Formal Interpretation and Analysis of Collective Intelligence as Individual Intelligence

Tibor Bosse and Jan Treur

Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081a, NL-1081 HV Amsterdam, The Netherlands
{tbosse, treur}@cs.vu.nl

Abstract. This paper addresses the question to what extent a process involving multiple agents that shows some form of collective intelligence can be interpreted as a single agent. The question is answered by formal analysis. It is shown for an example process how it can be conceptualised formalised and simulated in two different manners: from a single agent (or cognitive) and from a multi-agent (or social) perspective. Moreover, it is shown how an ontological mapping can be formally defined between the two formalisations, and how this mapping can be extended to a mapping of dynamic properties. Thus it is shown how collective behaviour can be interpreted in a formal manner as single agent behaviour.

1 Introduction

Many processes in the world can be conceptualised using an agent metaphor, as a single agent (or cognitive) process or a multi-agent (or social) process. Especially for processes that are distributed, it is natural to describe them as a group of interacting agents. If a group of agents acts in a coherent way, however, one is often tempted to intuitively and informally interpret the process in singular form as a collective, and, in fact, as one individual (super)agent. The question addressed in this paper is whether in certain cases such an informal interpretation of a multi-agent system, acting in a collective manner, as an individual can be supported by a formal analysis. The approach to address this question is by formally defining an interpretation mapping between a conceptualisation of a process as a multi-agent system and a conceptualisation of the same process as an individual.

The prerequisites to undertake such a formal analysis concern formalisations of the notion of agent, single agent behaviour and multi-agent behaviour, and the notion of interpretation mapping. More specifically, what is needed is a formal notion of what an agent is in the sense of

- distinctions between the agent's internal mental processes, the agent's body, and the agent's environment
- interactions and relationships between mental aspects and body aspects
- interactions and relationships between agent and environment, including interactions with other agents

Furthermore, formalisations of single agent behaviour and multi-agent behaviour are needed that cover

- the externally observable behaviour

- the underlying internal processes

Moreover, a formal notion of interpretation mapping of a single agent conceptualisation into a multi-agent conceptualisation is needed that

- maps ontological concepts describing a conceptualisation of a process from an individual perspective to ontological concepts describing a conceptualisation of the same a process from a multi-agent perspective
- covers mapping of individual mental state properties for the single agent conceptualisation to shared mental state properties for the multi-agent conceptualisation
- covers the mapping of dynamic aspects of single agent behaviour onto those of multi-agent behaviour

In this paper for these three notions formalisations are provided and used to indeed achieve a formalisation of how a collective can be formally interpreted as an individual.

The formalisation is evaluated for the case of collective behaviour of an ants colony. The intelligence shown by ant colonies are an interesting and currently often studied example of collective intelligence [1], [4], [6]. In this case by using pheromones the external world is exploited as a form of extended mind; cf. [2], [3], [5], [10], [11]. The analysis of this case study comprises on the one hand a multi-agent model, simulation based on identified local dynamic properties, and identification of dynamic properties for the overall process. On the other hand the same is done for an alternative model based on a single agent with internal mental states, and the two models are related to each other via the interpretation mapping.

In Section 2, a formalisation of basic agent concepts will be introduced. Section 3 explains, using a simple example, the idea of the basic formal ontology mapping between state properties in a single agent conceptualisation and state poperties in a multi-agent conceptualisation. In Section 4 this notion of basic interpretation mapping of state properties is applied to two conceptualisations of the more complex ant colony example, the central case study in the paper. Section 5 discusses the dynamics for the two conceptualisations of the ant colony example in more detail, which leads to formal specification of executable local dynamic properties that have been used for simulation. In Section 6 the basic interpretation mapping for state properties is extended to dynamic properties, thus obtaining an interpretation mapping between the two conceptualisations of the dynamics of the example ants colony process. Section 7 is a final discussion.

2 Basic Agent Concepts

The agent perspective entails a distinction between the following different types of ontologies:

- an ontology for *internal mental properties* of the agent A (MentOnt(A)),
- for properties of the agent's (physical) *body* (BodyOnt(A)),
- for properties of the (sensory or communication) *input* (InOnt(A))
- for properties of the (action or communication) *output* (OutOnt(A)) of the agent, and
- for properties of of the *external* world (ExtOnt(A)).

For example, the property 'the agent A feels pain' may belong to MentOnt(A), resp. BodyOnt(A), whereas 'it is raining' and 'the outside temperature is 7° C' may belong to

ExtOnt(A). The agent input ontology InOnt defines state properties for received perception or communication, as an in-between step from environment or body state properties to internal mental state properties, the agent output ontology OutOnt defines state properties that indicate initiations of actions or communications of the agent, as an in-between step from internal mental state properties to environment or body state properties. The combination of InOnt and OutOnt is the *agent interaction ontology*, defined by $\text{InteractionOnt} = \text{InOnt} \cup \text{OutOnt}$.

To formalise state property descriptions of the types introduced above, ontologies are specified in a (many-sorted) first order logical format: 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 itsraining, or by using n-nary predicates (with $n \geq 1$) like has_pain(A) and has_temperature(environment, 7).

For a given ontology Ont, the propositional language signature consisting of all *state ground atoms* 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 (using conjunction, negation, disjunction, implication) from the ground atoms. 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. In other words, a *state* S is an indication of which atomic state properties are true and which are false, i.e., a mapping $S: \text{APROP(Ont)} \rightarrow \{\text{true, false}\}$.

To describe the internal and externally observable dynamics of the agent, explicit reference is made to time. Dynamics will be described as evolution of states over 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 informally stated dynamic property for belief creation based on observation:

'if the agent observes at t1 that it is raining, then the agent will believe that it is raining'.

To express such dynamic properties, and other, more sophisticated ones, the sorted predicate logic *Temporal Trace Language* (TTL) is used [7]. Here, a *trace* over an ontology Ont is a time-indexed sequence of states over Ont. TTL is built on atoms referring to, e.g., traces, time and state properties. For example, 'in trace γ at time t property p holds' is formalised by $\text{state}(\gamma, t) \models p$. Here |= 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 dynamic property put forward above can be expressed in a more structured semiformal manner as:

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

In formalised TTL form it looks as follows:

$$\forall \gamma \, \forall t1 \, [\, \text{state}(\gamma, t1) \models \text{observes(A, itsraining)} \;\Rightarrow\; \exists t2 \geq t1 \;\; \text{state}(\gamma, t2) \models \text{belief(A, itsraining)} \,]$$

Based on TTL, a simpler temporal language has been defined to specify simulation models. This language (the *leads to* language) enables 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 $\alpha \bullet\!\!\twoheadrightarrow_{e, f, g, h} \beta$, 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 [7]. 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 The Basic Interpretation Mapping

In this section it is discussed how a conceptualisation based on a single agent and individual (internal) mental state properties can formally be mapped onto a conceptualisation based on multiple agents and shared (for the sake of simplicity assumed external) mental state properties. Here this ontological mapping is only given in its basic form, for the state properties. In Section 6 the basic mapping is extended to temporal expressions describing behaviour.

First, consider Figure 1. This figure depicts a simple case of a single agent A with behaviour based on an individual internal mental state property m1. The solid arrows depict temporal *leads to* relationships. Mental state property m1 (temporally) depends on observations of three world state properties c1, c2, c3. Moreover, action a1 depends on m1.

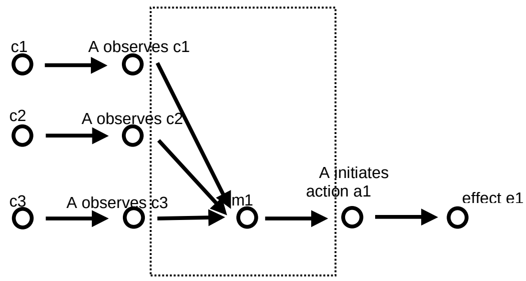

Figure 1 Single Agent behaviour based on an internal mental state

Now consider Figure 2. This figure depicts a group of agents A1, A2, A3, A4 with behaviour based on a physical external world state property m2 that serves as a shared external mental state property.

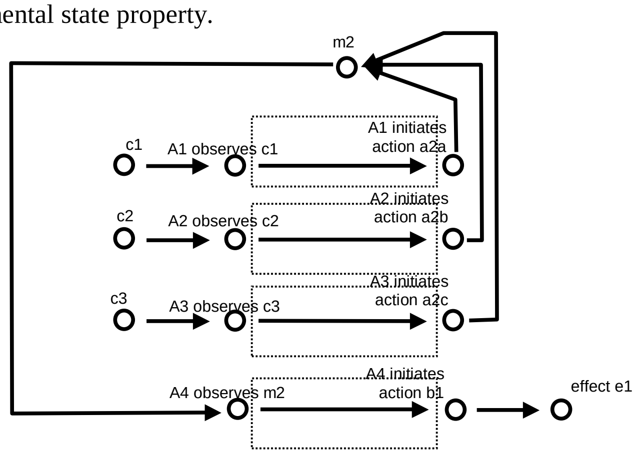

Figure 2 Multi-Agent behaviour based on a shared external mental state

To create this shared mental state property, actions a2a, a2b, a2c of the agents A1, A2, A3 are needed, and to show the behaviour, first an observation of m2 by agent A4 is needed. Note that here the internal processing is chosen as simple as possible: stimulus response. The interaction between agent and external world is a bit more complex: compared to a single agent perspective with internal mental state m1, extra actions of some of the agents needed to create the external mental state property m2, and additional observations are needed to observe it.

To make the similarity between the two different cognitive processes more precise, the following mapping from the nodes (state properties) in Figure 1 onto nodes in Figure 2 can be made (see Figure 3):

External world state properties

φ:	c1	→	c1
φ:	c2	→	c2
φ:	c3	→	c3
φ:	effect e1	→	effect e1

Observation state properties

φ:	A observes c1	→	A1 observes c1
φ:	A observes c2	→	A2 observes c2
φ:	A observes c3	→	A3 observes c3

Action initiation state properties

φ:	A initiates action a1	→	A4 initiates action b1

Mental state property to external world state property

φ:	m1	→	m2

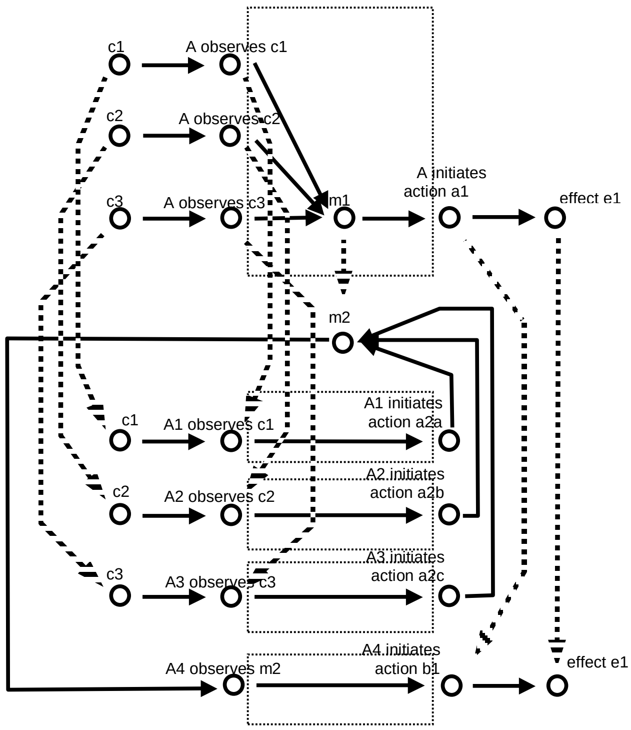

Figure 3 Isomorphism relationship between shared extended mind and individual mental state

Note that in this case, for simplicity it is assumed that each observation of A is an observation of exactly one of the Ai, and the same for actions.

This mapping φ, indicated by the vertical dotted arrows in Figure 3, preserves the temporal dependencies in the form of *leads to* relationships (the solid arrows) and provides an (isomorphic, in the mathematical sense) embedding of a cognitive process based on internal mind into a cognitive process based on extended mind.

In their paper about extended mind, Clark and Chalmers [3] point at the similarity between cognitive processes in the head and some processes involving the external world. This similarity can be used as an indication that these processes can be considered extended cognitive processes or extended mind:

> If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. Cognitive processes ain't (all) in the head! [3], Section 2. (…)
> One can explain my choice of words in Scrabble, for example, as the outcome of an extended cognitive process involving the rearrangement of tiles on my tray. Of course, one could always try to explain my action in terms of internal processes and a long series of "inputs" and "actions", but this explanation would be needlessly complex. If an isomorphic process were going on in the head, we would feel no urge to characterize it in this cumbersome way. (…) In a very real sense, the re-arrangement of tiles on the tray is not part of action; it is part of *thought*. [3], Section 3.

Clark and Chalmers [3] use the isomorphism to a process 'in the head' as one of the criteria to consider external and interaction processes as cognitive, or mind processes. As the shared mental state property m2 is modelled as an external state property, this 'isomorphism principle' is formalised in Figure 3 for a simple example of such an isomorphism. Note that the process from m1 to action a1, modelled as one step in the single agent, internal case, is mapped onto a process from m2 via A4 observes m2 to A4 initiates action b1, in the external case modelled as a two-step process. So the isomorphism is an embedding in one direction, not a bidirectional isomorphism, simply because on the multi-agent side, the observation state for A4 observing m2 has no counterpart in the single agent, internal case (and the same for the agents A1, A2, A3 initiating actions a2a, a2b, a2c).

Notice that the mapping φ is a (formal) mapping between state properties. However, it was already put forward that temporal *leads to* relations are preserved under φ, so the mapping can be extended to a mapping of *leads to* properties onto *leads to* properties. From a more general perspective, it can be analysed how far the mapping φ can be extended to a (formal) mapping from dynamic properties to dynamic properties expressed in TTL. This will be addressed in detail in Section 6.

4 Two Conceptualisations and their Mapping

The general formalisation perspective put forward in previous sections has been evaluated for a case study: a process of collective ant behaviour. For this example process two conceptualisations have been made, one from a multi-agent (or social) perspective, and one for a single agent (or cognitive) perspective.

The world in which the ants live is described by a labeled graph as depicted in Figure 4. Locations are indicated by A, B,…, and edges by E1, E2,… To represent such a graph the predicate connected_to_via(l0,l1,e1) is used. The ants move from location to location via edges; while passing an edge, pheromones are dropped. The same or other ants sense these pheromones and follow the route in the direction of the strongest concentration. Pheromones evaporate over time; therefore such routes can vary over time. The goal of the ants is to find food and bring this back to their nest. In this example there is only one nest (location A) and one food source (location F).

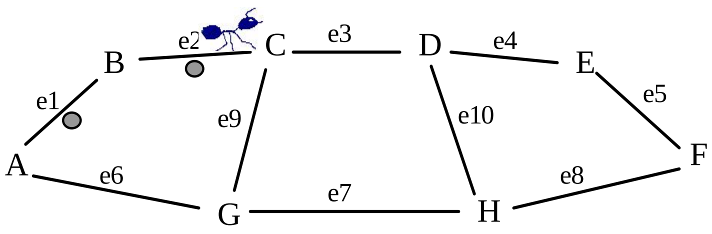

Figure 4 An ants world

4.1 Multi-Agent Conceptualisation

The example process conceptualised from a multi-agent perspective concerns multiple agents (the ants), each of which has input (to observe) and output (for moving and dropping pheromones) states, and a physical body which is at certain positions over time, but no internal mental state properties (they are assumed to act purely by stimulus-response behaviour). An overview of the formalisation of the state properties of this multi-agent conceptualisation is shown in Table 1.

	Multi-Agent Conceptualisation
	body positions in world:
pheromone level at edge e is i	pheromones_at(e, i)
ant a is at location l coming from e	is_at_location_from(a, l, e)
ant a is at edge e to l2 coming from location l1	is_at_edge_from_to(a, e, l1, l2)
ant a is carrying food	is_carrying_food(a)
	world state properties:
edge e connects location l1 and l2	connected_to_via(l1, l2, e)
location l has i neighbours	neighbours(l, i)
edge e is most attractive for ant a coming from location l	attractive_direction_at(a, l, e)
	input state properties:
ant a observes that it is at location l coming from edge e	observes(a, is_at_location_from(l, e))
ant a observes that it is at edge e to l2 coming from location l1	observes(a, is_at_edge_from_to(e, l1, l2))
ant a observes that edge e has pheromone level i	observes(a, pheromones_at(e, i))
	output state properties:
ant a initiates action to go to edge e to l2 coming from location l1	to_be_performed(a, go_to_edge_from_to(e, l1, l2))
ant a initiates action to go to location l coming from edge e	to_be_performed(a, go_to_location_from(l, e))
ant a initiates action to drop pheromones at edge e coming from location l	to_be_performed(a, drop_pheromones_at_edge_from(e, l))
ant a initiates action to pick up food	to_be_performed(a, pick_up_food)
ant a initiates action to drop food	to_be_performed(a, drop_food)

Table 1 Multi-Agent conceptualisation: state properties

4.2 Single-Agent Conceptualisation

The conceptualisation of the example process from a single agent perspective (Superant S), however, takes into account one body, of which each ant is part (for convenience we call them the ‘paws’ of this body). Also the pheromone levels at the edges are part of the body.

Single Agent Conceptualisation	
mental state properties:	
belief(S, relevance_level(e, i))	belief on the relevance level i of an edge e
body position in world:	
has_paw_at_location_from(S, p, l, e) has_paw_at_edge_from_to(S, p, e, l1, l2) is_carrying_food_with_paw(S, p)	position of paw p at location l coming from edge e position of paw p at edge e to l2 coming from location l1 paw p is carrying food
world state properties:	
connected_to_via(l1, l2, e) neighbours(l, i) attractive_direction_at(p, l, e)	edge e connects location l1 and l2 location l has i neighbours edge e is most attractive for paw p coming from location l
input state properties:	
observes(S, has_paw_at_location_from(p, l, e)) observes(S, has_paw_at_edge_from_to(p, e, l1, l2))	S observes that paw p is at location l coming from edge e S observes that paw p is at edge e to l2 coming from location l1
output state properties:	
to_be_performed(S, move_paw_to_edge_from_to(p, e, l1, l2)) to_be_performed(S, move_paw_to_location_from (p, l, e)) to_be_performed(S, pick_up_food_with_paw(p)) to_be_performed(S, drop_food_with_paw(p))	S initiates action to move paw p from location l1 to edge e to l2 S initiates action to move paw p from edge e to location l S initiates action to pick up food with paw p S initiates action to drop food with paw p

Table 2 Single Agent conceptualisation: state properties

The body position of this agent in the world is defined by the collection of positions of each of the paws. Mental state properties for this single agent occur in the form of beliefs that a certain edge has a certain relevance level (realised in the body by the pheromone levels). Input of the single agent is defined by the collection of inputs of the ants at each of the paws. Output is defined by initiation of movements of one or more of the paws. Notice that in this case dropping pheromones is not an action, but an internal body process to create or update the proper beliefs by creating or updating their realisation in the body. An overview of the formalisation of the state properties of the multi-agent conceptualisation is shown in Table 2. Note that there S stands for the Superant.

4.3 Mapping between Conceptualisations

The two conceptualisations described in Sections 4.1 and 4.2 are two conceptualisations of one and the same example process. A concept in any of the two conceptualisations in principle has a one-to-one correspondence to an aspect of this example process which can be considered the informal semantics of the concept (in our case the concept is formalised); see the double arrows in Figure 5.

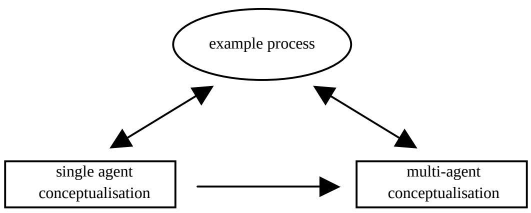

Figure 5 Two conceptualisations and their mapping

Given these one-to-one correspondences, a mapping from the single agent conceptualisation to the multi-agent conceptualisation can be made as follows:

1) Take any state property c belonging to the single agent conceptualisation
2) Identify to what aspect a of the example process this state property corresponds
3) Identify to which state property d in the multi-agent conceptualisation this aspect a corresponds
4) Map c to d.

If this approach works, then a mapping is obtained that is sincere with respect to the example process: the state property d to which c is mapped corresponds to the same aspect a of the process as c, and therefore will be true (for the informal semantics) if and only if c is. The approach can also fail. It can fail in 2) if state properties are used in the single agent conceptualisation that have no counterpart in the example process. It can fail in 3) if in the single agent conceptualisation aspects of the process are covered that are left out of consideration in the other conceptualisation. Actually such aspects exist the other way around: there are aspects of the process, such as observing the pheromones covered by the multi-agent conceptualisation, but not by the single agent conceptualisation. Therefore such a mapping is not possible from right to left in Figure 5 (see also Figure 3 in Section 3, where the mapping is not bijective either). However, a mapping from left to right (single agent to multi-agent conceptualisation), is possible. It is shown in Table 3. Note that there S stands for the Superant, and paw p corresponds to ant a.

Single Agent Conceptualisation	**Multi-Agent Conceptualisation**
belief(S, relevance_level(e, i))	pheromones_at(e, i)
has_paw_at_location_from(S, p, l, e)	is_at_location_from(a, l, e)
has_paw_at_edge_from_to(S, p, e, l1, l2)	is_at_edge_from_to(a, e, l1, l2)
is_carrying_food_with_paw(S, p)	is_carrying_food(a)
connected_to_via(l1, l2, e)	connected_to_via(l1, l2, e)
neighbours(l, I)	neighbours(l, i)
attractive_direction_at(p, l, e)	attractive_direction_at(a, l, e)
observes(S, has_paw_at_location_from(p, l, e))	observes(a, is_at_location_from(l, e))
observes(S, has_paw_at_edge_from_to(p, e, l1, l2))	observes(a, is_at_edge_from_to(e, l1, l2))
---	observes(a, pheromones_at(e, i))
to_be_performed(S, move_paw_to_edge_from_to(p, e, l1, l2))	to_be_performed(a, go_to_edge_from_to(e, l1, l2))
to_be_performed(S, move_paw_to_location_from (p, l, e))	to_be_performed(a, go_to_location_from(l, e))
---	to_be_performed(a, drop_pheromones_at_edge_from(e, l))
to_be_performed(S, pick_up_food_with_paw(p))	to_be_performed(a, pick_up_food)
to_be_performed(S, drop_food_with_paw(p))	to_be_performed(a, drop_food)

Table 3 Mapping between state properties

5 Two Simulation Models

The two conceptualisations introduced above have been used to create two simulation models for collective ant behaviour: one from a multi-agent (social) perspective and one from a single agent (cognitive) perspective. The basic building blocks of the model were dynamic properties in *leads to* format, specifying the local mechanisms of the process. Examples of such dynamic properties (for the *multi-agent case*) are the following:

LP5 (Selection of Edge)
"If an ant observes that it is at location l, and there are three edges connected to that location, then the ant goes to the edge with the highest amount of pheromones."
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))

LP6 (Arrival at Edge)
"If an ant goes to edge e from location l to location l1, then later the ant will be at this edge e."
to_be_performed(a, go_to_edge_from_to(e, l, l1)) •→ is_at_edge_from_to(a, e, l, l1)

LP9 (Dropping of Pheromones)
"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."
observes(a, is_at_edge_from_to(e, l, l1)) •→ to_be_performed(a, drop_pheromones_at_edge_from(e, l))

LP12 (Observation of Pheromones)
"If an ant is at a certain location l, then it will observe the number of pheromones present at all edges that are connected to location l."
is_at_location_from(a, l, e0) and connected_to_via(l, l1, e1) and pheromones_at(e1, i) •→ observes(a, pheromones_at(e1, i))

LP13 (Increment of Pheromones)
"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.
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)
"If an ant observes that it is at location F (the food source), then it will pick up some food."
observes(a, is_at_location_from(F, e)) •→ to_be_performed(a, pick_up_food)

To model the example from a single agent perspective, again a number of local dynamic properties are used. Most, but not all of these local properties have a 1:1 correspondence to those for the multi-agent case. For example, the properties for the *single agent case* that correspond to the properties above are as follows (see the next section for more information about this correspondence):

LP5' (Selection of Edge)
"If S observes that it has a paw p at location A, and there are three edges connected to that location, then S will move its paw to the edge of which it believes that it has the highest relevance level."
observes(S, has_paw_at_location_from(p, l, e0)) and neighbours(l, 3) and connected_to_via(l, l1, e1) and belief(S, relevance_level(e1, i1)) and connected_to_via(l, l2, e2) and belief(S, relevance_level(e2, i2)) and e0 ≠ e1 and e0 ≠ e2 and e1 ≠ e2 and i1 > i2 •→ to_be_performed(S, move_paw_to_edge_from_to(p, e1, l1))

LP6' (Paw Arrival at Edge)
"If S moves its paw p to an edge e from a location l to a location l1, then later this paw will be at this edge e."
to_be_performed(S, move_paw_to_edge_from_to(p, e, l, l1)) •→ has_paw_at_edge_from_to(S, p, e, l, l1)

LP11' (Increment of Belief)
"If S has exactly one paw at edge e, then the new number of pheromones at e becomes i*decay+incr."
observes(S, has_paw_at_edge_from_to(p1, e, l, l1)) and ∀l2 not observes(S, has_paw_at_edge_from_to(p2, e, l, l2)) and ∀l3 not observes(S, has_paw_at_edge_from_to(p3, e, l, l3)) and p1 ≠ p2 and p1 ≠ p3 and p2 ≠ p3 and belief(S, relevance_level(e, i)) •→ belief(S, relevance_level(e, i*decay+incr))

LP12' (Collecting of Food)
"If S observes that it has a paw p at location F (the food source), then it will pick up some food with that paw."
observes(S, has_paw_at_location_from(p, F, e)) •→ to_be_performed(S, pick_up_food_with_paw(p))

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, it can generate simulation traces. An example of (part of) such a trace can be seen in Figure 6. 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 was based on the multi-agent simulation model.

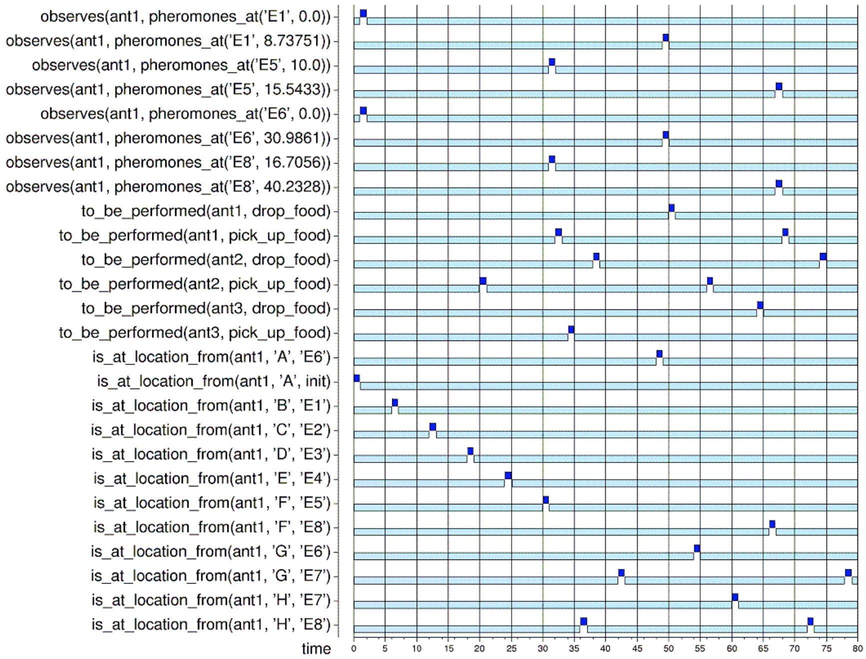

Figure 6 Multi-Agent Simulation Trace

Figure 7 depicts a similar trace as Figure 6, this time based on the single agent simulation model. Note that there are several differences between Figure 6 and 7. In the first place, all ants that are treated as separate agents in Figure 6, are considered as parts of Superant S in Figure 7. For example, is_at_location_from(ant1, A, E6)) in the multi-agent case corresponds to has_paw_at_location_from(S, paw1, A, E6)) in the single agent case. Another important difference is that in the single agent case, there is no explicit observation of pheromones. The reason for this is that the belief(S, relevance_level(e, i)) states (which are the single agent equivalent for the pheromones_at(e, i) states in the multi-agent case) are internal states of S, which do not have to be observed.

Altogether, the software environment has been used to successfully generate a large number of simulation traces on the basis of both simulation models. Because of space restrictions, not all resulting traces are shown completely here. However, all traces, as well as the complete sets of dynamic properties, are shown on the following URL: http://www.cs.vu.nl/~tbosse/isomorphism/.

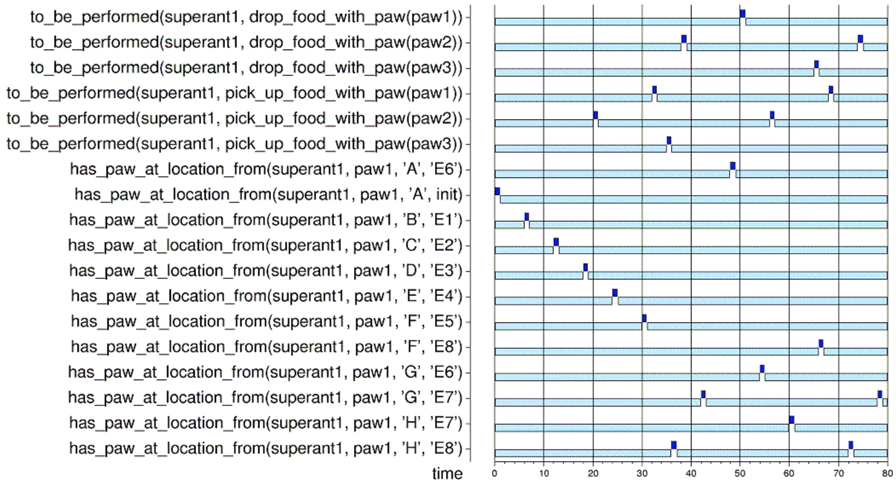

Figure 7 Single Agent Simulation Trace

6 The Extended Interpretation Mapping

In Section 3 it was shown how the basic interpretation mapping can be defined as a mapping between state properties. It was suggested that this mapping can be extended to a mapping between dynamic properties in *leads to* format. Therefore, the following interpretation mapping can be defined:

$$\varphi(\alpha \bullet\!\!\rightarrow \beta) \;=\; \varphi(\alpha) \bullet\!\!\rightarrow \varphi(\beta)$$

Using this interpretation mapping, combined with the basic mapping of the state ontology elements described in Section 4, mappings between the dynamic properties of the case study can be found, e.g.:

φ(LP6')
= φ(to_be_performed(S, move_paw_to_edge_from_to(p, e, l, l1)) •→ has_paw_at_edge_from_to(S, p, e, l, l1))
= φ(to_be_performed(S, move_paw_to_edge_from_to(p, e, l, l1))) •→
φ(has_paw_at_edge_from_to(S, p, e, l, l1))
= to_be_performed(a, go_to_edge_from_to(e, l, l1)) •→ is_ at_edge_from_to(a, e, l, l1))
= LP6

A mapping between all dynamic properties (in *leads to* format) of the case study is given in Table 4. Notice that in some cases a certain dynamic property is mapped to a dynamic property that is not literally in the multi-agent model, but actually is a combination of two other local properties present in the model. This shows where the single agent conceptualisation is simpler than the multi-agent conceptualisation.

Single Agent Conceptualisation	Multi-Agent Conceptualisation
LP1’	LP1
LP2’	LP2
LP3’	LP3
LP4’	LP4
LP5’	LP5 & LP12
LP6’	LP6
LP7’	LP7
LP8’	LP8
LP9’	LP10
LP10’	LP11
LP11’	LP9 & LP13
LP12’	LP14
LP13’	LP15
LP14’	LP16
LP15’	LP17
LP16’	LP9 & LP18

Table 4 Mapping between dynamic properties in *leads to*

In addition, it is possible to extend the mapping to the wider class of TTL expressions. Recall that TTL expressions are built on atoms of the form state(γ, t) |= p. By the basic mapping the state property p can be translated into φ(p), which is assumed to be part of the ontology of one of the agents Ai in the multi-agent conceptualisation. Moreover, the trace name γ can be mapped onto a trace name φ(γ) = γ’. Then the extended interpretation mapping for state(γ, t) |= p is defined by:

φ: state(γ, t) |= p = state(γ’, t) |= φ(p)

After these atoms have been mapped, TTL expressions as a whole can be mapped in a straightforward compositional manner:

φ(A&B) = φ(A) & φ(B)
φ(A ⇒ B) = φ(A) ⇒ φ(B)
φ(not A) = not φ(A)
φ(∀v A(v)) = ∀v’ φ(A(v’))
φ(∃v A(v)) = ∃v’ φ(A(v’))

For example, take the following TTL expression, which is a global property for the single agent case of the ant example:

GP1’ Food Discovery
“Eventually, one of the paws of S will be at the food location.”
∃t,p,l,e [state(γ, t) |= has_paw_at_location_from(S, p, l, e) & state(γ, t) |= food_location(l)]

This expression is mapped as follows:

φ(∃t,p,l,e [state(γ, t) |= has_paw_at_location_from(S, p, l, e) & state(γ, t) |= food_location(l)])
= ∃t’,p’,l’,e’ φ([state(γ, t’) |= has_paw_at_location_from(S, p’, l’, e’) & state(γ, t’) |= food_location(l’)])
= ∃t’,p’,l’,e’ [φ(state(γ, t’) |= has_paw_at_location_from(S, p’, l’, e’)) & φ(state(γ, t’) |= food_location(l’))]
= ∃t’,p’,l’,e’ [state(γ’, t’) |= φ(has_paw_at_location_from(S, p’, l’, e’)) & state(γ’, t’) |= φ(food_location(l’))]
= ∃t’,p’,l’,e’ [state(γ’, t’) |= is_at_location_from(p’, l’, e’) & state(γ’, t’) |= food_location(l’)]

7 Discussion

This paper addresses the question to what extent a process involving multiple agents that shows some form of collective intelligence can be interpreted as single agent behaviour. The question is answered by formal analysis. It is shown for an example process how it can be conceptualised and formalised in two different manners: from a single agent (or cognitive) and from a multi-agent (or social) perspective. Moreover, it is shown how a basic ontological mapping can be formally defined between the two formalisations, and how this mapping can be extended to a mapping of dynamic properties. Thus it is shown how the collective behaviour can be interpreted in a formal manner as single agent behaviour. For example, the fact that food is taken from the source to the nest can be explained by a sequence of actions of one agent, based on its beliefs.

Having such a mapping allows one to explain collective or social behaviour in terms of single agent concepts in the following manner. Behaviour often is explained by considering the basic underlying causal relations or mechanisms. The mapping and its formalisation allows to replace an explanation of behaviour in terms of basic mechanisms involving frequent interactions of the multiple agents (with each other and/or with the external world), by an explanation that leaves out these interactions and bases itself directly on mental states of the single agent conceptualisation. This explanation is simpler, more abstract and perhaps more elegant, than the more complicated explanation based on the interactions. This is made possible by introducing a new ontology for states involved. For example, considering part of the external world as extended mind allows one to give another interpretation to external physical processes and states. Physical state properties such as ‘pheromone is present at d’ are reconceptualised as, for example, ‘it is believed that d is a relevant path’. Why would one introduce extra language to refer to the same fact in the world? Given the literature on reduction, where often it is claimed that mental state properties can be and actually should be replaced by their physical realisers, at first sight such an opposite move may seem a bit surprising. For example, Kim [9] (pp. 214-216) claims that ontological simplification is one of the reasons to reduce mental state properties to physical state properties. In the extended mind case at hand the converse takes place; a question is what is the advantage of this ontological complication. A number of arguments in support of this can be given. By Clark and Chalmers [3], it is claimed that this allows application of other types of explanation and other methods of scientific investigation:

> (…) we allow a more natural explanation of all sorts of actions. (…) in seeing cognition as extended one is not merely making a terminological decision; it makes a significant difference to the methodology of scientific investigation. In effect, explanatory methods that might once have been thought appropriate only for the analysis of "inner" processes are now being adapted for the study of the outer, and there is promise that our understanding of cognition will become richer for it. [3], Section 3.

In [8] it is explained in some detail why in various cases in other areas (such as Computer Science) such an antireductionist strategy often pays off; some of the discussed advantages in terms of insight, transparency and genericity are: additional higher-level ontologies can improve understanding as they may allow simplification of the picture by abstracting from lower-level details; more insight is gained from a conceptually higher-level perspective; analysis of more complex processes is possible; finally, the same concepts have a wider scope of application, thus obtaining unification.

Future research will further analyse the interpretation mapping in the context of logic: the notion of an interpretation of one (formal) logical theory T in another logical theory T' has a formal definition in logic. It is an interesting question whether it can be proven logically that the conditions of this definition are fulfilled for the mapping defined in this paper. For example, a question is whether it can be proven that:

$$T \mid\!\!-\!\!-\ \alpha \Rightarrow T' \mid\!\!-\!\!-\ \varphi(\alpha)$$

for all formulae α, where T is a logical theory of single agent behaviour and T' a theory of multi-agent behaviour.

Acknowledgements

The authors are grateful to Catholijn Jonker for many valuable discussions on the topic of the paper, to Lourens van der Mey for his contribution to the development of the software environment, and to Martijn Schut for his contribution to parts of the simulations.

References

1. Bonabeau, E., Dorigo, M., and Theraulaz, G. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York.
2. Clark, A. *Being There: Putting Brain, Body and World Together Again*. MIT Press, 1997.
3. Clark, A. and Chalmers, D. The Extended Mind. In: *Analysis*, vol. 58, 1998, pp. 7-19.
4. Deneubourg, J.L., Aron, S., Goss, S., Pasteels, J.M., and Duerinck, G. Random Behavior, Amplification Processes and Number of Participants: How They Contribute to the Foraging Properties of Ants. In: *Evolution, Games and Learning: Models for Adaptation in Machines and Nature*, North Holland, Amsterdam, 1986, pp. 176-186.
5. Dennett, D.C. *Kinds of Mind: Towards an Understanding of Consciousness*, New York: Basic Books.
6. Drogoul, A., Corbara, B., and Fresneau, D.: MANTA: New experimental results on the emergence of (artificial) ant societies. In: Gilbert N. and Conte R. (eds.), *Artificial Societies: the computer simulation of social life*, UCL Press, 1995.
7. 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(3), 2003, pp. 191-210.
8. Jonker, C.M., Treur, J., and Wijngaards, W.C.A., Reductionist and Antireductionist Perspectives on Dynamics. *Philosophical Psychology Journal*, vol. 15, 2002, pp. 381-409.
9. Kim, J. *Philosophy of Mind*. Westview Press
10. Kirsh, D. and Maglio, P. On distinguishing epistemic from pragmatic action. *Cognitive Science*, vol. 18, 1994, pp. 513-49.
11. Menary, R. (ed.) (2005). *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, 2005, to appear.