The Social Robot Architecture:
Towards Sociality in a Real World Domain

C.F.B. Rooney, R.P.S. O'Donoghue, B.R. Duffy, G.M.P. O'Hare, R.W Collier

PRISM Laboratory, Dept. of Computer Science

University College Dublin (UCD)

Belfield, Dublin 4, Ireland

{Colm.Rooney, Ruadhan.ODonoghue, Brian.Duffy, Gregory.OHare, Rem.Collier}@ucd.ie

26.3.99

Abstract. This paper advocates the application of multi-agent techniques in the realisation of social robotic behaviour. We present an architecture which commissions agent-based deliberation without sacrificing the reactive qualities necessary in a real world domain, and which is situated within a social landscape through the use of an Agent Communication Language.

Proc. TIMR 99 "Towards Intelligent Mobile Robots", Bristol 1999.

Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 – 6161.

Report number UMCS-99-3-1. http://www.cs.man.ac.uk/cstechrep/titles99.html

The Social Robot Architecture:
Towards Sociality in a Real World Domain

C.F.B. Rooney, R.P.S. O'Donoghue, B.R. Duffy, G.M.P. O'Hare, R.W Collier

PRISM Laboratory, Dept. of Computer Science, University College Dublin (UCD), Belfield, Dublin 4, Ireland

{Colm.Rooney, Ruadhan.ODonoghue, Brian.Duffy, Gregory.OHare, Rem.Collier}@ucd.ie

Abstract. This paper advocates the application of multi-agent techniques in the realisation of social robotic behaviour. We present an architecture which commissions agent-based deliberation without sacrificing the reactive qualities necessary in a real world domain, and which is situated within a social landscape through the use of an Agent Communication Language.

Keywords: *multi-robot systems, agent-based robotics, agent architectures, agent communication languages, co-ordination and collaboration.*

1 Introduction

Research into co-operative teams of robots performing constructive tasks has only recently come into its own. Within this work the majority of focus has been on reactive systems using relatively simple communication and often none at all. This paper proposes another approach. Based on the concepts of *social intelligence*, we present an architecture which commissions agent-based deliberation without sacrificing the reactive qualities necessary in a real world domain, and which is situated within a social landscape through the use of an Agent Communication Language (ACL).

Section 2 provides an overview of agent-based robotics. Section 3 offers an introduction to BDI architectures. An introduction to Agent Communication Languages is presented in section 4. Section 5 describes the *Social Robot Architecture* while section 6 animates its use. A discussion and conclusions are presented in sections 7 and 8.

2 Agent-Based Robotics

Initial research in the field of robotics focused on the behaviour [Bro86] [Ste94] and navigation problems of single robots. More recently, the area of multiple robots has demanded considerable attention. There are numerous advantages in the use of multiple robots, which include *inter alia* distributed capabilities; parallelism; task and load distribution; increased functionality with minimal complexity and graceful degradation upon component failure. However, despite these advantages, achieving coherent behaviour presents considerable challenges, none more acute than overcoming problems of co-ordination and interference. A great deal of work in the field of Distributed Artificial Intelligence has focused on such problems [Jen96]. It is our belief that Multi-Agent Systems (MAS) research can usefully be harnessed through its transference to systems of multiple autonomous robots, with significant insights obtained into such issues as co-ordination and interference.

Initially, agent based robotics emerged from cellular robotics where the robots had limited functionality and relied on swarm like intelligence to achieve their desired task, typically exhibiting *emergent* capabilities. Fukuda *et al* [CF95] [FNK+89] conducted research in the early stages of robot agents on both multi-robot systems in DARS (Distributed Autonomous Robotic System) and reconfigurable robots within CEBOT (Cellular Robotic System).

More recently, Dudek *et al* in [DJM+96] presented a taxonomy that classifies multi-agent systems according to communication, computational and other capabilities. Arkin and Balch [AB98] developed their reactive strategy for multi-agent co-operation with robots searching for trash objects, which they grasp and carry to wastebaskets. Their research focused on the development of behaviours for formation maintenance in heterogeneous societies of mobile robots.

Balch *et al* investigates the impact of communication in reactive multi-agent robotic systems in [BA94]. Balch later uses the soccer scenario for evaluating agents in terms of performance, policy convergence, and behavioural diversity [Bal97].

The next section provides an introduction to one approach, which has proven popular in the development of agent-based systems.

3 Belief Desire Intention (BDI) Architectures

Much research work has been commissioned on Multi-Agent Systems (MAS) and Distributed Artificial Intelligence (DAI) [BG88], [DLC89], [OJ96]. A number of competing agent architectures have been proposed in the literature. Two major architectural schools have emerged, namely those of the *reactive system school* and the *deliberative system school*. The former has predominated in the arena of autonomous mobile robot control. In this paper we advocate the synthesis of reactive and deliberative reasoning.

In the delivery of computationally tractable models of deliberative reasoning, one approach that has gained wide acceptance is to represent the properties of an agent using mental attitudes such as *belief*, *desire*, and *intention*. In this terminology, an agent can be identified as having: a set of beliefs about its environment and about itself, a set of desires which are computational states which it wants to maintain, and a set of intentions which are computational states which the agent is trying to achieve. Multi-agent architectures that are based on these concepts are referred to as *BDI-architectures* (Belief-Desire-Intention) [RG91] [Jen93] [OJ96], and have recently been the subject of much theoretical research. Proponents of the BDI approach argue that the understanding of the dynamics of these mental attitudes and their intimate interdependencies is crucial in achieving *rational* agent behaviour. We shall see later how the deliberative component of our architecture embraces the BDI philosophy and is developed using Agent Factory [Col96].

Multi-agent systems generally imply the existence of some medium that permits inter-agent communication, which supports their inherent social behaviour. In the next section we present one means by which agents may exchange information.

4 Agent Communication Languages (ACLs)

Communication in multi-agent systems ranges from none at all to interaction through high level agent communication languages. Among the approaches for agent communication are:

- Simple semaphores or signals. For example, in a reactive agent system a certain stimulus may trigger an agent to emit an alarm.
- Blackboard systems for the exchange of information.
- Simple message passing.
- Agent communication languages.

It is the latter approach that we investigate in terms of the IMPACT[1] project. Our agents are designed along the lines of the Belief-Desire-Intention architecture. As such, a language is required with sufficient expressive power to represent and express the concepts of beliefs, intentions, requests for information and services, replies to such requests, etc.

"A language with a precisely defined syntax, semantics and pragmatics that is the basis of communication between …… agents." [FIP97]

Agent communication languages are high level languages used by agents for the exchange of information and requesting services. They generally have a syntax powerful enough to generate a wide range of communicative actions.

There are a number of features, which an ACL should possess, and criteria by which they may be judged. Mayfield *et al* [MLF95] presents one such list of *desiderata* for agent communication languages. In brief:

Form. An ACL should, ideally, be syntactically simple, concise, yet extensible.

[1] Initiative for Mobile Programmable Agent-based Computing Technologies (IMPACT)

Content. A distinction should be made between the language, which expresses communicative acts, and the content language which conveys the information.
Semantics. The semantics of a communication language should be grounded in theory, and it should be unambiguous.
Implementation. The implementation should be efficient, provide an intuitive interface, which hides the message transport details, and be amenable to partial implementation.
Networking. The language should be suitable for implementation on top of modern networking technologies.
Environment. An ACL should be designed with the possibility of a heterogeneous and highly dynamic environment in mind.
Reliability. The language must be able to provide secure and reliable message transport.

One theory widely used in the development of agent communication languages is *Speech Act Theory.*

4.1 Speech Act Theory

Speech Act Theory was developed by Austin [Aus69] as a new theory of language usage. Existing theories treated utterances as either true or false, i.e. *constatives*. Austin specified a new class of utterances, *performatives*, which were not just used to say things, but rather to actively do things. Performatives are considered speech *acts*, that is, like any physical act, e.g. pushing, they cannot be true or false but they can be infelicitous. That is, they can fail. Performative verbs are those verbs which may be used to perform actions, e.g., request, promise, assert. It is with this theory in mind that we developed our ACL Teanga. An introduction to Teanga is presented in section 5.3.

These last few sections have presented an overview of the research landscape within which our work is situated. In the following sections, we shall provide an introduction to the Social Robot Architecture and its constituent parts, together with its development and use.

5 The Social Robot Architecture

This section presents the Social Robot Architecture in detail. We examine first the Social Intelligence Hypothesis, which acts as a grounding for this work. Following this we describe the architecture. Teanga, the ACL which empowers the social aspects of the system is then presented, and finally we describe how the system binds together into one coherent whole.

5.1 Social Intelligence

The Social Intelligence Hypothesis or Machiavellian intelligence hypothesis (recent discussion in [KDG97]) promotes the theory that in order to achieve a degree of intelligent behaviour from an agent, the agent must be both embodied in a physical environment and embodied in a social environment. This agent will therefore be subjected to complex dynamic social interactions in a real world, and this is believed [DC95] a necessity for the development of an artificially intelligent agent.

As indicated by Edmunds [Edm97], it is argued that "social intelligence is not merely intelligence plus interaction, but should allow for individual relationships to develop between agents". While the complexity of a social structure can grow exponentially with the size of the society, our research deals with a collective of homogeneous robots with relatively limited functionality and the development and implementation of a social structure between these robots. This requires that each robot have the ability to distinguish, identify, model and address other agents in a flexible manner. To achieve this, a number of tools are required to build an architecture that supports this degree of social functionality. An in-depth discussion of these tools is presented in [ODC+99] and the architecture itself is examined in detail in the next section.

5.2 The Modular Architecture

The computational machinery needed to facilitate team building and collaborative behaviour is non-trivial. We describe the *Social Robot Architecture*, which goes some way toward achieving this through the judicious synthesis of the reactive model with that of the deliberative model.

The *Social Robot Architecture* (figure 1) is comprised of four discrete layers: physical, reactive, deliberative developed using Agent Factory [OA96] [Col96], and social. We consider each in more detail in the subsequent sections. The modular architecture offers basic reactive behaviours for fast robot reflex responses to unexpected or dangerous events. These constitute a set of primary *survival behaviours* for the robot.

Any process intensive action (i.e. planning) retards the real-time reflexive nature of an agent.

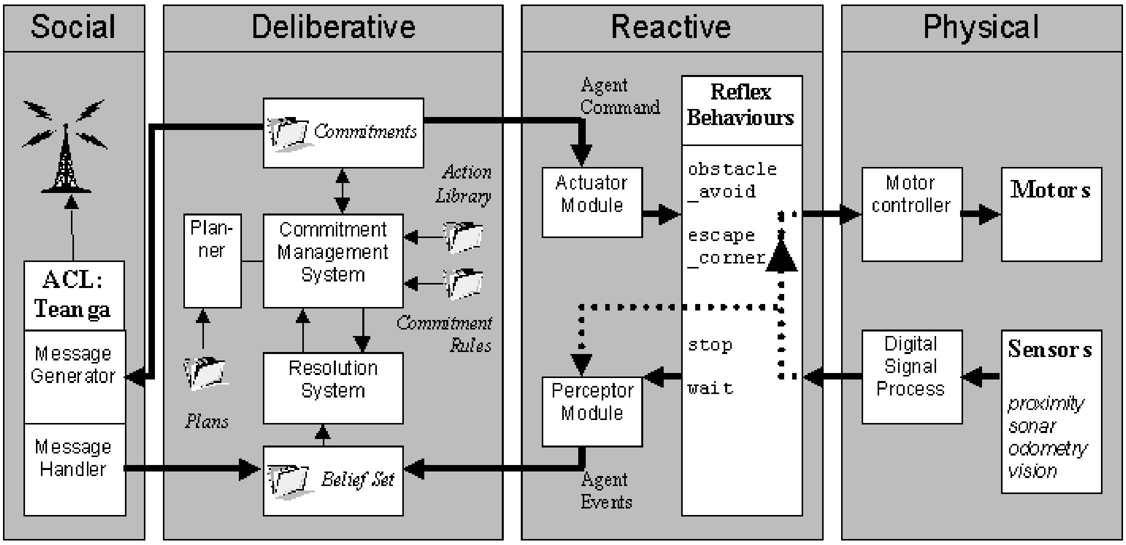

Additionally, more deliberative behaviour typically necessitates increased inter-robot communication, planning and deduction, which could result in a loss of perception granularity. Therefore, goal oriented behaviours are governed by the intentional agent structures of the deliberative machinery.

A library of robot behaviours exists which adopts a subsumption model [Bro86]. Higher levels provide increasing complexity and subsume lower level functionality (see figure 1). Reactive or reflex survival behaviours are implemented at the reactive level with more complex behaviours defined within the deliberative level. Each robot therefore has a degree of autonomy, which is independent of any inter-robot communication, and important in the realisation of robust complex robot behaviours.

We now consider the functional components of the architecture.

5.2.1 *The Physical Level*

Robots in our terminology may take the form of either that of a physical entity, specifically the Nomad Scout II from Nomadic Technologies or a simulated entity, provided by Nserver, a robotic simulator for the Nomad Scout II robots. Robots are situated in the world and exhibit behaviour based upon degrees of perception, action and interaction. Three such physical robots Aengus, Bodan and Bui[2] are used for experimentation.

5.2.2 *The Reactive Level*

A series of fundamental reflex behaviours are implemented at this level. They provide the basic *survival kit* of the robot necessary for dynamic and unpredictable environments. These *Survival Behaviours* include such behaviours as `avoid_obstacle`, `stop`, and `retreat`.

The sensory information received from the 16 ultrasound sensors, the bumper ring, and the odometry is processed at this level, resulting in clear agent-events being generated and communicated to the agent's deliberative level (i.e. `door_found(X,Y,Xi,Yi)`).

Agent-commands are also sent to this level from the deliberative layer. The actuator module acts as a behaviour arbitrator which decides the robot's actions depending on its input.

5.2.3 *The Deliberative Level*

The deliberative layer provides the deliberative machinery. This is achieved through a BDI architecture as described in section 3 and developed through Agent Factory.

---

[2] Aengus, Bodan and Bui are ancient characters in Irish legend and their graves can be found at Bru na Boinne, Ireland.

In essence, Agent Factory is a tool, which facilitates the rapid prototyping of Intelligent Agents. The Agent Factory System has been discussed more completely elsewhere in the literature [Col96] [OA96] [OCC+98]. The system offers an integrated toolset that supports the developer in the instantiation of generic *agent structures* that are subsequently utilised by a pre-packaged agent interpreter that delivers the BDI machinery. Other system tools support interface customisation and agent community visualisation.

An agent's deliberative actions are governed by its belief set, which is updated with the receipt of new agent states or robot sensor events and communication with other robots. An agent starts by gathering *perceptions* (raw sensor data). Low level robot processing filters these and thus achieves attention focusing, generating a set of agent events. If emergency action is required, a reflex behaviour is activated to overcome the immediate problem and this action is communicated as an agent event to the deliberative level. The perception process at this deliberative level deals with converting these agent events into beliefs and adding them to the belief set providing the agent with an up to date model of its current *perceived* situation. This situation is then analysed and the agents' commitments are updated accordingly. Pre-existing commitments are analysed and those pertaining to the current time frame honoured resulting in either a communicative act being sent to the social level or a physical act being passed to the actuators via the reactive level.

By way of example, a team of robots may have a goal of mapping their environment. Decomposing this problem, the individual robots might have to map constituent rooms, which in turn are made up of walls and a door. By analysing the data gathered from the sonar's these components might be identified. If a wall is found the robot then generates an agent event of the form, `wall_from(X,Y,Xi,Yi)`. Upon receipt the robot agent will revise it's belief set and accordingly add the belief `Bel(wall_from(X,Y,Xi,Yi))`.

5.2.4 The Social Level

Our agents interact via an Agent Communication Language (ACL), Teanga, which is described in more detail in section 5.3. The social level comprises of a message generator and a message handler. When a message is received the message handler decodes, checks and passes it on to the deliberative layer where it is dealt with accordingly, e.g. the belief set is updated or a commitment adopted. When a commitment to send a message is acted upon, the commitment management system utilises the services of the message generator to correctly generate and send the appropriate message. The hardware required to achieve this inter-robot communication is the Mercury-EN wireless unit from Nomadic Communications Inc, and the RangeLAN from Proxim Inc.

5.3 Teanga

Our ACL, which is called Teanga[3], is based upon Speech Act Theory. Teanga consists of 4 basic categories of *communicative acts* (messages): informatives, directives, commissives and declarations. This categorisation was developed from a classification of performatives proposed by [Sea76]. Within each category there are more specialised types of communicative acts, e.g. `drop_commitment`. Figure 2 presents an overview of the language in terms of its functionality. It illustrates the four distinct classes of speech act along with their sub-types. The language is designed as a carrier for an application-dependent content language. A content language is used to convey the information in the message, while the ACL provides context. For example, done(map_room) could be a statement of a content language. As it stands if it were sent as a message its meaning would be ambiguous at best. We therefore couch it in our ACL e.g. **COMMIT ACHIEVE** done(map_room). Changing the performatives also alters the context, e.g. **REQUEST ACHIEVE** done(map_room). The content language has a syntax and semantics separate to those of the ACL. We do, however, place some constraints on prospective content languages. In the context of the Social Robot Architecture the content language must be able to allow, amongst other things, the representation of actions (including speech acts) and their status e.g. done, doing; and an agent's mental states, e.g. beliefs and commitments.

[3] Teanga is Irish for "language" and interestingly is an anagram for a agent

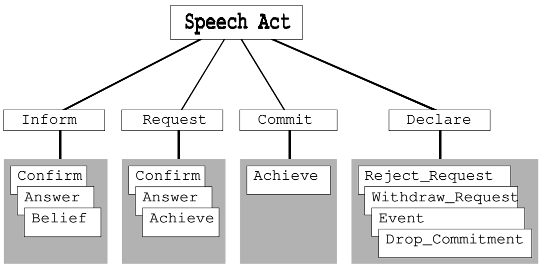

Figure 2. Functionality overview of Teanga

One of the main reasons why we chose to develop our own ACL rather than work with an established language such as KQML [LF94] or FIPA's ACL [FIP97] is that our language should be compositional. We wish to be able to support nested speech acts (and speech acts contained inside composite actions, e.g. plans). To implement a nested communicative act in KQML requires that the content be a KQML message. However, in KQML the content language is independent and there is therefore no content checking. *"A disadvantage of content independence is that it prevents the content from being checked for compatibility with the speech act type"* [CL95].

Examples of the communication structure follow. The first shows the generic form of a speech act represented in EBNF. The other examples are fully formed communicative acts.

<speech act>**::=SPEECH ACT (**<sender><recipients>**)**
(<structure>**)**

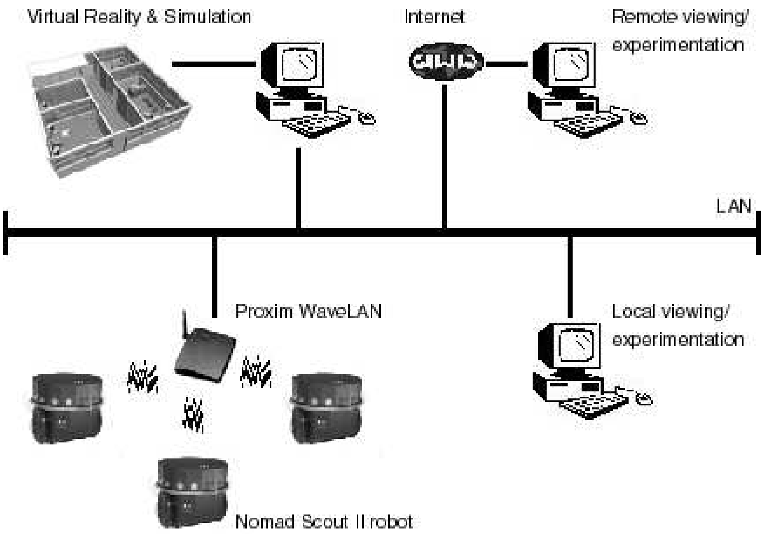

SPEECH ACT(Aengus Bodan) (**INFORM BELIEF**
Bel(wallfrom(x1y1;x2y2)) now)
SPEECH ACT(Bodan Bui) (**REQUEST ACHIEVE**
followwall(x1y1;x2y2) now)
SPEECH ACT(Bui Bodan) (**COMMIT ACHIEVE**
followwall(x1y1;x2y2) now)

5.4 A Coherent Whole

The Social Robot Architecture allows the seamless integration of a group of autonomous mobile robots into a larger, web-accessible system via the use of Teanga, an expressive agent communication language as shown in figure 3. This facilitates the rapid-prototyping of multi-robot communities within a social metaphor.

In addition to providing local and remote access to physical robots, this system provides a powerful development and experimentation environment, with rich visualisation and simulation support. Virtual reality provides the medium for remote observation of robot experiments [ODR+99]. Nserver, a robotic simulator for the Nomad Scout II robots, allows for the initial testing of robot behaviours before implementation on the physical robots. The subsequent section animates this architecture by considering an appropriate example.

6 The Robot Waltz

One experiment, which exercises the degree of architectural and behavioural functionality required to demonstrate the applicability of the *social robotics architecture*, is that of a waltz duet. The particular waltz consists of a simple square shape repetitive step. The objective is for two robots, Aengus and Bui to perform the waltzing behaviour side by side in a space restricted environment. It is a simple formation maintenance experiment similar to [AB98].

Motion Mechanics	Physical
Avoiding wall & obstacles	Reactive
Individual Waltz	Deliberative
Group Waltz	Social

Figure 4. Architecture Levels applied to Waltz

Figure 4 illustrates what occurs at each level of the architecture during the co-operative waltz task. Initially, the deliberative components of both robot agents adopt a commitment to start a waltz. When this commitment is evaluated, dialogue results at the social level through which the robot agents agree to start. The deliberative level then generates an agent motion command, which is sent through the reactive level and from there on to the physical level. This command results in the robot driving forward, sending position updates back to the deliberative level as agent events until one of two discrete situations are recognised. The first is that the deliberative level adopts the belief that the robot agent has driven far enough and that it must turn and begin the next step. In the second, sonar information being fed back from the physical level through the reactive level may indicate an obstacle. In this case the reactive level halts the robot and an appropriate agent event is sent to the deliberative level. If either of these events occurs the robot agent adopts commitments to stop, turn and begin the next step. As this is a social task, namely a duet, the robot agents also adopt a commitment to co-ordinate these actions. This results in another dialogue consisting of messages being generated and handled at the social level. Only when the robot agent receives confirmation from its partner it may deal with the commitment to turn and start the next step of the waltz. This is achieved through an agent command sent through the reactive level to the physical level indicating that the robot should turn and the cycle begins again.

This experiment highlights the judicious accommodation of reactive and deliberative based control in a single architecture. We have been able to demonstrate in simulation and in the real world the ability of robots to synchronise their behaviour thus testifying to the adequacy of Teanga as a communication language and wireless Ethernet as its medium.

7 Discussion

In both the fields of DAI and mobile robotics there are two major camps, the reactive and the deliberative. A number of researchers have proposed that neither of these approaches is sufficient [Ark89][Gat92]. Rather they advocate a synthesis of reactivity and deliberative reasoning. We follow such an approach with the addition of a high-level social component. This allows us to investigate a wider range of tasks, and also to apply work from the domain of MAS.

The Social Robot Architecture facilitates the development of teams of robots which may be heterogeneous, both in terms of behaviour and hardware. Behavioural heterogeneity is achieved through the use of Agent Factory to develop BDI agents with differing capabilities, e.g. commitment adoption strategies, knowledge resources, planning facilities, etc. The use of a standard interface between the deliberative component and the reactive layer (i.e. the agent-events and commands) means that robots with different sensors and actuators can be readily integrated into the system. The use of a standard interface for interaction, i.e. the ACL Teanga, allows these heterogeneous agents to communicate and coordinate.

8 Conclusion

This research has discussed *Social Robots* and an accompanying *Social Robot Architecture.* The architecture provides firstly, a powerful robot control mechanism, which integrates reactive and deliberative features and secondarily, a rich experimentation and visualisation medium, which offers seamless, contrasting, yet consistent views of social robotic behaviour. Subsequently *social robot* experiments may be observed and analysed via these mediums. This paper has focused upon the former, however the latter is considered elsewhere in the literature [ODC+99] [ODR+99].

The Social Robot Architecture offers a modular control architecture, which integrates separable layers into a coherent whole. In particular, a BDI deliberative layer provided through Agent Factory is integrated with a reactive and physical layer.

To date, experimental evidence has demonstrated that *constructive* and *coherent* collaboration is achievable in restricted task domains. We are confident that our social robot architecture is scalable and current research aims at demonstrating its applicability to different real-world problem domains.

On-going research seeks to extend our experimentation. More complex collaborative tasks are being investigated including collaborative grazing, and in the near future collaborative search between teams of robots equipped with vision systems. As robot soccer playing is an inherently collaborative and team based activity, it provides a suitable testbed for the development of the *Social Robot Architecture*. Layer based behaviours are currently under development to achieve suitable reactive and competitive behaviours within the context of the RoboCUP initiative [RoboCUP].

Acknowledgements

We gratefully acknowledge the support of Enterprise Ireland through grant No.'s SC/96/621 COMEDY and SP/97/08 IMPACT (Initiative for Mobile Programmable Agent-based Computing Technologies).

References

[AB98] Arkin, R.C., Balch, T. "Cooperative Multi-agent Robotic Systems", *Artificial Intelligence and Mobile Robots,* MIT/AAAI Press, 1998.

[Ark89] Arkin, R.C. "Towards the Unification of Navigational Planning and Reactive Control", *Working Notes of the AAAI Spring Symposium on Robot Navigation, Stanford University, 1989.*

[Aus69] Austin, J. L., *"How To Do Things With Words"*, Oxford University Press, 1969.

[BA94] Balch, T., Arkin, R.C. "Communication in reactive multi-agent robotic systems", *Autonomous Agents, 1, P1-25, 1994.*

[Bal97] Balch, T. "Learning Roles: Behavioural diversity in robot teams", AAAI workshop on Multi-Agent Learning, 1997.

[BG88] Bond, A., Gasser, L. (eds.), *Readings in Distributed AI*, San Mateo, CA., 1988.

[Bro86] Brooks, R. A. "A Robust Layered Control System for a Mobile Robot,", *IEEE Jour. Rob. and Autom., 2(1) 1986*

[CF+95] Cai, A, Fukuda, T., et al. "Hierarchical Control Architecture for Cellular Robotic System – Simulation and Experiments", *IEEE Int. Conf. On Robotics and Automation, 1995,*

[Col96] Collier, R. "The realisation of Agent Factory: An environment for the rapid prototyping of intelligent agents", *M.Phil., Dept of Computation, Univ. of Manchester Institute of Science and Technology(UMIST), 1996.*

[CL95] Cohen, P., & Levesque, H., "Communicative Actions for Artificial Agents". *Proc. Int. Conference on Multi-Agent Systems, AAAI Press, San Francisco, June, 1995.*

[DC95] Dautenhahn, K., Christaller, T., "Remembering, rehearsal and empathy – Towards a social and embodied cognitive psychology for artifacts", *Proc. AISB-95*

[DJM+96] Dudek, G., Jenkin, M., Milios, E., and Wilkes, D., "A Taxonomy for multi-agent robotics", *Autonomous Robots, Vol 3, 375-397, 1996*

[DLC89] Durfee, E. H., Lesser, V. R., Corkhill, D., "Trends in co-operative problem solving", *IEEE: Knowl. Data Eng. 11(1), 63-8, 1989.*

[Edm97] Edmonds, B. "Modelling Socially Intelligent Agents in Organisations" *at the AAAI Fall Symposium on Socially Intelligent Agents, Cambridge, MA, November, 1997.*

[FIP97] FIPA 97 Specification Part 2: Agent Communication Language. Foundation for Intelligent Physical Agents, 1997, http://drogo.cselt.stet.it /ufv/leonardo/fipa/spec/

[FNK+89] Fukuda, T., Nakagawa, S., Kawauchi, Y., Buss, M., "Structure Decision for Self Organising Robots Based on Cell Structures", *IEEE: Rob. & Autom., Scottsdale Arizona, 1989.*

[Gat92] Gat, E. "Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots" *appears in AAAI 1992.*

[Jen93] Jennings, N.R. "Specification and implementation of a Belief-Desire joint intention architecture for collaborative problem solving" *Int. Jour. of Intel. and Co-op. Info. Sys. Vol. II no3, 1993.*

[Jen96] Jennings, N.R. "Coordination Techniques for Distributed Artificial Intelligence", in O'Hare, G.M.P. and Jennings, N.R. (Editors.), Foundations of Distributed Artificial Intelligence, *Wiley Interscience Pub., New York, 1996, 296 pages. ISBN 0-471-00675*

[KDG+97] Kummer, H., Daston, L., Gigerenzer, G. and Silk, J., "The social intelligence hypothesis", *Weingart et al. (eds), Human by Nature: between biology and social sciences.*

Hillsdale, NJ: Lawrence Erlbaum Associates, P157-179.

[LF94] Labrou, Y. & Finin, T., "A Semantics Approach for KQML- A General Purpose Communication Language for Software Agents", 1994.

[MLF95] Mayfield, J., Labrou, Y., & Finin, T. "Desiderata for Agent Communication Languages", *Proc. of 1995 AAAI Spring Symposium on Information Gathering in Distributed Environments, March 1995.*

[OA95] O'Hare, G.M.P. and Abbas, S., "Commitment Manipulation within Agent Factory", *Proc. of Decent. Intel. & MAS, Cracow, Poland, 1995*

[OCC+98] O'Hare, G.M.P., Collier, R., Conlon, J. and Abbas, S., "Agent Factory: An Environment for Constructing and Visualising Agent Communities", *Proc. AICS98., 1998*

[ODC+99] O'Hare, G.M.P., Duffy, B.R., Collier, R., Rooney, C., O'Donoghue, R., "Agent Factory: Towards Social Robots", *PRISM Technical Report TR99-1, Dept. of Computer Science, UCD.*

[ODR+99] O'Hare, G.M.P., Duffy, B.R., Rooney, C., Collier, R., O'Donoghue, R., "Social Robotics: Reality and Virtuality in Muli-Agent Robotics", *PRISM Technical Report TR99-2, Dept. of Computer Science, UCD.*

[OJ96] O'Hare, G.M.P. and Jennings, N.R. (Editors.), Foundations of Distributed Artificial Intelligence, *Wiley Interscience Pub., New York, 1996, 296 pages. ISBN 0-471-00675*

[RG91] Rao, A.S. and Georgeff, M.P., "Modelling Rational Agents within a BDI Architecture", *Prin. of Knowl. Rep. & Reas., San Mateo, CA., 1991*

[RoboCUP] http://www.robocup.v.kinotrope.co.jp/

[Sea76] J. R. Searle "The Classification of Illocutionary Acts". *Lang. in Society, 5, p1-24, 1976*

[Ste94] Steels, L., "Building Agents with Autonomous Behaviour Systems", The `artificial life' route to `artificial intelligence'. Building situated embodied agents. *Lawrence Erlbaum Associates, New Haven. 1994*