simulated) environment, we can further our theories and models of control mechanisms for use in real environments.

These are very long-term motivations and we can expect to make slow progress in these directions. The following (non-exhaustive) list of qualities and types of processing that the (human) mind exhibits shows how far short we are from developing truly "intelligent" computational agents: to perceive the world and learn; to remember and control actions; to cogitate and learn new ideas; to control communication with others; to create the experience of feelings, intentions and self-awareness [11].

⋆ Work carried while author a member of the Cognition and Affect Group, School Of Computer Science, University of Birmingham, Birmingham, B15 2TT, UK

We would like to address a wide range of issues, so many issues are addressed at a relatively coarse grain. Like Bates and colleagues [1], our high level aims include understanding natural and artificial complete agent architectures. This broad approach necessarily requires an initially shallow approach in designing computational agencies. The exploration of deeper, more complete, implementations will follow. In the short term, we hope that our investigations will shed light upon a number of interesting research questions such as: what makes a complete agent; what kinds of motivation and related processing does a complete agent have; what are appropriate architectures and the control issues for such an agent; do complete agents have a number of levels of processing (perhaps of differing modalities of operation) and if so, how can they integrated; does sensory data need be structured for different levels of internal processing (behaviours); and what are the appropriate decision mechanisms to use in determining which among (possibly conflicting) behaviours are to be preferred in certain circumstances. Many of these questions impinge on that addressed by Franklin and Graesser [8]: i.e. what sort of thing is an agent?

Earlier work in the Cognition and Affect group has focused on differing agent architectures (based on different simplistic models of biologically and psychologically plausible mechanisms) in a number of simulated environments. Parallel to this has been the development of an information processing architecture [19, 6], that allows many different coexisting components with complex interactions. Some processes are automatic (pre-attentive) in the sense that all changes are triggered directly; for example reflexes (whether learnt or innate) that bypass 'normal' processing. Other processes at a similar level are responsible for the activation of motivational states but pass control to further levels, where 'reflective' or 'attentive' knowledge-based management processes explicitly consider and evaluate options before selection. The management processes are resource-limited and may require attention filters to protect them from disturbance; and meta-management processes, involving some sort of self-monitoring, to regulate and direct them. The internal management (and meta-management) processes involve a certain amount of (asynchronous) parallelism. Resource limits restricting parallelism in high level processes may lead to emotional and other characteristically human states involving partial loss of control of attention and thought processes [16, 17]. This requires an architecture combining a wide variety of types of mechanisms.

We see our work as analogous to that being pursued by a number of other research groups on complete agents; for example the SOAR community [13], the behaviour-based subsumption work of Brooks [4, 5], and others such as Georgeoff [9]. However none of these researchers really try to integrate reactive and deliberative processing levels; indeed Brooks has taken a position denying the need for explicit representation. Hayes-Roth [10] suggests that peripheral reflexes can be incorporated into the Guardian architecture, along with reactive and deliberative processes; however it remains unclear how conflicts between high and low level behaviours are resolved. Other researchers take a variety of stances. Perhaps the reason for this huge differentiation in approaches to agents is due to the various researchers addressing just part of a larger research area.

2 Overview of the Agent Design

Here we consider a development of an ongoing scenario; that of an artificial crèche with a minder looking after a number of babies (or charges) in a dynamic, and possibly hazardous, environment. The environment can be seen as representing a type of factory floor, with agents (of differing sophistication) performing different kinds of tasks; ranging from simple conveyance to collaborative maintenance. The base level agency (a kind of mobile plinth) is allowed some degree of autonomy and must navigate the environment without resource to higher level deliberative processes; thus freeing them for other tasks. The minder is an abstract entity (with its own representations of the world and its tasks within that world) which relies upon perceptual information to update its model of the environment. It makes use of higher level deliberative agents to perform reasoning tasks over its model of the world and other representations related to agent goals and possible actions.

One thread that we are currently pursuing draws on the ideas of Minsky (expounded in [12]); i.e. the minder does not exist as a completely independent entity but is a collection of the reactive and deliberative agencies. Each agent can be given slightly different capabilities, and monitored by the minder. To simplify initial development work, and ensure that the different layers to our architecture function in their own right, we allow each agent to use a separate "computational" body. Conflicts between the actions and goals of the (situated) agents will provide a framework for experimenting with resource bounds on the management-level processing. Here we will also try to deepen and make more generic, the perceptual, reactive and goal handling processes within our architectures. Initially this will be at the expense of ignoring some of the other higher level attentive and resource management processes.

The current experimental scenario can make use of up to four agents:

reflexives: simple instinctual and reflexive agents with no explicit motivational states but the need to avoid colliding into agents and other objects within the environment.

reactives: these extend the mechanisms used in the reflexive agents with extra capabilities and behaviours including explicit goal-oriented motivational states such as *hunger* (or recharge) and *avoid danger*. This type of behaviour is similar to those expounded by Brooks in his description of behaviour-based agents [4, 5] and can be modelled using a subsumption-type architecture.

reflectives: Here the mechanisms of the reactives are further extended with abstract and deliberative (management) processing capabilities. This will include cognitive behaviours such as explicit planning, the consideration of multiple *surfaced* goals and the resolution of conflict between proposed actions arising from the different processing layers of the agents.

minder the (multi-agent) agency that is (collectively) responsible for the monitoring of the environment and the initiation of actions upon it to care for the agents that it contains. While it may contain its own private representational schemes (e.g. an explicit model of the environment or descriptions of tasks), it will make use of reflective agents: to provide perceptual information about the environment; to reason about incoming and persistent information; to develop and maintain motivational states; and to perform actions upon the world, for example, to collect an effector and retrieve some object or agent from a possibly hazardous situation. In a full implementation of the architecture discussed in the introduction, it could monitor (and perhaps control) the types of deliberative processes of the reflective agents, this is what our earlier work [6] has termed *inner perception* and *inner action*. This meta-deliberation arises from an interaction of the behaviour of the overall agent, and its designated (or acquired) niche role. We consider this meta-management level processing of the reflectives to be the most abstract level of agent processing, and so avoid the recursive abstraction problem of meta-meta-deliberation to monitor the meta-deliberative processes, and yet further levels monitoring them.

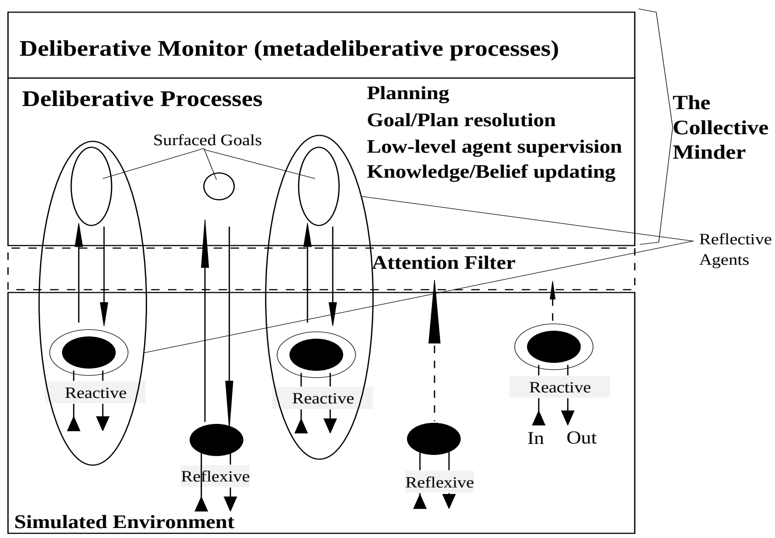

Fig. 1. The different levels of agency in the collective minder scenario.

Figure 1 provides a sketch of the relation between the different levels, with two reflexive, one reactive and two reflective agents. Here, the minder exists only as the collective processing of the reflective agents and their actions upon the environment. The motivating goals of one reflexive and one reactive are currently being ignored by the deliberative processes. The needs of one reflexive agent and the reactive goals of two reflective agents are being attended to as surfaced goals. The exact relationship and means of communication between higher levels of agency and the reflexives (and reactives) is an open issue; currently the reflective agents make inferences related to the needs of reflexive (and other) agents as part of the instantiation of caring goals.

3 A Design for a Crèche Simulation

The crèche consists of one large room of four walls: the north wall contains an entrance (or intake door); the west wall contains a recycle door through which 'broken' or spent agents are ejected; while the east wall contains an exit (or discharge door), and a hazard, in this case a ditch into which agents can fall. Other, multiple room, environments have been used. Ditches can be seen by all agents; however this does not mean that an agent can avoid falling into one. The base level agents can be thought of as frictionless platforms moving around the environment, driven by a single directed lateral (impulse); the other agents are complexes built upon this platform. An energy source is present. Normally it is static, but when it needs recharging it moves (towards the recycle door) and acts as a memory hazard. All agents can visually sense the moving energy source, and so attempt to avoid collisions. However, only one subtype of the reactive (but all the reflective) agents can sense the memory danger and so give it a wide berth; this is a benefit that other reactive agents accrue through teaming up with instances of this agent.

3.1 General description of the agents

All agents enter the environment through the intake door or are already present. They leave through the discharge door, or if ‘terminally’ damaged through the recycle door. Three classes of agents can be present in the crèche: the semi-autonomous (implicitly motivated) agents (reflexives); explicitly self-motivated agents (reactives); and explicitly motivated caring agents (reflectives). There is also the potential for the agents to team up and form attachments. The agents can be given up to three senses (visual. auditory and memory danger detectors) so that they can negotiate their way around a hazardous environment. The auditory and memory danger senses have a 360 degree field; while vision is restricted to 180 degrees, centred on the current direction of the agent. When they are deemed to have achieved some high level goal, agents can be discharged from the environment.

Reflexive agents The base level (reflexive) agents combine internal processing with perceptual and action processes. These agents can move in one of four directions (north, south, east and west) and are given an initial energy level, velocity and direction. Any change in velocity or direction causes an energy unit to be consumed. When the energy level is reduced below a certain level, they must be recharged. When their energy reaches zero, they become static until ‘rescued’. Reflexive agents can be given implicit low-level goals (such as move to location(X,Y)), by the minder or reflective agents. A limited set of behaviours define how the reflexive agents move around the environment. The *stop* behaviour brings the agent to an immediate halt; the *start* behaviour gives the agent a velocity in its current direction; and the *turn* behaviours actuate a change in direction either to the left or right. Further (hierarchically organised) behaviours such as accelerate, reverse and wander have been experimented with. The impulse to perform a specific behaviour is activated through the use of perceptual information. For example, if an agent senses objects in front, behind and to the right, turning left will be the most appropriate behaviour. If the agent senses objects in front and to the left and right, then stopping will be an appropriate behaviour. A default behaviour, relying upon no perceptual information, is sanctioned if no other action is initiated. This causes the agent to continue moving in the current direction with the current velocity; if stationary, the agent remains stationary. Initially these behaviours will be implemented in a shallow (instantaneous) manner; further deepening of the architecture, the environment and the behaviours will require a more natural modelling of these behaviour forms (i.e. starting and stopping can become more gradual temporal processes through acceleration and deceleration).

Reactive agents The second level reactives are reflexive agents that make use of extra (perceptual, internal and action-related) processing level related to explicitly modelled motivational (or goal-oriented) states. The goals of these agents are self-centred reactions to internal (hunger) and external (danger) states. The behaviours associated with these motivational states can be built up from the more primitive (reflexive) actions already described; for example, to feed the agent needs to move to a specific location (the energy source) while avoiding collisions using the stop, start, turn and default behaviours. Two subtypes (placid and excitable agents) demonstrate extra motivational behaviours, with differing preferences in the population density of their immediate environment; for example, one class of agent (placid) prefers less populated areas. Like the reflexive agents, there is no persistent representation of the external environment. However, goals, their status and their processing structures, do persist over time. The processing and structure of goals in discussed further later in this paper.

Reflective agents The third level reflective agents combine the mechanisms used in the reflexive and reactive agents with deliberative processes. They extend the motivational states associated with the reactive agents to include more selfless goals such as feed reflexive agents, rescue any trapped agent or investigate non-moving agents. The cognitive processes are also further extended to allow a persistent model of the world (incorporating a simple memory model), an attention filter and a number of deliberative processes including the explicit consideration of goals. The agent must reason over its representation of the environment to generate these potential selfless goals. The simple memory model allows conflicts between beliefs and makes use of fuzzy valued memory strength parameters. In the current work the meta-deliberative processes are very (very) shallow. Reflective agents cannot (presently) communicate with each other but attempt to co-operate by occupying different areas of the environment.

3.2 The information processing ontology

We use a precise specification of the information processing ontology when designing our agents. Though the formal notation (a propositional calculus) does not currently map onto our agent toolkit [20], it does ensure that a consistent view of epistemic events is maintained, at least at the design level. Examples of propositional statements related to perception are given in figure 2.

We make use of symbolic processing rules and rule-sets to simulate the various internal and external behaviours in all the agents. There are rule-sets for: the management of incoming information; the various behaviour activation modules; the generation, instantiation and management of goals; and for making decisions at various levels in the more complex agents. For example rule *l0_start_rule* in figure 3 states that if the agent

new_sense_data is used by the sensing methods in our toolkit. It is used to derive all information about objects in an agent's environment: e.g. [new_sense_data (object:wall4) 100]
percept0 propositions relate not only to the perceptual modality but also give a sensed object's relative position: $percept0(mode, object\_identifier, position, range, x, y)$
where $position \in [Front, Back, Left, Right]$ and $mode \in [vision, magnetic]$
e.g. [percept0 vision (object:wall4) Left 100 -210 -150]
belief used to reference beliefs (at the deliberative level) about the perceived world:
$belief(timestamp, object\_identifier, x, y)$

Fig. 2. Examples of propositional statements used in the agents

is active, has zero velocity and cannot sense anything in front of it, generate a level zero (behaviour) goal to start moving. The second rule generates a high level (carer) goal when the agent can infer, from information pertaining to current and past environmental states, that another agent has not moved over the last two cycles.

Low level (0) behaviour : RULE l0_start_rule1
[Status Active]
[Velocity 0]
[NOT percept0 ?sense ?object Front ?distance ==]
==>
[gen_goal 0 start []]
High level (2) goal generation : RULE care_goal2
[belief ?cycle1 ?object:isagent ?posx ?posy]
[memory ?cycle2 ?object ?posx ?posy] [WHERE cycle1 > (cycle2+2)]
==>
[gen_goal 2 investigate [[Object ?object ?posx ?posy] [Detect ischarge] [Importance low]]]

Fig. 3. Examples of goal generating rules used in the agents

4 Architectural Considerations

Here we will consider how to design a computational architecture for the simple reflexive agents that can be extended for the processes needed for the reactive and reflective agents (and ultimately the entire minder). There are a number of architectures that we might consider. In investigating what are appropriate agent architectures for the scenarios we are interested in, we should aim to address such questions as:

How can different kinds of learning be integrated in these architectures.
Are there different types of actions, responses and situations at different levels in an information processing (agent) architecture and if so in which circumstances is one (or more) behaviours more appropriate? Also how does an agent choose between behaviours, given that some may be incompatible?

What types of control systems (e.g. feedback) do we require to model the required agent behaviours?

Related to the control issue are questions such as how do certain behaviour sets become over-ruled or interact. For example, in response to certain perceptual stimuli, the internal processing nodes may specify that accelerate and turn behaviours are appropriate. Do we allow the agent to choose just one of these (and if so how?) or an interaction resulting in more complex behaviours?

One possibility is an integration of the behaviour-based approach with a more classical AI blackboard approach (in a manner not dissimilar to that often proposed by Hayes-Roth [10]). For example, a number of (simulated) concurrent behaviours can be given access to sensory data, each (possibly) producing its own action potential (see 4). This differs from orthodox subsumptive architectures (where there are direct links between different behaviours) in that an explicit (symbol based) potential action blackboard is used. The agent decides on the most appropriate (set of) posted behaviours on the basis of some decision process (for example which behaviour subsumes most action potentials). An alternative approach is to use a Touring Machine type architecture [7]. Here a decision (processing) node, for example context-activated control rules, is used to override certain behaviours. This high level module (or agency) is responsible for switching behaviour nodes on and off. One criticism of this type of control model is it tends to preclude low-level (instinctual/reflexive) behaviours that need to bypass higher level control mechanisms.

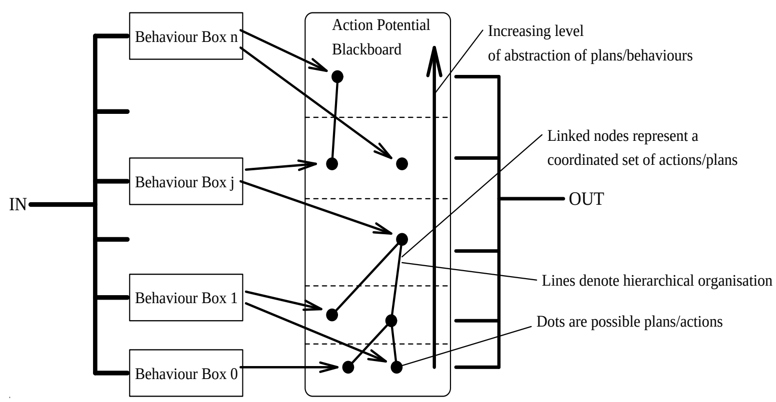

Fig. 4. Hybrid Behaviour-Based Architecture (with an action potential blackboard)

A criticism of both of these classes of architectures is their use of a flat perceptual system, with no discrimination between the types of perceptual information needed by the different behaviour nodes. Even simple biological mechanisms tune the perceptual information to the type of input appropriate to or expected by the processors responsible for different behaviours. For example, the sensory information passed to a frog's fly-catching behaviour will be insufficient (and different) to that of use to its predator avoidance behaviour (and vice versa). Similar analogies can be drawn with the sensory processing of higher level organisms. To simply stop and not collide with another object in the environment, it is only necessary to consider its relative location (and direction); to determine something more about that object (for example, its potential danger or state of distress) requires further processing. A close read of Brooks [4] shows that different sensor types are used at different levels within his subsumptive architectures; he is in fact simply bypassing structuring (and re-representation) of perceptual information by placing this structuring at the sensor level! A further, and more serious criticism, is that the discussed architectures do not seem to offer the mechanisms necessary for goal processing, or only allow one specific form of motivational processing. We may therefore need to combine both these approaches with something like that described in [9] and [2].

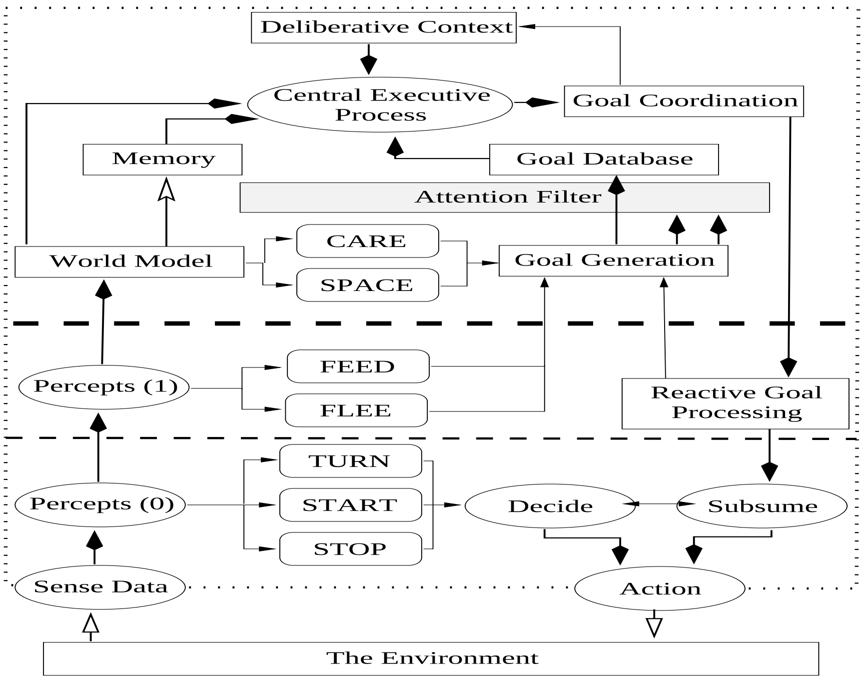

Fig. 5. A hybrid subsumption architecture for simple agents.

The architecture that we are proposing is a combination of these ideas and the behaviour-based approach, but addresses the above criticism of flat perceptual architectures. Figure 5 shows how a small set of different behaviour or information processing levels can be put together. We can model the different sets of behaviours (described above) as individual (and independent) processing nodes (implemented using symbolic rule-sets). Although we have placed similar levels of processing together in distinct layers, these are not connected. Their commonality is the use of a specific perceptual processing output. These processing nodes can be viewed as concurrent activities (although in our toolkit we can only simulate concurrency). For example, each of the behaviour nodes can be considered to be locally concurrent and each level of the architecture similarly so. The lowest level may provide more rapid reflexes as it is placed earlier in the perceptual processing chain and requires fewer control decisions to map perceptual information though to actions. As we ascend the architecture, the processing layers become slower. A full implementation of this design would require a toolkit capable of supporting asynchronous concurrency. The attention filter is not dissimilar to the (numerically) quantitative mechanism for use in selecting between goals in [15]; it's function, however, is not to activate goals but to protect the rational processes from being overloaded with (currently) unimportant potential goals.

5 Goals, Plans and Motivational States

What are goals? Dependent upon your perspective, this term can mean a number of things (see [2, 7, 9, 15, 16]). Here we consider how goals fit into the possible design for a motivated, autonomous agent. To do this we will have to consider not only goals but a number of other phenomena (e.g. motivation, attitudes and attention). This will lead into a consideration of what is required to process goals in an information processing architecture for a mind. A major problem in this area of research is that many different research groups use different terms to relate to similar phenomena and as perplexingly the same term to refer to different phenomena. The starting place for Sloman and Beaudoin's consideration of goals [2, 19] is the conjecture that the mind can be viewed as a control system, non-exclusively composed of belief and desire-like control states, and that goals are equivalent to some forms of control states.

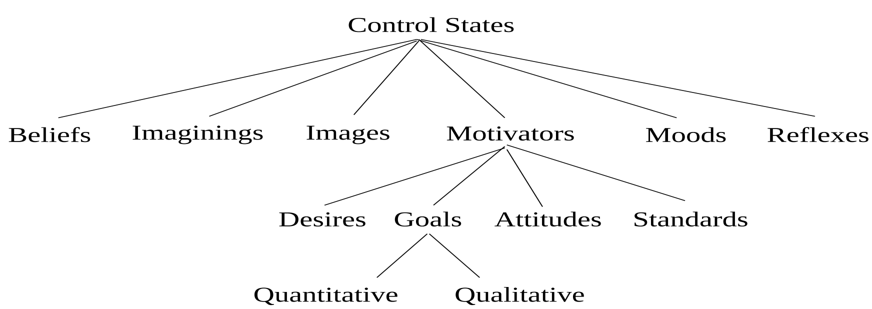

Fig. 6. A simple taxonomy for different motivational states, based on [2]

A taxonomy of six different forms of control states can be described (see figure 6); one of which (*Motivators*) consists of four further categories. Briefly these control states, which are categorisation of mental phenomena, can be described:

Beliefs are internal models of the world, possibly inferred from perceptual acts (e.g. dangerous object close and to the left); these need not have a rational basis;

Imaginings can be characterised as hypothetical "what-if" processes;

Images are control states using mental images (for example, used in spatio-temporal reasoning or thinking about the work of different visual artists);

Motivators are dispositions to assess situations in a certain way; i.e. a context for reasoning about epistemological events.

Moods are persistent states; they can be viewed as emergent states that pervade the entirety of cognitive processing or a side-effect of other control states. Certain moods favour certain motivators and inhibit others; i.e. they are closely related to predispositions and attitudes.

Reflexes are ballistic mappings from input (i.e. perception) to output (i.e. behavioural response) and can involve single actions (e.g. eye blink) or multiple actions (e.g. moving head and eyes in the direction of an unexpected sound).

Motivators can take several forms :

Desires are aims or goals that may not be realistic or achievable; however these can still influence the behaviour of an agent;

Quantitative goals are those type of goals talked about in control theory, and tend to involve negative feedback;

Qualitative goals are similar to most artificial intelligence goals (especially in the planning literature) and can involve relations, predicates, states, behaviours etc.;

Attitudes are predispositions to respond or act (either internally or externally) to specific (perceptual or cognitive) cues and can involved intricate collections of beliefs, motivators etc.;

Standards are prescriptive or relative states that embody ethical, social or personal rules.

This taxonomy, in practice, is very fuzzy and there is much overlap between these different categories. It is also incomplete; an obvious omission is how personality traits influence control states, and effectively act as higher order motivators. The current work addresses only a few of these control states.

5.1 Conceptual structure of goals and other control states

Control states can be implicit or distributed among co-existing processes and memory structures; for example, the behaviours (reflexes) associated with the lowest architectural level in figure 5 are implemented in such a manner. Goals, however, are explicitly represented. The core of a goal is some descriptor. This can be as simple as a single attitude (e.g. *make true* towards a single proposition (e.g. *move_to(location(X,Y))*). In many cases there is a need for multiple attitudes towards multiple propositions, and these attitudes can be varying in their nature. Goals may be nested; for example, the top-level goal associated with the motivational state of hunger is to feed from the energy source; this may require subgoals such as locate and move to the energy source.

Among the more important attributes of goals (see [2, 18]) are

A list of preconditions for the goal to be generated. Related to this a list of satisfied preconditions, for example ”reactive1 has low energy level” and a motivational attitude related to the propositions (e.g. *make false*).

A set of fuzzy values for goal importance (e.g. *high, medium, low*); goal urgency (e.g. *within 5 cycles*); and goal intensity (e.g. *high, medium, low*).

Sub-goal or plan factors such as a list of sub-goals, or plans, and the other agents involved. In the current implementation each class of goal has a set of pre-compiled (or reactive) plans; this is discussed in further detail below.

Status information such as commitment status, e.g. one of [*unknown, adopted, rejected, ignored*], and its dynamic state, e.g. one of [*passive, postponed, active, failed, successful*].

plan1 :
IF Energy(V1) *AND* EnergyThreshold(V2) *AND* Greater(V1,V2+VT)
THEN SATISFIED(Goal(Hunger))
plan2 :
IF Energy(V1) *AND* EnergyThreshold(V2) *AND* Lower(V1,V2+ VT)
AND Sense(EnergySource) *AND* Close(EnergySource)
THEN SubsumeBehaviours(Feed)
plan3 :
IF Energy(V1) *AND* EnergyThreshold(V2) *AND* Lower(V1,V2+ VT)
AND Sense(EnergySource) *AND NOT*(Close(EnergySource))
THEN GenerateGoal(Move, [EnergySource, Insistence])
plan4 :
IF Energy(V1) *AND* EnergyThreshold(V2) *AND* Lower(V1,V2+ VT)
AND NOT(Sense(EnergySource))
THEN GenerateGoal(Detect, [EnergySource, Insistence])

Fig. 7. A set of plans associated with the hunger goal

In formulating plans for achieving goals we map our formal specification onto a technique devised by Nilsson [14]. Figure 7 shows a set of shallow plans for the hunger behaviour; other goals have similarly defined plans. At their simplest, these plans provide an exclusive serial choice; if plan1 can be applied the (hunger) goal is satisfied; else if plan2 can be applied the agent will feed; else the agent generates (sub-)goals to either move towards the energy_source (plan3) or explore the environment and find the energy source (plan4). If none of these plans can be used, there is a generic mechanism (used for all goal plans) that activates a “goal-plans-failed” process; this may in turn activate a “abandon goal” process.

5.2 Making decisions about actions and goals

Our adopted architectural approach to modelling agents, with competing behaviour modules acting independently of each other, can cause problems in that we need to provide some means of deciding between conflicting potential actions. For example, a reflexive agent cannot turn left and turn right at the same time, but could turn left and stop, or start and turn right. A number of possible solutions exist including trainable decision nets and associating (static or dynamic) weights with behaviours. The decision and subsume boxes (in figure 5) contain rules which define compatible behaviours allied to a set of (currently static) weights. We use a 'winner-takes-all' strategy for conflicting actions, whereby the potential action(s) with the highest weight is preferred. Sometimes different actions are not in conflict and parallel or sequential combinations can be adopted. An alternative approach [20] is to use a simple additive rule to produce a resultant action.

For reactive and reflective agents with a number of levels of processing a more sophisticated approach is required. If no goals are generated, the agent acts as if it is a reflexive agent. For reactive agents, with one or more generated goals we use the same strategy to choose a goal (i.e. the goal with the highest insistence value wins). A goal insistence provides a quickly defined indicator of a goal's possible importance and urgency. Its calculation involves heuristics suggesting a synergy between a goal's relative abstract status (for example, a flee danger goal has a *high* status) and the current perceived state of the agent's environment (e.g. the source of danger is very close). An adopted goal is then executed at the reactive (and reflexive) levels. Depending upon the current internal state of the agent, and what it "knows", this may spawn further nested goals or map into actions to be performed on itself or other agents and objects in the environment. Where the mapping of plans onto low-level actions occurs, actions sanctioned by goal plans are given temporary weights reflecting the insistence value of the goal. It should be noted that an agent need not necessarily prefer the potential actions related to reactive (or other) goals over reflexive actions. The agent then decides from its set of (possibly conflicting) potential behaviours using the subsumption and decision processes at the reflexive level.

For reflective agents, if a goal is generated and its insistence value is higher than the threshold of the attention filter, then it can be placed on the goal database. If a goal is subsequently adopted, then the attention filter threshold can be set to its insistence value and the deliberative context will reflect the context of this goal. The adoption of a goal ought to include an evaluation of its current plans and whether it is currently achievable (i.e. a goal can be accepted but deferred). The plans associated with an adopted and active goal are (currently) executed at the reactive level as described above. If further goals are generated (in subsequent time intervals) and their insistence value is greater than the attention filter threshold, then they are added to the goal database; potential goals not meeting this criteria are simply deleted. The central executive process on noticing goals in the goal database calls a process to rank these goals in terms of their intensity (a combination of importance and urgency); the most intense goal can replace the current goal. The goal database is managed in other ways too; deadlines (detailed from within the representational structure for goals) can cause goals to be deleted from the database. Goals can also be stricken from the deliberative levels if they are abandoned during execution. We are looking at ways of combining and co-ordinating the execution of multiple goals, for example allowing an agent to feed itself and another agent at same time, without pre-compiled plans to do both. Currently it is possible for an agent to reduce it's energy level to zero while ensuring another agent is recharged!

6 Conclusion

Our developing theory does provide an architecture that allows a clean integration of reflexive, reactive and deliberative behaviours. However, providing an objective evaluation of a computational architecture, that is related to a slowly emerging theory of possible mechanisms for mind is not a simple task. One means of evaluating our work, is to compare the ratio of simple to carer agents for different minder architectures. For example, we have found that providing the minder with a central co-ordinating database increases this ratio over that necessary for a completely distributed model.

Future computational developments of the work presented here fall into two main areas: extending the complexity of the scenario by developing further the reflective agents and other processes associated with minder agents; and investigations into how learning can be incorporated into our current architecture. Further work on the "attention filter" ought to allow some contextual (or qualitative) filtering (similar to that used in [3]). For example, potential goals related to an agent's current deliberative context should need a lower filter threshold than unrelated goals. The attention filter could be used to filter perceptual information, as well as goals, with its threshold related to the context for the current deliberative processing. We also need to consider how the differing perceptual information (from the reflectives) can combined to give (the minder) a globally centred descriptive model of the environment rather than a number of agent centred descriptions.

We will need to address the nature of the higher (meta-management) reflective processes. Even though we have made only a small progress in our stated aims, our simple model implementation has shown emergent properties that are of concern; in particular, perturbant states resulting from competitive and conflicting motivations. Perhaps further work on the deliberative (management and meta-management) processes will highlight ways in which these problems can be resolved. This may include features such as evaluation of behaviour and long-term goals with regard to niche roles, what niche roles are and how they develop and influence cognitive behaviour.

7 Acknowledgements

This research was funded by a University Of Birmingham internal grant. Members of the Cognition and Affect group, in particular Aaron Sloman and Ian Wright, for discussions related to this work.

References

1. J. Bates, A. B. Loyall, and W. S. Reilly. Broad agents. *SIGART BULLETIN, 2(4)*, pages 38–40, August 1991.
2. L. P. Beaudoin. *Goal processing in autonomous agent*. PhD thesis, School of Computer Science, University Of Birmingham, 1994.
3. L. M. Botelho and H. Coelho. Emotion-based attention shift in autonomous agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages*

(ATAL-96), Lecture Notes in Artificial Intelligence. Springer-Verlag, Heidelberg, 1996. In this volume.
4. R. A. Brooks. How to build complete creatures rather than isolated cognitive simulators. In K. VanLehn, editor, *Architectures For Intelligence*, pages 225–239. LEA Pubs, Hove and London, 1991.
5. R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
6. D. N. Davis, A. Sloman, and R. Poli. Simulating agents and their environments. *AISB Quarterly*, October 1995.
7. I. A. Ferguson. Integrated control and coordinated behaviour: a case for agentmodels. In M. J. Wooldridge and N. R. Jennings, editors, *Intelligent Agents — Proceedings of the ECAI94 Workshop on Agent Theories, Architectures, and Languages (ATAL-94)*, Lecture Notes in Artificial Intelligence, pages 203–218. Springer-Verlag, Heidelberg, 1995.
8. S. Franklin and A. Graesser. Is it an agent, or just a program?: A taxonomy for autonomous agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Heidelberg, 1996. In this volume.
9. M. P. Georgeoff and A. L. Lansky. Reactive reasoning and planning. In *Proc the Sixth National Conference on Artificial Intelligence*, volume 2, pages 677–682, Seattle, WA, 1987. AAAI.
10. B. Hayes-Roth. Intelligent control. *Artificial Intelligence*, 59:213–220, 1993.
11. P. N. Johnson-Laird. *The Computer and the Mind*. Fontana Press, second edition, 1993.
12. M. L. Minsky. *The Society of Mind*. William Heinemann Ltd., London, 1987.
13. A. Newell. *Unified Theories of Cognition*. Harvard University Press, 1990.
14. N. J. Nilsson. Teleo-reactive programs for agent control. *Journal of Artificial Intelligence Research*, 1:139–158, 1994.
15. T. .J. Norman and D. Long. Alarms:an implementation of motivated agency. In M. J. Wooldridge, J. P. Müller, and M. Tambe, editors, *Intelligent Agents II — Proceedings of the Second International Workshop on Agent Theories, Architectures, and Languages (ATAL-95)*, volume 1037 of *Lecture Notes in Artificial Intelligence*, pages 219–234. Springer-Verlag, Heidelberg, 1996.
16. H. A. Simon. Motivational and emotional controls of cognition. In *Models of Thought*, pages 29–38. Yale University Press, 1979.
17. A. Sloman. The mind as a control system. In C. Hookway and D. Peterson, editors, *Philosophy and the Cognitive Sciences*, pages 69–110. Cambridge University Press, 1993.
18. A. Sloman. What sort of control system is able to have a personality? In R. Trappl, editor, *Proceedings of Workshop on Designing Personalities for Synthetic Actors*, Vienna, June 1995.
19. A. Sloman, L. Beaudoin, and I. Wright. Computational modeling of motive-management processes. In N.Frijda, editor, *Proceedings of the Conference of the International Society for Research in Emotions*. ISRE Publications, July 1994.
20. A. Sloman and R. Poli. Sim_agent: A toolkit for exploring agent designs. In M. J. Wooldridge, J. P. Müller, and M. Tambe, editors, *Intelligent Agents II — Proceedings of the Second International Workshop on Agent Theories, Architectures, and Languages (ATAL-95)*, Lecture Notes in Artificial Intelligence, pages 304–316. Springer-Verlag, Heidelberg, 1995.

This article was processed using the LaTeX macro package with LLNCS style