

Making and Breaking Engagements: An Operational Analysis of Agent Relationships

Mark d’Inverno¹ and Michael Luck²

- ¹ School of Computer Science, University of Westminster, London, W1M 8JS, UK.
Email: dinverm@westminster.ac.uk
- ² Department of Computer Science, University of Warwick, Coventry, CV4 7AL, UK.
Email: mikeluck@dcs.warwick.ac.uk

Abstract. Fundamental to the operation of multi-agent systems is the concept of cooperation between individual agents by which the overall system exhibits significantly greater functionality than the individual components. Not only is it important to understand the structure of such relationships in multi-agent systems, it is also important to understand the ways in which cooperative relationships come about. This is particularly true if such an analysis is to be relevant to, and useful for, the construction of real systems for which the invocation and destruction of such relationships is critical. This paper extends previous work concerned with formalising the structure of inter-agent relationships to provide an operational analysis of the invocation and destruction of engagement and cooperation within our formal agent framework.

1 Introduction

Fundamental to the operation of multi-agent systems is the concept of cooperation between individual agents. If single agent systems can cooperate, they can exploit the capabilities and functionality of others to achieve their own individual goals. This moves beyond the advantages of robustness in traditional distributed systems in the face of individual component failure since components can be replaced and cooperation configurations realigned. It allows the specific expertise and competence of different agents to complement each other so that in addition to general resilience, the overall system exhibits significantly greater functionality than individual components.

Durfee, for example, defines a multi-agent system as a collection of problem solvers that “work together” to achieve goals that are beyond the scope of their individual abilities [2]. This notion of agents helping each other, working together or cooperating in some way is common. Huberman and Clearwater similarly characterise multi-agent systems by the interaction of many agents trying to solve problems in a cooperative fashion [3], and Lesser describes a multi-agent system as a computational system in which several semi-autonomous agents interact or work together to perform some tasks or satisfy some goals [4].

Since cooperation underlies the structure of multi-agent systems, it is important to be able to model and analyse it in detail in the context of a well-founded framework. Moreover, the ways in which cooperative relationships come about

are also important for a complete understanding of the nature of cooperation. This is especially true if such an analysis is to be relevant to real systems for which the invocation and destruction of such relationships is critical.

Previous work has focussed on taxonomising the types of interactions that occur between agents in a multi-agent system, distinguishing in particular between *engagements* of non-autonomous agents and *cooperation* between autonomous agents [7]. This work was based on our agent hierarchy which defined *agency* and *autonomy*, and specified how autonomy is distinct but is achieved by *motivating* agency and generating goals [5]. In this paper we aim to extend this work on cooperation and engagement by describing the operations necessary for their invocation and destruction, and providing a formal specification in the context of the existing framework. First, we briefly review the agent framework and the nature of engagement and cooperation within it. We then address the invocation and destruction of such structures, enumerating the variety of situations in which they arise. Finally, we discuss what benefits such an analysis provides.

2 Background

To present a mathematical description of our definitions, we use the specification language, Z [8]. Space constraints prohibit a detailed explanation of the notation here, but our use should be clear from the combination of text and formalism. (A brief introduction to Z can be found in the appendix of [1].) In this section, we sketch previous work on the agent framework. Details may be found in [5, 7].

An *object* comprises a set of *motivations*, a set of *goals*, a set of *actions*, and a set of *attributes* such that the attributes and actions are non-empty.

<p><i>Object</i></p> <p><i>attributes</i> : \mathbb{P} <i>Attribute</i></p> <p><i>capableof</i> : \mathbb{P} <i>Action</i></p> <p><i>goals</i> : \mathbb{P} <i>Goal</i></p> <p><i>motivations</i> : \mathbb{P} <i>Motivation</i></p> <hr/> <p>$attributes \neq \{ \} \wedge capableof \neq \{ \}$</p>
--

An object, described by its features and capabilities, is just an automaton and cannot engage other entities to perform tasks, nor is it itself used in this way. However, it serves as a useful abstraction mechanism by which it is regarded as distinct from the remainder of the environment, and can subsequently be transformed into an agent, an augmented object, engaged to perform some task or satisfy some goal. Viewing something as an agent means that we regard it as satisfying or pursuing some goal. Furthermore, it means that it must be doing so either for another agent, or for itself, in which case it is autonomous. If it is for itself, it must have generated the goal through its motivations. Autonomous agents are thus just agents that can generate their own goals from motivations.

<i>Agent</i>
<i>Object</i>
$goals \neq \{ \}$

<i>AutonomousAgent</i>
<i>Agent</i>
$motivations \neq \{ \}$

For ease of exposition, we further refine the agent hierarchy with two new definitions. A *neutral-object* is an object that is *not* an agent, and a *server-agent* is an agent that is *not* an autonomous agent.

<i>NeutralObject</i>
<i>Object</i>
$goals = \{ \} \wedge motivations = \{ \}$

<i>ServerAgent</i>
<i>Agent</i>
$motivations = \{ \}$

Now, we consider a multi-agent system as a collection of objects, agents and autonomous as defined below. In this view, a multi-agent system contains autonomous agents which are all agents, and in turn, all agents are objects. An agent is either a server-agent or an autonomous agent, and an object is either a neutral-object or an agent.

<i>MultiAgentSystemComponents</i>
$objects : \mathbb{P} Object$
$agents : \mathbb{P} Agent$
$autonomousagents : \mathbb{P} AutonomousAgent$
$neutralobjects : \mathbb{P} NeutralObject$
$serveragents : \mathbb{P} ServerAgent$
$autonomousagents \subseteq agents \subseteq objects$
$agents = autonomousagents \cup serveragents$
$objects = neutralobjects \cup agents$

3 Engagement and Cooperation

In this section, we briefly review previous work that categorised inter-agent relationships into two distinct classes, engagements and cooperations. We provide formal descriptions and brief explanations. For a full exposition on the nature of these relationships, see [7]. Note, however, that there have been several minor modifications to the earlier formal specification.

3.1 Engagement

A direct engagement occurs when a neutral-object or a server-agent adopts some goals. In a direct engagement, a *client*-agent with some goals uses another *server*-agent to assist them in the achievement of those goals. A server-agent either exists already as a result of another engagement, or is instantiated from a neutral-object for the current engagement. No restriction is placed on a client-agent.

We define a *direct engagement* below to consist of a client agent, *client*, a server agent, *server*, and the goal that *server* is satisfying for *client*. An agent cannot engage itself, and both agents must have the goal of the engagement.

<i>DirectEngagement</i> <i>client</i> : Agent <i>server</i> : ServerAgent <i>goal</i> : Goal
<i>client</i> ≠ <i>server</i> <i>goal</i> ∈ (<i>client</i> .goals ∩ <i>server</i> .goals)

The set of all *direct* engagements in a system is given by *direengagements* in the following schema. For any direct engagement in *direengagements*, there can be no intermediate *direct* engagements of the goal, so there is no other agent, *y*, where *client* engages *y* for *goal*, and *y* engages *server* for *goal*.

<i>SystemEngagements</i> MultiAgentSystemComponents <i>direengagements</i> : P <i>DirectEngagement</i>
$\forall eng : direengagements \bullet \neg (\exists y : Agent; e_1, e_2 : direengagements \mid$ $e_1.goal = e_2.goal = eng.goal \bullet e_1.server = eng.server \wedge$ $e_2.client = eng.client \wedge e_1.client = e_2.server = y)$

An *engagement chain* represents a sequence of *direct engagements*. Specifically, an *engagement chain* comprises a goal, *goal*, the autonomous client that generated the goal, *autoagent*, and a sequence of server-agents, *agentchain*, where each agent in the sequence directly engages the next. For any engagement chain, there must be at least one server-agent, all agents must share *goal*, and each agent can only be involved once. seq₁ represents a non-empty sequence.

<i>EngagementChain</i> <i>goal</i> : Goal <i>autoagent</i> : AutonomousAgent <i>agentchain</i> : seq ₁ Agent
<i>goal</i> ∈ <i>autoagent</i> .goals <i>goal</i> ∈ $\bigcup \{s : Agent \mid s \in \text{ran } agentchain \bullet s.ggoals\}$ $\#(\text{ran } agentchain) = \#agentchain$

The set of all engagement chains in a system is given in the schema below by *engchains*. For every engagement chain, *ec*, there must be a direct engagement between the autonomous agent, *ec.autoagent*, and the first client of *ec*, *head ec.agentchain*, with respect to the goal of *ec*, *ec.goal*. There must also be a direct engagement between any two agents which follow each other in *ec.agentchain* with respect to *ec.goal*. In general, an agent *engages* another agent if there is some engagement chain in which it precedes the server agent.

<i>SystemEngagementChains</i>
<i>SystemEngagements</i> <i>engchains</i> : \mathbb{P} <i>EngagementChain</i>
$\forall ec : engchains; s_1, s_2 : Agent \bullet$ $(\exists d : direngagements \bullet d.goal = ec.goal \wedge d.client = ec.autoagent$ $\wedge d.server = head\ ec.agentchain) \wedge$ $(s_1, s_2) \text{ in } ec.agentchain \Rightarrow (\exists d : direngagements \bullet$ $d.client = s_1 \wedge d.server = s_2 \wedge d.goal = ec.goal)$

3.2 Cooperation

Two autonomous agents are said to be *cooperating* with respect to some goal if one of the agents has adopted goals of the other. This notion of autonomous goal acquisition applies both to the *origination* of goals by an autonomous agent for its own purposes, and the *adoption* of goals from others, since in each case the goal must have a positive motivational effect [6]. For autonomous agents, the goal of another can only be adopted if it has such an effect, and this is also exactly why and how goals are originated. Thus goal adoption and origination are related forms of goal generation. The term *cooperation* can be used only when those involved are autonomous and, potentially, capable of resisting. If they are not autonomous, nor capable of resisting, then one simply *engages* the other.

A *cooperation* describes a goal, the autonomous agent that generated the goal, and those autonomous agents who have adopted that goal from the generating agent. In addition, all the agents involved have the goal of the cooperation, an agent cannot cooperate with itself, and the set of cooperating agents must be non-empty. Cooperation cannot, therefore, occur unwittingly between agents, but must arise as a result of the motivations of an agent and that agent recognising the goal in another.

<i>Cooperation</i>
<i>goal</i> : <i>Goal</i> <i>generatingagent</i> : <i>AutonomousAgent</i> <i>cooperatingagents</i> : \mathbb{P} <i>AutonomousAgent</i>
$goal \in generatingagent.goals$ $\forall aa : cooperatingagents \bullet goal \in aa.goals$ $generatingagent \notin cooperatingagents$ $cooperatingagents \neq \{ \}$

The set of cooperations in a multi-agent system is given by the *cooperations* variable in the schema, *SystemCooperations*. The predicate part of the schema states that for any cooperation, the union of the cooperating agents and the generating agent is a *subset* of the set of all autonomous agents which have that goal. As a consequence, two agents sharing a goal are not necessarily cooperating. In addition, the set of all cooperating agents is a subset of all autonomous agents of the system since not all are necessarily participating in cooperations.

<p style="text-align: center;"><i>SystemCooperations</i></p> <hr/> <p><i>MultiAgentSystemComponents</i> <i>cooperations</i> : P <i>Cooperation</i></p> <hr/> <p>$\forall c : \text{cooperations} \bullet (c.\text{cooperatingagents} \cup \{c.\text{generatingagent}\}) \subseteq$ $\{a : \text{autonomousagents} \mid c.\text{goal} \in a.\text{goals} \bullet a\}$</p> <p>$\cup \{c : \text{cooperations} \bullet c.\text{cooperatingagents}\} \subseteq \text{autonomousagents}$</p>
--

4 Operational Aspects of Engagement and Cooperation

So far we have provided an analysis of the *structure* of multi-agent systems based on inter-agent relationships. In this section, we provide an operational analysis of how these cooperations and direct engagements are created and destroyed, and how this affects the configuration of inter-agent relationships.

There are four principal operations to be considered, as follows:

- a server-agent adopting the goals of an agent giving rise to a new engagement;
- a server-agent being released from some or all of its agency obligations by an engaging agent, thus destroying a direct engagement;
- an autonomous agent adopting the goal of another, so that either a new cooperation is formed, or the agent joins an existing cooperation;
- an autonomous agent destroying the goals of an existing cooperation, resulting in either the destruction of the cooperation, or the removal of the autonomous agent from the cooperation.

To provide an operational account of these relationships we must specify how they are affected when new cooperations and engagements are invoked and destroyed. Before considering the operations in detail we first specify some general functions to create relationships from individual components. Thus, *MakeEng*, *MakeEngChain* and *MakeCoop* below simply construct the schema types, *Engagement*, *EngagementChain* and *Cooperation* respectively. The functions, *MakeEng* and *MakeCoop*, are partial; *MakeEng* is not defined if the two agents involved in the cooperation are the same, and *MakeCoop* is not defined if the single autonomous agent initiating the cooperation is in the set of other autonomous agents. Note that $\mathbb{P}_1 \text{AutonomousAgent}$, represents non-empty sets

of elements of type *AutonomousAgents*. The schema also makes use of the *mu-expression*. In general, the value of a in the mu-expression, $\mu a : A \mid p$, is the unique value of type A which satisfies predicate p .³

$$\begin{array}{l}
\hline
\textit{MakeEng} : (\textit{Goal} \times \textit{Agent} \times \textit{ServerAgent}) \rightsquigarrow \textit{DirectEngagement} \\
\textit{MakeEngChain} : (\textit{Goal} \times \textit{AutonomousAgent} \times \textit{seq}_1 \textit{ServerAgent}) \\
\hspace{15em} \rightsquigarrow \textit{EngagementChain} \\
\textit{MakeCoop} : (\textit{Goal} \times \textit{AutonomousAgent} \times \mathbb{P}_1 \textit{AutonomousAgent}) \\
\hspace{15em} \rightarrow \textit{Cooperation} \\
\hline
\forall g : \textit{Goal}; a : \textit{Agent}; aa : \textit{AutonomousAgent}; s : \textit{ServerAgent}; \\
\hspace{2em} ch : \textit{seq}_1 \textit{ServerAgent}; aas : \mathbb{P}_1 \textit{AutonomousAgent} \bullet \\
\textit{MakeEng}(g, a, s) = (\mu d : \textit{DirectEngagement} \mid \\
\hspace{2em} d.\textit{goal} = g \wedge d.\textit{client} = a \wedge d.\textit{server} = s) \wedge \\
\textit{MakeEngChain}(g, a, ch) = (\mu ec : \textit{EngagementChain} \mid \\
\hspace{2em} ec.\textit{goal} = g \wedge ec.\textit{autoagent} = a \wedge ec.\textit{agentchain} = ch) \wedge \\
\textit{MakeCoop}(g, aa, aas) = (\mu c : \textit{Cooperation} \mid \\
\hspace{2em} c.\textit{goal} = g \wedge c.\textit{generatingagent} = aa \wedge c.\textit{cooperatingagents} = aas)
\end{array}$$

Next, we define the generic function, *cut*, which takes an injective sequence (where no element appears more than once) and an element, and removes all the elements of the sequence which appear after this element. If the element does not exist in the sequence, the sequence is left unaltered.

$$\begin{array}{l}
\hline
[X] \\
\hline
\textit{cut} : (\textit{iseq} X \times X) \rightsquigarrow \textit{iseq} X \\
\hline
\forall x : X; s, t : \textit{seq} X \bullet \textit{cut}(s, x) = t \Leftrightarrow \\
\hspace{2em} \textit{last} t = x \wedge (\exists u : \textit{seq} X \bullet s = t \frown u)
\end{array}$$

We also define a total function that creates a new object by ascribing a set of goals to an existing object. It is valid for any object and any set of goals. This is detailed elsewhere with a complete formalisation of goal generation [6].

$$\begin{array}{l}
\hline
\textit{ObjectAdoptGoals} : (\textit{Object} \times \mathbb{P} \textit{Goal}) \rightarrow \textit{Object} \\
\hline
\forall gs : \mathbb{P} \textit{Goal}; \textit{old}, \textit{new} : \textit{Object} \bullet \\
\textit{ObjectAdoptGoals}(\textit{old}, gs) = \textit{new} \Leftrightarrow \textit{new}.\textit{goals} = \textit{old}.\textit{goals} \cup gs \\
\hspace{2em} \wedge \textit{new}.\textit{capableof} = \textit{old}.\textit{capableof} \wedge \textit{new}.\textit{attributes} = \textit{old}.\textit{attributes}
\end{array}$$

This allows us to define two further functions, *ExtendChain* and *CutChain*. The first takes an engagement chain and an object, and extends the engagement chain to include the object. The second function cuts an engagement chain after the occurrence of an object.

³ For example, the expression $\mu n : \mathbb{N} \mid (n * n) \bmod 2 = 0 \wedge 0 < (n * n) \leq 10l$ binds the variable n to the value 10.

$ExtendChain : EngagementChain \times Object \rightarrow EngagementChain$ $CutChain : EngagementChain \times Agent \rightarrow EngagementChain$
$\forall c : EngagementChain; e : Object \bullet$ $ExtendChain(c, e) = (\mu new : EngagementChain $ $new.goal = c.goal \wedge new.autoagent = c.autoagent$ $\wedge new.agentchain = c.agentchain \hat{\ } \langle ObjectAdoptGoals(e, \{c.goal\}) \rangle) \wedge$ $CutChain(c, e) = (\mu new : EngagementChain $ $new.goal = c.goal \wedge new.autoagent = c.autoagent$ $\wedge new.agentchain = cut(c.agentchain, e))$

Formally, the structure of multi-agent systems is defined below. It states that a multi-agent system must contain at least one relationship between two agents.

$MultiAgentSystemStructure$
$MultiAgentSystemComponents$ $SystemEngagements$ $SystemEngagementChains$ $SystemCooperations$
$\#cooperations + \#direngagements \geq 1$

4.1 Making Engagements

When a new direct engagement is formed between an agent and a server agent, the associated engagement chain may be altered in several ways. The different possibilities depend on whether the engaging agent is at the tail, the head, or in the middle of the chain. Consider an engagement chain where A is an autonomous agent at the head of the chain directly engaging the server-agent, $S1$, which is directly engaging server-agent, $S2$, in turn directly engaging $S3$, as in Figure 1(a).

- If the last agent in the chain, $S3$, engages a neutral-object, O , the chain is extended to include the engagement between $S3$ and O , as in Figure 1(b).
- If the autonomous agent, A , directly engages O , a new engagement chain is created solely comprising A and O , as in Figure 1(c).
- If any *server*-agent, other than that at the tail of the engagement chain, engages O , then a new engagement chain is formed between them. Thus if $S1$ engages O , the existing chain is unchanged, but a new chain is formed from the engagements up to and including $S1$ in the original chain, with the addition of the new engagement of O by $S1$, as in Figure 1(d).

The aspects of forming a new engagement common to all three scenarios are described in the next schema. Here, the engaging agent, $agent?$, the engaged object, $e?$, the goal of the engagement, $goal?$, and an optional engagement chain, $chain?$, are given as input to the operation, and the structure of the multi-agent system changes. The predicate part states that the object is a system object, and the agent is a known agent with the goal, $goal?$. If no engagement

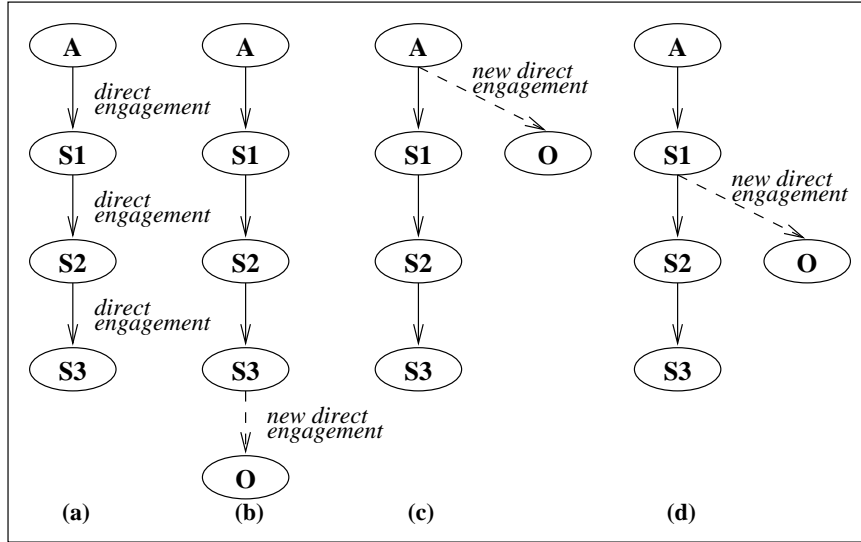


Fig. 1. Alternative ways to make engagements

chain already exists, so that *chain?* is not defined, then *agent?* must be autonomous. Conversely, if *chain?* is defined, *agent?* must be a server-agent, the goal of *chain?* must be *goal?*, and *agent?* must be part of *chain?*. (Formal definitions for optional, defined, undefined and the are given in the appendix.) There is no change to the set of cooperations, but the set of direct engagements is updated to include the new engagement between the agent and the object.

GeneralEngage

agent? : Agent
e? : Object
goal? : Goal
chain? : optional [EngagementChain]
 Δ MultiAgentSystemStructure

$e? \in \text{objects}$
 $agent? \in \text{agents}$
 $goal? \in agent?.goals$
undefined *chain?* $\Leftrightarrow agent? \in \text{autonomousagents}$
defined *chain?* \Rightarrow
 $(agent? \in \text{serveragents} \wedge$
 $(\text{the } chain?).goal = goal? \wedge$
 $agent? \in \text{ran}(\text{the } chain?).agentchain)$

cooperations' = *cooperations*
direngagements' = *direngagements* \cup
 $\{ \text{MakeEng}(goal?, agent?, (\text{ObjectAdoptGoals}(e?, \{goal?\}))) \}$

The distinct aspects of the ways in which the set of engagement chains are affected in each scenario are detailed below. First, the engaging agent is a server-agent at the end of the chain so that the chain is extended to include this new direct engagement.

$\begin{array}{l} \textit{EngageExtendChain} \\ \textit{GeneralEngage} \\ \hline \text{defined } chain? \wedge e? = \text{last}(\text{the } chain?).agentchain \Rightarrow \\ \text{engchains}' = \text{engchains} \setminus chain? \cup \\ \{ \textit{ExtendChain}(\text{the } chain?, e?) \} \end{array}$

Second, the engaging agent is autonomous, and a new engagement chain is formed from $goal?$, $agent?$, and the sequence consisting of $agent?$ and the newly instantiated agent.

$\begin{array}{l} \textit{StartNewChain} \\ \textit{GeneralEngage} \\ \hline agent? \in \textit{autonomousagents} \Rightarrow \\ \text{engchains}' = \text{engchains} \cup \\ \{ \textit{MakeEngChain}(goal?, agent?, \langle agent?, \textit{ObjectAdoptGoals}(e?, \{goal?\}) \rangle) \} \end{array}$

Third, if $agent?$ is not at head or tail of the chain, then the original chain is unchanged, and a new chain is formed from the direct engagements in the original chain up to the agent, plus the new direct engagement between $agent?$ and the newly instantiated object.

$\begin{array}{l} \textit{CutandAddChain} \\ \textit{GeneralEngage} \\ \hline (\text{defined } chain?) \wedge (e? \neq \text{last}(\text{the } chain?).agentchain) \\ \wedge e? \neq (\text{head}(\text{the } chain?).agentchain) \Rightarrow \\ \text{engchains}' = \text{engchains} \cup \\ \{ \textit{ExtendChain}(\textit{CutChain}(\text{the } chain?, agent?), e?) \} \end{array}$
--

The operation of making an engagement can then be defined using schema disjunction. Thus, the *Engage* operation is applied when any of the following three occur: *CutandAddChain*, *EngageExtendChain* or *StartNewChain*.

$$\textit{Engage} \hat{=} \textit{EngageExtendChain} \vee \textit{StartNewChain} \vee \textit{CutandAddChain}$$

4.2 Breaking Engagements

If an autonomous agent or a server-agent in an engagement chain *disengages* another server-agent, either through destroying the goal itself or because the agent is no longer required to achieve it, all subsequent engagements in the chain

agents are autonomous. The sets of direct engagements and engagement chains remain unchanged.

$\begin{array}{l} \text{GeneralCooperate} \\ \text{goal?} : \text{Goal} \\ \text{genagent?}, \text{coopagent?} : \text{AutonomousAgent} \\ \Delta \text{MultiAgentSystemStructure} \end{array}$
$\begin{array}{l} \text{goal?} \in \text{genagent?.goals} \\ \text{goal?} \notin \text{coopagent?.goals} \\ \{\text{genagent?}, \text{coopagent?}\} \subseteq \text{autonomousagents} \\ \text{direngagements}' = \text{direngagements} \\ \text{engchains}' = \text{engchains} \end{array}$

Then a new cooperation is formed when there is no existing cooperation for $goal?$ with $genagent?$ as the generating agent. This is described formally below.

$\begin{array}{l} \text{NewCooperation} \\ \text{GeneralCooperate} \end{array}$
$\begin{array}{l} \neg (\exists c : \text{cooperations} \bullet c.\text{goal} = \text{goal?} \wedge c.\text{generatingagent} = \text{genagent?}) \wedge \\ \text{cooperations}' = \text{cooperations} \cup \\ \{ \text{MakeCoop}(\text{goal?}, \text{genagent?}, \{ \text{coopagent?} \}) \} \end{array}$

If such a cooperation does exist, then $coopagent?$ is added to it. Specifically, it adopts the goal, $goal?$, and joins the cooperation, as specified in the *JoinCooperation* schema.

$\begin{array}{l} \text{JoinCooperation} \\ \text{GeneralCooperate} \end{array}$
$\begin{array}{l} \exists c : \text{cooperations} \bullet c.\text{goal} = \text{goal?} \wedge c.\text{generatingagent} = \text{genagent?} \wedge \\ c.\text{goal} = \text{goal?} \wedge c.\text{generatingagent} = \text{genagent?} \wedge \\ \text{cooperations}' = \text{cooperations} \setminus \{c\} \cup \\ \{ \text{MakeCoop}(\text{goal?}, \text{genagent?}, c.\text{cooperatingagents} \cup \\ \{ \text{ObjectAdoptGoals}(\text{coopagent?}, \{ \text{goal?} \}) \}) \} \} \end{array}$

4.4 Breaking or Leaving a Cooperation

There are three cases for autonomous agents destroying the goal of a cooperation in which they are involved, illustrated in Figure 2. First, the generating agent, G , destroys the goal of a cooperation with the result that the cooperation is itself destroyed. This does not imply that $C1$ and $C2$ have destroyed the goal. Second, the cooperation is also destroyed when the only cooperating agent destroys the cooperation goal. Finally, when there are many cooperating agents, one of which destroys the cooperation goal, the cooperation is not destroyed but modified so that only one agent *leaves* the cooperation.


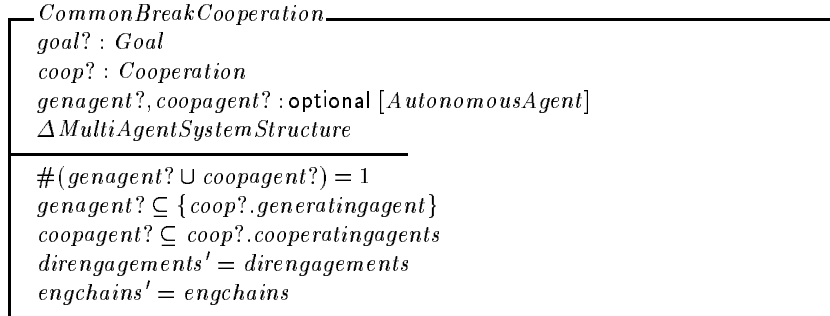



Fig. 2. Breaking a cooperation

In all three cases, however, the state is changed and the set of engagements is unaltered as defined by *CommonBreakCooperation*. A goal, *goal?*, a cooperation, *coop?*, and, optionally, two autonomous agents, *genagent?*, and *coopagent?* are inputs. The preconditions state that either *genagent?* or *coopagent?* is input. In addition, the schema checks that *genagent?* is the generating agent and that *coopagent?* is a cooperating agent of *coop?*. The sets of direct engagements and engagement chains are unchanged.



Each of the three different scenarios can now be specified formally as refinements to this general operation schema. First, the generating agent of the cooperation destroys the goal of the cooperation. The cooperation, *coop?*, is destroyed and removed from *cooperations*.



Second, the only cooperating agent destroys the goal of the cooperation. In this case, the cooperation is similarly destroyed and removed from *cooperations*.

<i>CooperatingAgentDestroysCooperation</i>
<i>CommonBreakCooperation</i>
(defined $coopagent? \wedge coopagent? \subset coop?.cooperatingagents$) $\Rightarrow cooperations' = cooperations \setminus \{coop?\}$

Finally, a cooperating agent which is not the only cooperating agent destroys the goal of the cooperation. It is removed from the cooperation and the remaining cooperation is added to *cooperations*.

<i>CooperatingAgentLeavesCooperation</i>
<i>CommonBreakCooperation</i>
(defined $coopagent? \wedge coopagent? \subset coop?.cooperatingagents$) $\Rightarrow cooperations' = cooperations \setminus \{coop?\} \cup$ $\{MakeCoop(goal?, coop?.generatingagent,$ $(coop?.cooperatingagents \setminus coopagent?))\}$

Schema disjunction is then used to define *BreakCooperation*.

$$BreakCooperation \hat{=} GeneratingAgentDestroysCooperation \vee CooperatingAgentDestroysCooperation \vee CooperatingAgentLeavesCooperation$$

5 Discussion

Identifying the structures and relationships between agents in multi-agent systems provides a way of understanding the nature of the system, its purpose and functionality. This is typical of existing analyses. However, if we are to build systems based on a recognition of these relationships, so that prior relationships are not destroyed when new ones are created, we must extend such analyses into *operational* areas. This paper has provided just such an analysis, explicating the different kinds of structures that can arise, and showing how different configurations of agents and relationships can evolve by invoking and and destroying them. As part of the analysis, we have had to consider a range of scenarios and the effects of changing relationships upon them.

This is particularly important for system developers aiming to build programs that are capable of exploiting a dynamic multi-agent environment. An operational analysis is vital in providing a link from the structural account to methods for accessing and manipulating these structures. If such operational analyses are avoided, then the merit of research that aims to lay a foundation for the practical aspects of multi-agent systems is limited. Indeed, most existing research has concentrated on addressing either theoretical or practical problems, and has not managed to cross the boundary between the two. Our work strives to place a foot firmly in both camps. We have constructed a formal specification of the structures necessary for an understanding of multi-agent systems and the relationships therein, and we have also shown how it is possible to set about using this understanding in the development of practical systems. Certainly, more

work is required for a complete move from theory to practice, but the current work takes us a large part of the way down that road.

References

1. M. d’Inverno and M. Luck. A formal view of social dependence networks. In *Proceedings of the First Australian DAI Workshop, Lecture Notes in Artificial Intelligence, 1087*, pages 115–129. Springer Verlag, 1996.
2. E. H. Durfee and T. Montgomery. MICE: A flexible testbed for intelligent coordination experiments. In *Proceedings of the 1989 International Workshop on Distributed Artificial Intelligence (IWDAI-89)*, 1989.
3. B. A. Huberman and S. H. Clearwater. A multiagent system for controlling building environments. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 171–176, San Francisco, CA, June 1995.
4. V. Lesser. Preface. In *Proceedings of the First International Conference on Multi-Agent Systems*, page xvii, 1995.
5. M. Luck and M. d’Inverno. A formal framework for agency and autonomy. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 254–260. AAAI Press / MIT Press, 1995.
6. M. Luck and M. d’Inverno. Goal generation and adoption in hierarchical agent models. In *AI95: Proceedings of the Eighth Australian Joint Conference on Artificial Intelligence*. World Scientific, 1995.
7. M. Luck and M. d’Inverno. Engagement and cooperation in motivated agent modelling. In *Proceedings of the First Australian DAI Workshop, Lecture Notes in Artificial Intelligence, 1087*, pages 70–84. Springer Verlag, 1996.
8. J. M. Spivey. *The Z Notation*. Prentice Hall, Hemel Hempstead, 2nd edition, 1992.

Appendix: Z Extensions

We have found it useful in this specification to be able to assert that an element is optional. The following definitions provide for a new type, `optional T`, for any existing type, `T`, along with the predicates `defined` and `undefined` which test whether an element of `optional T` is defined or not. The function, `the`, extracts the element from a defined member of `optional T`.

$$\text{optional } [X] == \{xs : \mathbb{P} X \mid \# xs \leq 1\}$$

$[X]$
<code>defined</code> $_$, <code>undefined</code> $_$: $\mathbb{P}(\text{optional } [X])$ <code>the</code> : <code>optional</code> $[X] \rightarrow X$
$\forall xs : \text{optional } [X] \bullet \text{defined } xs \Leftrightarrow \# xs = 1 \wedge$ $\qquad \qquad \qquad \text{undefined } xs \Leftrightarrow \# xs = 0$
$\forall xs : \text{optional } [X] \mid \text{defined } xs \bullet$ $\qquad \qquad \qquad \text{the } xs = (\mu x : X \mid x \in xs)$