Modelling Shared Extended Mind and Collective Representational Content

Tibor Bosse[1], Catholijn M. Jonker[1], Martijn C. Schut[1], and Jan Treur[1,2]
[1]Vrije Universiteit Amsterdam, Department of Artificial Intelligence
{tbosse, jonker, schut, treur}@cs.vu.nl
http://www.cs.vu.nl/~{tbosse, jonker, schut, treur}

[2]Universiteit Utrecht, Department of Philosophy

Abstract

Some types of animals exploit the external environment to support their cognitive processes, in the sense of patterns created in the environment that function as external mental states and serve as an extension to their mind. In the case of social animals the creation and exploitation of such patterns can be shared, thus obtaining a form of shared mind or collective intelligence. This paper explores this shared extended mind principle for social animals in more detail. The focus is on the notion of representational content in such cases. Proposals are put forward and formalised to define collective representational content for such shared external mental states. A case study in social ant behaviour in which shared extended mind plays an important role is used as illustration. For this case simulations are described, representation relations are specified and are verified against the simulated traces.

1. Introduction

Behaviour is often not only supported by internal mental structures and cognitive processes, but also by processes based on patterns created in the external environment that serve as external mental structures; cf. [5, 6, 7 & 8]. Examples of this pattern of behaviour are the use of ‘to do lists’ and ‘lists of desiderata’. Having written these down externally (e.g., on paper, in your diary, in your organizer or computer) makes it unnecessary to have an internal memory about all the items. Thus internal mental processing can be kept less complex. Other examples of the use of extended mind are doing mathematics or arithmetic, where external (symbolic, graphical, material) representations are used; e.g., [4 & 12]. In [16] a collection of papers can be found based on presentations at the conference ‘The Extended Mind: The Very Idea’ that took place in 2001. Clark [6] points at the roles played by both internal and external representations in describing cognitive processes: ‘Internal representations will, almost certainly, feature in this story. But so will external representations, …’[6, p. 134]. From another, developmental angle, also Griffiths and Stotz [9] endorse the importance of using both internal and external representations; they speak of ‘a larger representational environment which extends beyond the skin’, and claim that ‘culture makes humans as much as the reverse’ [9, p. 45].

Allowing mental states, which are in the external world and thus accessible for any agent around, opens the possibility that other agents also start to use them. Indeed, not only in the individual, single agent case, but also in the social, multi-agent case the extended mind principle can be observed, e.g., one individual creating a pattern in the environment, and one or more other individuals taking this pattern into account in their behaviour. For the human case, examples can be found everywhere, varying from roads, and traffic signs to books or other media, and to many other kinds of cultural achievements. Also in [17] it is claimed that part of the total team knowledge in distributed tasks (such as air traffic control) comprises external memory in the form of artefacts. In this multi-agent case the extended mind principle serves as a way to build a form of social or collective intelligence, that goes beyond (and may even not require) social intelligence based on direct one-to-one communication.

Especially in the case of social animals external mental states created by one individual can be exploited by another individual, or, more general, the creation and maintenance, as well as the exploitation of external mental states can be activities in which a number of individuals participate. For example, presenting slides on a paper with multiple authors to an audience. In such cases the external mental states cross, and in a sense break up, the borders between the individuals and become *shared extended mental states*. An interesting and currently often studied example of collective intelligence is the intelligence shown by an ant colony [2]. Indeed, in this case the external world is exploited as an extended mind by using pheromones. While they walk, ants drop pheromones on the ground. The same or other ants sense these pheromones and follow the route in the direction of the strongest sensing. Pheromones are not persistent for long times; therefore such routes can vary over time.

In [3] the shared extended mind principle is worked out in more detail. The paper focusses on formal analysis and formalisation of the dynamic properties of the processes involved, both at the local level (the basic mechanisms) and the global level (the emerging properties of the whole), and their relationships. A case study in social ant behaviour in which shared extended mind plays an important role is used as illustration.

In the current paper, as an extension to [3], the notion of *representational content* is analysed for mental processes based on the shared extended mind principle. The analysis of notions of representational content of *internal* mental state properties is well-known in the literature on Cognitive Science and Philosophy of Mind. In this literature a relevant internal mental state property m is taken and a representation relation is identified that indicates in which way m relates to properties in the external world or the agent’s interaction with the external world; cf. [1, 10 & 15, pp. 184-210]. For the case of extended mind an extension of the analysis of notions of representational content to *external* state properties is needed. Moreover, for the case of external mental state properties that are *shared*, a notion of *collective* representational content is needed (in contrast to a notion of representational content for a single agent).

Thus, by addressing the ants example and its modelling from an extended mind perspective, a number of challenging new issues on cognitive modelling and representational content are encountered:

- How to define representational content for an *external* mental state property
- How to handle *decay* of a mental state property
- How can *joint* creation of a *shared* mental state property be modelled
- What is an appropriate notion of *collective* representational content of a shared external mental state property
- How can representational content be defined in a case where a behavioural choice depends on *a number of mental state properties*

In this paper these questions are addressed. To this end the shared extended mind principle is analysed in more detail, and a formalisation is provided of its dynamics. It is discussed in particular how a notion of collective representational content for a shared external mental state property can be formulated. In the literature notions of representational content are usually restricted to internal mental states of one individual. The notion of collective representational content developed here extends this in two manners: (1) for external instead of internal mental states, and (2) for groups of individuals instead of single individuals. It is reported how in a case study of social behaviour based on shared extended mind (a simple ant colony) the proposals put forward have been evaluated. The analysis of this case study comprises multi-agent simulation based on identified local dynamic properties, identification of dynamic properties that describe collective representational content of shared extended mind states, and verification of these dynamic properties.

2. State Properties and Dynamic Properties

Dynamics will be described in the next section as evolution of states over time. The notion of state as used here is characterised on the basis of an ontology defining a set of physical and/or mental (state) properties that do or do not hold at a certain point in time. For example, the internal state property 'the agent A has pain', or the external world state property 'the environmental temperature is 7° C', may be expressed in terms of different ontologies. To formalise state property descriptions, an ontology is specified as a finite set of sorts, constants within these sorts, and relations and functions over these sorts. The example properties mentioned above then can be defined by nullary predicates (or proposition symbols) such as pain, or by using n-ary predicates (with $n \geq 1$) like has_temperature(environment, 7). For a given ontology Ont, the propositional language signature consisting of all *state ground atoms* (or *atomic state properties*) based on Ont is denoted by APROP(Ont). The *state properties* based on a certain ontology Ont are formalised by the propositions that can be made (using conjunction, negation, disjunction, implication) from the ground atoms. A *state* S is an indication of which atomic state properties are true and which are false, i.e., a mapping S: APROP(Ont) → {true, false}.

To describe the internal and external dynamics of the agent, explicit reference is made to time. Dynamic properties can be formulated that relate a state at one point in time to a state at another point in time. A simple example is the following dynamic property specification for belief creation based on observation:

‘at any point in time t1 if the agent observes at t1 that it is raining, then there exists a point in time t2 after t1 such that at t2 the agent believes that it is raining’.

To express such dynamic properties, and other, more sophisticated ones, the temporal trace language TTL is used; cf. [11]. To express dynamic properties in a precise manner a language is used in which explicit references can be made to time points and traces. Here *trace or trajectory* over an ontology Ont is a time-indexed sequence of states over Ont. The sorted predicate logic temporal trace language TTL is built on atoms referring to, e.g., traces, time and state properties. For example, ‘in the output state of A in trace γ at time t property p holds’ is formalised by $state(\gamma, t, output(A)) \models p$. Here $\models$ is a predicate symbol in the language, usually used in infix notation, which is comparable to the Holds-predicate in situation calculus. Dynamic properties are expressed by temporal statements built using the usual logical connectives and quantification (for example, over traces, time and state properties). For example the following dynamic property is expressed:

‘in any trace γ, if at any point in time t1 the agent A observes that it is raining, then there exists a point in time t2 after t1 such that at t2 in the trace the agent A believes that it is raining’.

In formalised form:

$$\begin{array}{l}\forall t1\ \ [\ state(\gamma, t1, input(A)) \models agent\_observes\_itsraining\ \ \Rightarrow \\ \qquad \exists t2 \geq t1\ \ state(\gamma, t2, internal(A)) \models belief\_itsraining\ \ ]\end{array}$$

Language abstractions by introducing new (definable) predicates for complex expressions are possible and supported.

A simpler temporal language has been used to specify simulation models. This language (the *leads to* language) offers the possibility to model direct temporal dependencies between two state properties in successive states. This executable format is defined as follows. Let α and β be state properties of the form ‘conjunction of atoms or negations of atoms’, and e, f, g, h non-negative real numbers. In the *leads to* language $\alpha \twoheadrightarrow_{e, f, g, h} \beta$, means:

If state property α *holds for a certain time interval with duration g,*
then after some delay (between e and f) state property β *will hold*
for a certain time interval of length h.

For a precise definition of the *leads to* format in terms of the language TTL, see [14]. A specification of dynamic properties in *leads to* format has as advantages that it is executable and that it can often easily be depicted graphically.

3. Representation for Shared Extended Mind

Originally, the different types of approaches to representational content that have been put forward in the literature on Cognitive Science and Philosophy of Mind, [1, 13 & 15, pp. 191-193, 200-202] are all applicable to internal (mental) states. They have in common that the occurrence of the internal (mental) state property m at a specific point in time is related (by a representation relation) to the occurrence of other state properties, at the same or at different time points. For the *temporal-interactivist* approach [1 & 13] a representation relation relates the occurrence of an internal state property to sets of past and future interaction traces. The *relational specification* approach to representational content is based on a specification of how a representation relation relates the occurrence of an internal state property to properties of states distant in space and time; cf. [15, pp. 200-202]. As mentioned in the Introduction, one of the goals of this paper is to apply these approaches to *shared extended* mental states instead of *internal* mental states.

Suppose p is an external state property used by a collection of agents in their shared extended mind, for example, as an external belief. At a certain point in time this mental state property is created by performing an action a (or maybe a collection of actions) by one or more agents to bring about p in the external world. Given the thus created occurrence of p, at a later point in time any agent can observe p and take this mental state property into account in determining its behaviour. For a representation relation, which indicates representational content for such a mental state property p two possibilities are considered: (1) a representation relation relating the occurrence of p to one or more events in the past (backward), or (2) a representation relation relating the occurrence of p to behaviour in the future (forward). Moreover, for each category, the representation relation can be described by referring to external world state properties, independent of the agent (using the *relational specification* approach), or referring to interaction state properties (e.g., observing, initiating actions) for the agent (using the *temporal-interactivist* approach). In this paper only the relational specification approach is addressed. This approach is applied both backward and forward. For reasons of presentation, first in the upcoming section the (qualitative) case is considered that p is the result of the action of one agent, e.g., the presence of pheromone. Next, the (quantitative) case that p is the result of actions of multiple agents is considered. Here p has a certain degree or level, e.g., a certain accumulated level of pheromone; in decisions levels for a number of such state properties p are taken into account. For the ants case study, the world in which the ants live is described by a labeled graph as depicted in Figure 1. Locations are indicated by A, B,…, and edges by e1, e2,… To represent such a graph the predicate connected_to_via(l0,l1,e) is used. The ants move from location to location via edges; while passing an edge, pheromones are dropped.

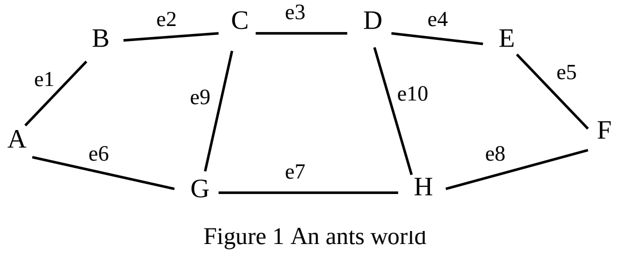

Figure 1 An ants world

3.1 The Qualitative Case

In this section representational content is addressed for the qualitative case. This means that an external state property p is the result of the action of one agent, e.g., the presence of pheromone.

Looking Backward

Looking backward, for the qualitative case the preceding state is the action a by an arbitrary agent, to bring about p. This action a is not an external state property but an interaction state property of this agent. However, this action was performed due to certain circumstances in the world that made the agent do the action. So, the chain of processes can be followed further back to the agent's internal state properties. Still further back it can be followed to the agent's observations that in the past formed the basis of these internal state properties. As these observations concern observations of certain state properties of the external world, we finally arrive at other external world state properties. These external world state properties will be used for the representation relation (conform the relational specification approach). It may be clear that if complex internal processes come between, such a representation relation can become complicated. However, if the complexity of the agent's internal processes is kept relatively simple (as is one of the claims accompanying the extended mind principle), this amounts in a feasible approach.

For the relational specification approach a representation relation can be specified by temporal relationships between the presence of the pheromone (at a certain edge), and other state properties in the past or future. Although the relational specification approach as such does not explicitly exclude the use of state properties related to input and output of the agent, in our approach below the state properties will be limited to external world state properties. As the mental state property itself also is an external world state property, this implies that temporal relationships are provided only between external world state properties.

The pheromone being present at edge e is temporally related to the existence of a state at some time point in the past, namely an agent's presence at e:

If at some time point in the past an agent was present at e,
then after that time point the pheromone was present at edge e.

If the pheromone is present at edge e,
then at some time point in the past an agent was present at e,

Note here that the sharing of the external mental state property is expressed by using explicit agent names in the language and quantification over (multiple) agents. In the usual single agent case of a representation relation, no explicit reference to the agent itself is made. A formalisation is as follows:

$$\begin{aligned}&\forall t1\ \forall l\ \forall e\ \forall a\ \ [\ \text{state}(\gamma, t1) \models \text{is\_at\_edge\_from}(a, e, l)\\&\quad \Rightarrow \exists t2>t1\ \ \text{state}(\gamma, t2) \models \text{pheromone\_at}(e)\ ]\end{aligned}$$

$$\begin{aligned}&\forall t2\ \forall x\ \forall e\ \ [\ \text{state}(\gamma, t2) \models \text{pheromone\_at}(e)\\&\quad \Rightarrow \exists a,\ t1<t2\ \ \text{state}(\gamma, t1) \models \text{is\_at\_edge\_from}(a, e, l)\ ]\end{aligned}$$

Looking Forward

Looking forward, in general the first step is to relate the extended mind state property p to the observation of it by an agent (under certain circumstances c). However, to reach external state properties again, the chain of processes can be followed further through this agent's internal processes to the agent's actions and their effects on the external world.

For the example, the effect of an agent's action based on its observation of the pheromone is that it heads for the direction of the pheromone. So, the representation relation relates the occurrence of the pheromone (at edge e) to the conditional (with condition that it is at the location) fact that the agent heads for the direction of e. The pheromone being present at edge e is temporally related to a conditional statement about the future, namely if an agent later arrives at the location, coming from any direction e', then he will head for direction e:

If the pheromone is present at edge e1,
then if at some time point in the future,
an agent arrives at a location involving e1, coming from any direction e2 ≠ e1,
    then the next direction he will choose is e1.

If a time point t1 exist such that
  at t1 an agent arrives at a location involving e1, coming from any direction e2 ≠e1,
 and if at any time point t2 ≥ t1
  an agent arrives at a location involving e1 coming from any direction e2 ≠ e1,
    then the next direction he will choose is e1,
  then at t1 the pheromone is present at direction e1.

A formalisation is as follows:

$$\begin{aligned}
&\forall t1\ \forall l\ \forall e1\ [\ \text{state}(\gamma, t1) \models \text{pheromone\_at}(e1) \Rightarrow \\
&\quad \forall t2{>}t1\ \forall e2, a\ [e2 \neq e1\ \&\ \text{state}(\gamma, t2) \models \text{is\_at\_location\_from}(a, l, e2) \Rightarrow \\
&\qquad \exists t3{>}t2\ \text{state}(\gamma, t3) \models \text{is\_at\_edge\_from}(a, e1, l)\ \&\ [\forall t4\ t2{<}t4{<}t3 \Rightarrow \text{is\_at\_location\_from}(a, l, e2)]\ ]\ ]
\end{aligned}$$

$$\begin{aligned}
&\forall t1\ \forall l\ \forall e1\ [\ \exists a, e2\ \ e2 \neq e1\ \&\ \text{state}(\gamma, t1) \models \text{is\_at\_location\_from}(a, l, e2)\ \& \\
&\quad [\forall t2{\geq}t1\ \forall a, e2\ [e2 \neq e1\ \& \\
&\qquad \text{state}(\gamma, t2) \models \text{is\_at\_location\_from}(a, l, e2) \Rightarrow \\
&\qquad\quad \exists t3{>}t2\ \text{state}(\gamma, t3) \models \text{is\_at\_edge\_from}(a, e1, l)\ \&\ [\forall t4\ t2{<}t4{<}t3 \Rightarrow \text{is\_at\_location\_from}(a, l, e2)]\ ]\ ] \\
&\quad \Rightarrow\ \text{state}(\gamma, t1) \models \text{pheromone\_at}(e1)\ ]
\end{aligned}$$

3.2 The Quantitative Case

The quantitative, accumulating case allows us to consider certain levels of a mental state property p; in this case a mental state property is involved that is parameterised by a number: it has the form p(r), where r is a number, denoting that p has level r. This differs from the above in that it is now possible to model: (1) joint creation of p: multiple agents together bring about a certain level of p, each contributing a part of the level, (2) by decay levels may decrease over time, (3) behaviour may be based on a number of state properties with different levels, taking into account their relative values, e.g., by determining the highest level of them. For the ants example, for each choice point multiple directions are possible, each with a different pheromone level; the choice is made for the direction with the highest pheromone level (ignoring the direction the ant just came from).

Looking Backward

To address the backward quantitative case (i.e., the case of joint creation of a mental state property), the representation relation is analogous to the one described above, but now involves not the presence of one agent at one past time point, but a summation over multiple agents at different time points. Moreover a decay rate r with $0 < r < 1$ is used to indicate that after each time unit only a fraction r is left. Thus for the ants example in mathematical terms the following property is expressed:

There is an amount v of pheromone at edge e, if and only if there is a history such that at time point 0 there was ph(0, e) pheromone at e, and for each time point k from 0 to t a number dr(k, e) of ants dropped pheromone, and $v = ph(0, e) * r^{t} + \Sigma_{k=0}^{t} dr(t-k, e) * r^{k}$

A formalisation of this property in the logical language TTL is as follows:

∀t ∀e ∀p state(γ, t) |= pheromones_at(e, v) ⇔
$\Sigma_{k=0}^{t} \Sigma_{a=ant1}^{ants}$ case(state(γ, k) |= is_at_edge(a, e), 1, 0) * $r^{t-k} = v$

Here for any formula f, the expression case(f, v1, v2) indicates the value v1 if f is true, and v2 otherwise.

Looking Forward

The forward quantitative case involves a behavioural choice that depends on the relative levels of multiple mental state properties. This makes that at each choice point the representational content of the level of one mental state property is not independent of the level of the other mental state properties involved at the same choice point. Therefore it is only possible to provide representational content for the combined mental state property involving all mental state properties involved in the behavioural choice. Thus for the ants example the following property is specified:

If at time t1 the amount of pheromone at edge e1 is maximal with respect to the amount of pheromone at all other edges connected to that location, except the edge that brought the ant to the location,
then, if an ant is at that location l at time t1,
then the next direction the ant will choose at some time t2 > t1 is e1.

If at time t1 an ant is at location 1 and
for every ant arriving at that location 1 at time t1,
the next direction it will choose at some time t2 > t1 is e1,
then the amount of pheromone at edge e1 is maximal with respect to the amount of pheromone at all other edges connected to that location l, except the edge that brought the ant to the location.

A formalisation of this property in TTL is as follows:

∀t1,l,l1,e1,e2,i1
[e1≠e2 &
state(γ, t1) |= connected_to_via(l, l1, e1) &
state(γ, t1) |= pheromones_at(e1, i1) &
[∀l2≠l1, e3≠e2 [state(γ, t1) |= connected_to_via(l, l2, e3) ⇒
∃i2 [0≤i2<i1 & state(γ, t1) |= pheromones_at(e3, i2)]]
⇒ ∀a [state(γ, t1) |= is_at_location_from(a, l, e2) ⇒
∃t2>t1 state(γ, t2) |= is_at_edge_from(a, e1, l) & [∀t3 t1<t3<t2 ⇒ is_at_location_from(a, l, e2)]]]]

∀t1, l,l1,e1,e2
[e1≠e2 &
state(γ, t1) |= connected_to_via(l, l1, e1) &
∃a state(γ, t1) |= is_at_location_from(a, l, e2) &
∀a [state(γ, t1) |= is_at_location_from(a, l, e2) ⇒
∃t2>t1 state(γ, t2) |= is_at_edge_from(a, e1, l) & [∀t3 t1<t3<t2 ⇒ is_at_location_from(a, l, e2)]]]
⇒ ∃i1 [state(γ, t1) |= pheromones_at(e1, i1) &
[∀l2≠l1, e3≠e2 [state(γ, t1) |= connected_to_via(l, l2, e3) ⇒
∃i2 [0≤i2≤i1 & state(γ, t1) |= pheromones_at(e3, i2)]]]]

4. A Simulation Model of Shared Extended Mind

In [3] a simulation model of an ant society is specified in which shared extended mind plays an important role. This model is based on local dynamic properties, expressing the basic mechanisms of the process. In this section, a selection of these local properties is presented, and a resulting simulation trace is shown. In the next section it will be explained how the representation relations specified earlier can be verified against such simulation traces. Here a is a variable that stands for ant, l for location, e for edge, and i for pheromone level.

LP5 (Selection of Edge)

This property models (part of) the edge selection mechanism of the ants. It expresses that, when an ant observes that it is at location l, and there are two edges connected to that location, then the ant goes to the edge with the highest amount of pheromones. Formalisation:

observes(a, is_at_location_from(l, e0)) and neighbours(l, 3) and connected_to_via(l, l1, e1) and observes(a, pheromones_at(e1, i1)) and connected_to_via(l, l2, e2) and observes(a, pheromones_at(e2, i2)) and e0 ≠ e1 and e0 ≠ e2 and e1 ≠ e2 and i1 > i2 •→ to_be_performed(a, go_to_edge_from_to(e1, l1))

LP9 (Dropping of Pheromones)

This property expresses that, if an ant observes that it is at an edge e from a location l to a location l1, then it will drop pheromones at this edge e. Formalisation:

observes(a, is_at_edge_from_to(e, l, l1)) •→ to_be_performed(a, drop_pheromones_at_edge_from(e, l))

LP13 (Increment of Pheromones)

This property models (part of) the increment of the number of pheromones at an edge as a result of ants dropping pheromones. It expresses that, if an ant drops pheromones at edge e, and no other ants drop pheromones at this edge, then the new number of pheromones at e becomes i*decay+incr. Here, i is the old number of pheromones, decay is the decay factor, and incr is the amount of pheromones dropped. Formalisation:

to_be_performed(a1, drop_pheromones_at_edge_from(e, l1)) and ∀l2 not to_be_performed(a2, drop_pheromones_at_edge_from(e, l2)) and ∀l3 not to_be_performed(a3, drop_pheromones_at_edge_from(e, l3)) and a1 ≠ a2 and a1 ≠ a3 and a2 ≠ a3 and pheromones_at(e, i) •→ pheromones_at(e, i*decay+incr)

LP14 (Collecting of Food)

This property expresses that, if an ant observes that it is at location F (the food source), then it will pick up some food. Formalisation:

observes(a, is_at_location_from(l, e)) and food_location(l) •→ to_be_performed(a, pick_up_food)

LP18 (Decay of Pheromones)

This property expresses that, if the old amount of pheromones at an edge is i, and there is no ant dropping any pheromones at this edge, then the new amount of pheromones at e will be i*decay. Formalisation:

pheromones_at(e, i) and ∀a,l not to_be_performed(a, drop_pheromones_at_edge_from(e, l)) •→ pheromones_at(e, i*decay)

A special software environment has been created to enable the simulation of executable models. Based on an input consisting of dynamic properties in *leads to* format, the software environment generates simulation traces. An example of such a trace can be seen in Figure 2. Time is on the horizontal axis, the state properties are on the vertical axis. A dark box on top of the line indicates that the property is true during that time period, and a lighter box below the line indicates that the property is false. This trace is based on all local properties identified.

Because of space limitations, in the example situation depicted in Figure 2, only three ants are involved. However, similar experiments have been performed with a population of 50 ants. Since the abstract way of modelling used for the simulation is not computationally expensive, also these simulations took no more than 30 seconds.

As can be seen in Figure 2 there are two ants (ant1 and ant2) that start their search for food immediately, whereas ant3 comes into play a bit later, at time point 3. When ant1 and ant2 start their search, none of the locations contain any pheromones yet, so basically they have a free choice where to go. In the current example, ant1 selects a rather long route to the food source (via locations A-B-C-D-E-F), whilst ant2 chooses a shorter route (A-G-H-F). Note that, in the current model, a fixed route preference (via the attractiveness predicate) has been assigned to each ant for the case there are no pheromones yet. After that, at time point 3, ant3 starts its search for food. At that moment, there are trails of pheromones leading to both locations B and G, but these trails contain exactly the same number of pheromones. Thus, ant3 also has a free choice among location B and G, and chooses in this case to go to B. Meanwhile, at time point 18, ant2 has arrived at the food source (location F). Since it is the first to discover this location, the only present trail leading back to the nest, is its own trail. Thus ant2 will return home via its own trail. Next, when ant1 discovers the food source (at time point 31), it will notice that there is a trail leading back that is stronger than its own trail (since ant2 has already walked there twice: back and forth, not too long ago). As a result, it will follow this trail and will keep following ant2 forever. Something similar holds for ant3. The first time that it reaches the food source, ant3 will still follow its own trail, but some time later (from time point 63) it will also follow the other two ants. To conclude, eventually the shortest of both routes is shown to remain, whilst the other route evaporates. Other simulations, in particular for small ant populations, show that it is important that the decay parameter of the pheromones is not too high. Otherwise, the trail leading to the nest has evaporated before the first ant has returned, and all ants get lost!

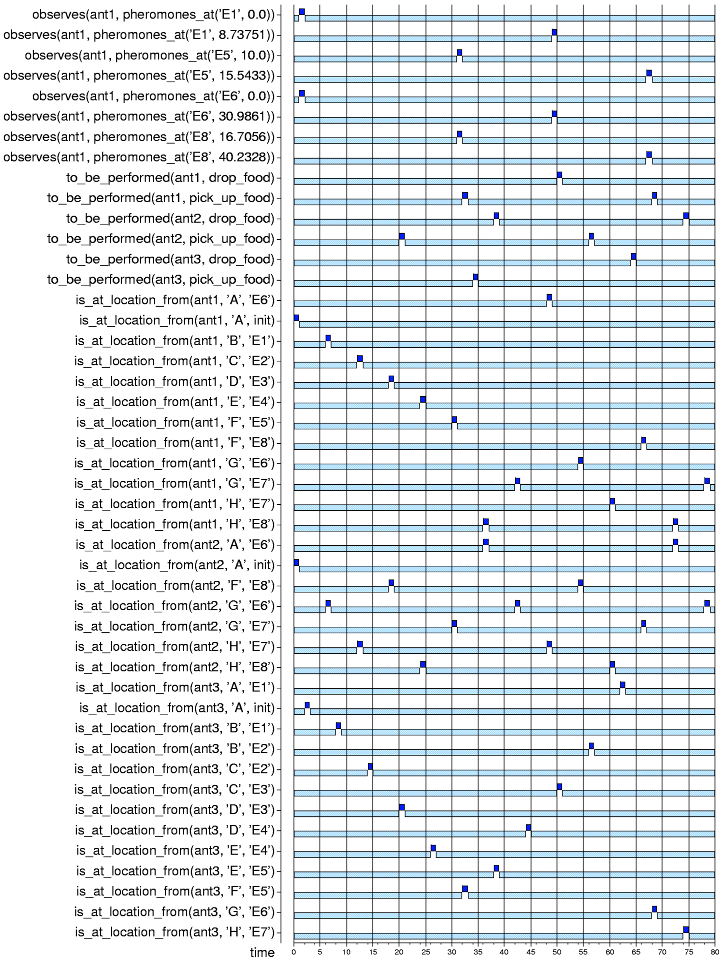

Figure 2 Simulation trace of the dynamics of the ants behaviour

5. Verification

In addition to the simulation software, a software environment has been developed that enables to check dynamic properties specified in TTL against simulation traces. This software environment takes a dynamic property and one or more (empirical or simulated) traces as input, and checks whether the dynamic property holds for the traces. Traces are represented by sets of Prolog facts of the form

```
holds(state(m1, t(2)), a, true).
```

where `m1` is the trace name, `t(2)` time point 2, and `a` is a state formula in the ontology of the component's input. It is indicated that state formula `a` is true in the component's input state at time point t2. The programme for temporal formula checking basically uses Prolog rules for the predicate `sat` that reduce the satisfaction of the temporal formula finally to the satisfaction of atomic state formulae at certain time points, which can be read from the trace representation. Examples of such reduction rules are:

```
sat(and(F,G)) :- sat(F), sat(G).

sat(not(and(F,G))) :- sat(or(not(F), not(G))).

sat(or(F,G)) :- sat(F).

sat(or(F,G)) :- sat(G).

sat(not(or(F,G))) :- sat(and(not(F), not(G))).
```

Using this environment, the formal representation relations presented in Section 3.2 have been automatically checked against traces like the one depicted in Section 4. The duration of these checks varied from 1 to 10 seconds, depending on the complexity of the formula (in particular, the backward representation relation has a quite complex structure, since it involves reference to a large number of events in the history). All these checks turned out to be successful, which validates (for the given traces at least) our choice for the representational content of the shared extended mental state property `pheromones_at(e, v)`. However, note that these checks are only an empirical validation, they are no exhaustive proof as, e.g., model checking is. Currently, the possibilities are explored to combine TTL with existing model checking techniques.

In addition to simulated traces, the checking software allows to check dynamic properties against other types of traces as well. In the future, the representation relations specified in this paper will be checked against traces resulting from other types of ants simulations, and possibly against empirical traces.

6. Discussion

The extended mind perspective introduces a high-level conceptualisation of agent-environment interaction processes. By modelling the ants example from an extended mind perspective, the following challenging issues on cognitive modelling and representational content were encountered:

1. How to define representational content for an external mental state property
2. How to handle decay of a mental state property
3. How can joint creation of a shared mental state property be modelled
4. What is an appropriate notion of collective representational content of a shared external mental state property
5. How can representational content be defined in a case where a behavioural choice depends on a number of mental state properties

These questions were addressed in this paper. For example, modelling joint creation of mental state properties (3.) was made possible by using relative or leveled mental state properties, parameterised by numbers. Each contribution to such a mental state property was modelled by addition to the level indicated by the number. Collective representational content (4.) from a looking backward perspective was defined by taking into account histories of such contributions. Collective representational content from a forward perspective was defined taking into account multiple parameterised mental state properties, corresponding to the alternatives for behavioural choices, with their relative weights. In this case it is not possible to define representational content for just one of these mental state properties, but it is possible to define it for their combination or conjunction (5.).

The high-level conceptualisation has successfully been formalised and analysed in a logical manner. The formalisation enables simulation and automated checking of dynamic properties of traces or sets of traces, in particular of the representation relations.

For future research, it is planned to define the general concept of extended mind in a more precise way. This will make the distinction between extended mind states and other external world states, which is currently not always clear, more concrete. In addition, the approach will be applied to several other cases of extended mind. For example, can the work be related to AI planning representations, traffic control, knowledge representation of negotiation, and to the concept of “shared knowledge” in knowledge management?

Acknowledgements

The authors are grateful to Lourens van der Meij for his contribution to the development of the software environment, and to an anonymous referee for some valuable comments on an earlier version of this paper.

References

1. Bickhard, M.H. Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 5, 1993, pp. 285-333.
2. Bonabeau, J. Dorigo, M. and Theraulaz, G. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.
3. Bosse, T., Jonker, C.M., Schut, M.C., and Treur, J. Simulation and Analysis of Shared Extended Mind. *Proceedings of the First Joint Workshop on Multi-Agent and Multi-Agent-Based Simulation, MAMABS’04*. To appear, 2004.
4. Bosse, T., Jonker, C.M., and Treur, J. Simulation and analysis of controlled multi-representational reasoning processes. *Proc. of the Fifth International Conference on Cognitive Modelling, ICCM’03*. Universitats-Verlag Bamberg, 2003, pp. 27-32.
5. Clark, A. *Being There: Putting Brain, Body and World Together Again*. MIT Press, 1997.
6. Clark, A. Reasons, Robots and the Extended Mind. In: *Mind & Language*, vol. 16, 2001, pp. 121-145.

7. Clark, A., and Chalmers, D. The Extended Mind. In: *Analysis*, vol. 58, 1998, pp. 7-19.
8. Dennett, D.C. *Kinds of Mind: Towards an Understanding of Consciousness*, New York: Basic Books, 1996.
9. Griffiths, P. and Stotz, K. How the mind grows: a developmental perspective on the biology of cognition. *Synthese*, vol.122, 2000, pp. 29--51.
10. Jacob, P. *What Minds Can Do: Intentionality in a Non-Intentional World*. Cambridge University Press, Cambridge, 1997.
11. Jonker, C.M., and Treur, J. Compositional Verification of Multi-Agent Systems: a Formal Analysis of Pro-activeness and Reactiveness. *International Journal of Cooperative Information Systems,* vol. 11, 2002, pp. 51-92.
12. Jonker, C.M., and Treur, J. Analysis of the Dynamics of Reasoning Using Multiple Representations. In: W.D. Gray and C.D. Schunn (eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society, CogSci 2002.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc., 2002, pp. 512-517.
13. Jonker, C.M., and Treur, J. A Temporal-Interactivist Perspective on the Dynamics of Mental States. *Cognitive Systems Research Journal*, vol. 4, 2003, pp. 137-155.
14. Jonker, C.M., Treur, J., and Wijngaards, W.C.A., A Temporal Modelling Environment for Internally Grounded Beliefs, Desires and Intentions. *Cognitive Systems Research Journal*, vol. 4, 2003, pp. 191-210.
15. Kim, J. *Philosophy of Mind*. Westview Press, 1996.
16. Menary, R. (ed.) *The Extended Mind,* Papers presented at the Conference *The Extended Mind - The Very Idea: Philosophical Perspectives on Situated and Embodied Cognition*, University of Hertfordshire, 2001. John Benjamins, 2004, to appear.
17. Scheele, M. Team Action: A Matter of Distribution. Distributed Cognitive systems and the Collective Intentionality they Exhibit. *Third International Conference on Collective Intentionality*, 2002.