

# Where Should Complexity Go? Cooperation in Complex Agents with Minimal Communication

Joanna J. Bryson

Department of Computer Science, University of Bath, Bath BA2 7AY, United Kingdom

J.J.Bryson@bath.ac.uk

**Abstract.** The ‘Radical Agent Concept’ in this chapter is that communication between agents in a MAS should be the simplest part of the system. When extensive real-time coordination between modules is required, then those modules should probably be considered elements of a single modular agent rather than as agents themselves. The advantage of this distinction is that system developers can then leverage standard software-engineering practices and more centralized coordination mechanisms to reduce the over-all complexity of the system. In this chapter I provide arguments for this point and also examples, both from nature and from my own research in building modular agents.

## 1 Introduction

Animal intelligence is complex, but the semantic content of animal communication tends to be simple. In contrast, artificial agents tend to be relatively simple, yet a great proportion of the current agent research deals with communication and negotiation between agents in Multi-Agent Systems (MAS) [e.g. 34, 48].

In this chapter, I question the wisdom of this approach. I have been concerned for some time that the MAS coordination community is overlooking fifteen years of research in action selection for individual modular agents (IMA). I am taking the opportunity of the Workshop on Radical Agent Concepts to air work in progress on a fairly basic question: what parts of an AI system should be decomposed into agents, what parts should be decomposed into modules in a single agent, and what the difference between these might be. I will propose that the main difference is whether the interactions between modules is really best modeled as communication, or more standard software conventions such as method calls or pipes.

I begin with a discussion of agents which are clearly individuated: mammals. I then consider the problems that action-selection in IMA address, and how these relate to coordination techniques in MAS. I then give an example of a MAS agent where all the agents are IMA, and show how increasing the coordination between agents can be done with very minimalist signal communication, embedding the complexity in the IMA instead. The IMA in this system are built with an AI methodology based on object-oriented design, and can thus exploit standard software engineering for debugging. Finally, I summarize my claims.

## 2 Societies Without Language

We have many animal models of social agents capable of complex behavior, both individually and systematically, while communicating only by signal, not by language [24, 25]. A signal is bound to a single meaning, while a language implies arbitrary links between symbols and semantics, as well as other powerful mechanisms such as compositionality and grammar. Despite extensive research, to date only humans have been shown to have languages, yet even humans communicate a great deal through signaling.

Many such signals, when expressed by humans, are generally accepted to be without intention. For example, facial expression is generally (though not always) unintentional, and some signals, such as blushing are nearly impossible to control. Other signals we often consider deliberate, such as physical threats or bowing. However, these actions may also be often automatic — reflexive responses tuned by social conditioning, exhibited in response to emotionally salient stimuli [3].

Although their communication is simple, the intelligence and behaviors of animals are incredibly complicated, particularly when compared to artificial agents. Animals manage enormous numbers of sensors and actuators, and pursue many conflicting goals that require intricate behavior in unpredictable environments. Why in nature is the complexity of individual agent behavior consistently higher than that of inter-agent communication? What are the circumstances that might lead to the evolution of such an arrangement?

## 3 Motivations from Nature for Simple Communication

Although the underlying causes of evolutionary trends can seldom be proved deductively, they are still open to the scientific method. For example, if we find any hypotheses of interest, we can test their predictive power in the performance of artificial agents. The following are four hypotheses for explaining why non-human animal communication seems to occur at a very high, abstract level of signaling, rather than at a potentially more informative level of detail and complexity such as languages provide.

1. *The environment demands attention.* Attentional constraints are a fact of animal intelligence [17] — for whatever reason, animals simply can't attend to very many things simultaneously. Infrequently-occurring combinations of semantically meaningful percepts take time and resources to process [29], and such time and resources are often required for conventional perception and action [42]. Thus signals need to be sufficiently easy to recognize and sufficiently persistent so that other multi-tasking agents have a good opportunity to observe them.
2. *Detailed direction requires significant two-way information exchange.* The purpose of communication is to influence the action of another agent. But at what level should this influence take place? It would be very difficult for one agent to give another detailed directions (such as precise path plans of foot placement) in a complex dynamic environment. Each agent is itself best situated for detailed observation of its opportunities, dangers, and requirements, and has the most experience for guiding its own actions and learning their probable outcomes. For example, each individual deer in a herd will have different jumping ranges and current energy levels.

Thus if one agent wanted to micro-manage the behavior of another, it would require communicating current models *to* the directing agent, as well as transmitting rapid and detailed instructions *from* it. As stated in the previous point, such transmission is difficult to accommodate in nature.

3. *Perception is unreliable.* This means signals must be very distinct, which in turn provides interdependent limits on how many different signals can exist, and how long processing or generating them might take. (Human language has required a large number of evolutionary innovations in both generation and perception [20, 41].)
4. *Conspecifics share a great deal of intelligence.* This fact does not limit how detailed communication *can* be, rather it limits how detailed it *needs* to be. Species tend to largely share both perceptual abilities and behavior patterns. Further, a species' innate behavior repertoire is likely to include behaviors that benefit other individuals, such as their offspring or kin. Consequently, in many common situations there would be no need for one agent's detailed management of another. It could be sufficient to communicate only a high-level goal or even simply an emotion in order to trigger appropriate behavior in a conspecific.

Either individually or taken in combination, these hypotheses could explain why communication in animals is generally limited to the high level, such as directing attention to goals, rather than managing a goal's completion in detail. Humans offer the only possible counter-example to this argument, yet even between humans, the true efficacy of communication is often overestimated. Even with our exceptionally complex speech production apparatus [20] and high-speed perceptual-event processing [41], complex, structured language is not used much in urgent or simply time-critical situations. For example, during a sporting event, utterances are generally limited in length and highly ritualized.

Humans must hold enormous systems of shared beliefs in order to communicate new ideas verbally [18], yet even between close, collaborating colleagues or carefully trained working crews misunderstandings frequently occur [47]. On the other hand, very effective real-time task coordination can be generated with very simple, even non-linguistic utterances [1]. Language is used primarily in social contexts [30], where it can alter or update another person's knowledge (and therefore behavioral repertoire) during times of relative calm.

## 4 Communication vs. Control

The extent to which the previous four observations about animal communication apply to an artificial MAS will vary from system to system. A few agent systems are biomimetic for the sake of being biomimetic — that is, to serve as models of animals. The emphasis of this chapter is more general. Here I am focusing on the extent to which limited communication is desirable from a purely engineering perspective.

There is no longer any debate that it makes good engineering sense to break a large, complex program or system into smaller, simpler elements which communicate only in a limited way through a well-defined interface. This is widely accepted in conventional software engineering [4, 15, 40] as well as in artificial intelligence [7, 32, 37].

The question I am asking is about the role of communication in control. This can also be seen as the question of agent decomposition: what entities need to be agents? I have previously spent some time addressing the question of *behavior* / module decomposition within modular (behavior-based) agents [6, 10, 14]. I took the answers from standard software engineering and object-oriented design. First, modular decomposition takes place around variable state. A module's core is one or a few related pieces of variable state with specialized representations to support learning correct values as efficiently and reliably as possible. A module consists of that state, any sensing required to create that state, and whatever actions are dependent on it.

The 'radical concept' I am proposing in this chapter is that, to the extent possible, the complexity of an intelligent system belongs *within* agents, where it can be designed, addressed and debugged exploiting the advances of more-or-less conventional software engineering. What this implies about agent decomposition is that a module of a system should be an agent precisely when the hypotheses in Section 3 hold — when communication is costly and needs to be minimized.

In order for this claim to be true, coordination (also known as action selection or behavior arbitration) within modular agents must be easier than coordination within MAS. The next two sections examine the extent to which this is true.

## 5 Coordination in Individual Modular Agents (IMA)

Historically, there have been two dominant strategies for action selection in autonomous agents. The first is *constructive* (search-based) *planning* [e.g. 38, 39]. The second is simply hard-coding the relationship or prioritization between modules [e.g. 5, 23]. This latter strategy is sometimes referred to as *reactive planning*, because it replaces planning in a conventional agent with a system that reacts reflexively given the current context. In the last decade, a large amount of research has focused on combining these two strategies [e.g. 22, 27, 31, 33] by having constructive planning create the structures underlying reactive planning when there is time available to do so.

One strategy that has been almost universally adopted, at least within robotics and virtual reality, is the modular (often called behavior-based) approach [32, 43]. Behavior-based AI originally decomposed all of the agent's intelligence into semi-autonomous modules which coupled motor control (or more generally, action) to whatever sensing was required by that action. The system also provided for minimalist, decentralized arbitration confined to connections between modules' input and output streams [5]. Subsequent research has generally favored more centralized and programmatic arbitration but nevertheless continues to emphasize the importance of modularity in simplifying the design and coordination tasks [7, 26, 32]. I will refer to such agents incorporating modularity as Individual Modular Agents (IMA) to contrast with Multi-Agent Systems (MAS).

In contrast to the hybrid of constructive and reactive planning, relatively little work has been done to control autonomous agents via techniques anything like those employed in the coordination of multi-agent systems. One exception is the Agent Network Architecture (ANA) [35], which uses a system of spreading activation between modules. The activation sources are modules linked to perception and to goals. All modules

in the system are interconnected for passing activation; the connections are weighted to reflect relevance. The first module that passes a particular threshold activation is allowed to operate the agent. This mechanism is somewhat akin to MAS auction-style coordination. However, it still relies on significant engineering for both the weights (though these are sometimes learned) and the modular decomposition. Further, the Maes architecture is generally considered impractical for large-scale systems due to difficulties maintaining behavioral coherence [46].

A few other autonomous agent architectures also use MAS coordination strategies. For example, Humphrys [28] uses voting, and Correia and Steiger-Garção [16] use chains of pair-wise negotiation. Overall, however, I am unaware of any such architecture other than ANA that has a wider user base than its developers<sup>1</sup>. Lack of current use does not *necessarily* imply that these architectures or strategies are not potentially useful, but it does increase the probability that they are being out-competed for some good reason.

## 6 Could MAS Coordination be Useful for IMA?

The concept presented in Section 4 claimed that a module should only be treated as an agent if constraints on communicating information between that and other modules are animal-like – that is, if perception is slow and noisy so signaling takes a long time relative to the action-selection cycle of the agent. If this is true, then it implies that there is no reason to use MAS coordination within a single agent.

In other words, I have no difficulty understanding when MAS technology is used on the World-Wide Web or across physically distributed problem solvers. But this paper is more concerned with the fact that MAS coordination technology is increasingly being applied into areas that have traditionally seen the application of other sorts of modularity, such as robotics [44] or shop-floor planning and scheduling [2]. In this context, it seems worthwhile to consider whether the arbitration schemes developed for IMA may be more efficient than coordination schemes developed for MAS. Or in other words, does it ever make sense to use MAS strategies to coordinate IMA?

At least two conditions have to hold in order for MAS coordination strategies to be useful for IMA.

1. The coordination process (voting, negotiation, bartering or such) must conclude quickly enough that its outcome (e.g. the action selected) can be expressed in the time-window when it is needed. Time has been a considerable problem for constructive planning, and underlies the wide-spread acceptance of reactive planning as at least one element (if not *the only* element) of on-line autonomous agent action selection. Unless a MAS coordination strategy can guarantee a good solution in a short, fixed length of time, it cannot be used for time-critical action selection.
2. The engineering process is shorter for developing the MAS coordination system than for specifying a reactive (hard-coded) plan for doing the same thing. Unlike the previous point, this is a pragmatic argument rather than a logical necessity — one could choose the slower development strategy, but it would be a waste.

---

<sup>1</sup> Another exception might be systems like the Open Agent Architecture [36] or Retsina [45], which are to some extent agent architectures, but I do not include them here since they are fully-fledged MAS.

For an IMA with heterogeneous modules, it seems likely that the cost of programming MAS coordination will be at least of the same order as that of encoding prioritizations for reactive planning. This is because each heterogeneous module must be described in a set of parameters meaningful to each other module. This presumably amounts to roughly the same amount of work that is necessary for reactive planning as specifying priorities per context. Given that this prioritization must take place anyway, the aforementioned timing advantages of reactive planning give it the edge.

The above points do not necessarily imply that MAS coordination can never be useful for IMA, but they do provide criteria for characterizing an IMA for which MAS coordination would be useful:

1. Part of the agent's remit must include some processes which do not require response times so quick that they violate the first point above. The prevalence of agent architectures that still include constructive planners as one of their elements is an indication this condition holds for some agents [see further 22].
2. The agent must contain a sufficient number of homogeneous modules that the second point above does not apply.

It is the second of these conditions that almost never holds in current IMA. Most modular decomposition strategies focus on functional or behavioral criteria. Even where behaviors are largely homogeneous, the slight difference between them is critical to action selection. For example the modules controlling each leg of a six-legged robot operate identically, but at different times (e.g. during the gait) or in different contexts (e.g. left vs. right turns).

This is also not to say that MAS coordination can't be useful for MAS. In many cases, individual MAS agents may be homogeneous from a software perspective, and differ only in terms of the hardware on which they are running or the users they represent. In this case, it may be that establishing prioritization can be done through existing standards (e.g. using money), while communication might be too slow for conventional IMA control. In this case, the arguments in Section 3 would come into play.

## **7 Example: Where Complexity goes in a MAS of IMA**

I will now illustrate the approach I am advocating — placing complexity within modular agents in order to simplify communication — with a simulation which does happen to be biomimetic. The system is a pilot study on a simulation being built to explore a hypothesis about social organization across species of non-human primates. Again, although this example is biomimetic, I believe the engineering approach taken may make sense for more standard, industrial modular intelligent systems.

The research shown here derives from the work of de Waal [19] on the evolution of specifically social behaviors such as those used in conflict resolution. In particular, this research is intended as part of the thesis work of Flack, who is seeking to understand the interaction between the number of conflict resolution behaviors expressed by a species and the social structure of its colonies [11]. Flack's research explores a hypothesized trade-off between egalitarian social models and complex conflict-resolution behaviors. This research has relevance to issues such as behavioral specialization within communities and time spent monitoring other agents.

## 7.1 Approach: Behavior Oriented Design

The IMA in this simulation were developed under Behavior Oriented Design (BOD) [14]. BOD modules are called behaviors [5, as in]. BOD uses hierarchical reactive plans for action selection / behavior arbitration. Behaviors are encoded in any standard object-oriented language; the reactive plan representation is more specialized, though also relatively standard [13]. The interface between plans and behaviors is an encoding of plan primitives in terms of methods on the objects which encode the behaviors.

As mentioned earlier, BOD takes inspiration from Object-Oriented Design [15, 40] in that variable state serves as a starting point for determining modular decomposition. Further, BOD provides a set of heuristics for simplifying an agent by shifting the specification of intelligent control between its basic representations, the behavior modules and the reactive plans. Details of BOD are available elsewhere [9, 14]. To understand the argument relating to the theme of this paper it is sufficient to see the increase in complexity and/or number of behaviors (where complexity is determined both by the amount of state contained and the amount of program code, as roughly indicated by the number of primitives / methods supported), and the size of the reactive plans.


## 7.2 Simulation Results — Impact of Increased Communication

The simulation concerns simple agents with two conflicting drives. One is social: the desire to groom. The other is individual: the desire to wander in relative isolation — a stand-in for slightly more complex behavior such as foraging. Each drive is placated when the agent engages for a while in activities consummatory to that drive (grooming or wandering). If the agent has no active desires, it rests in place. The agents also have a desire to avoid bumping into each other which is constant — it invokes avoidance behavior whenever another agent blocks its path.


The results shown in Figure 2 are for 16 rectangular agents co-inhabiting a walled, rectangular enclosure 140 times larger than each agent. For each condition shown in the figure the simulation was run until 11,000 behavior transitions had been recorded, or approximately 690 per agent. By ‘behavior transition’ I mean changing from one expressed behavior (e.g. sitting) to another (e.g. approaching a grooming partner.)

The results show the amount of time spent engaged in three activities: grooming, attempting to groom, and avoiding jostling. Grooming is a consummatory (fulfilling) action for the grooming drive — the agents would prefer (are driven) to engage in it 14 percent of their time. Attempting to groom covers a range of behavior including approaching another agent and aligning with their body in a way to facilitate grooming. These behaviors are not in themselves rewarding, but are motivated by the grooming drive. Similarly, avoiding jostling is not exactly consummatory, it simply facilitates other, more useful behavior. Thus the optimal time use for any individual agent would show the grooming bar at 14 percent and the other two at zero. Of course, this ideal cannot be realized, but the simulations examine how it can be approached.

The difference between the three experimental conditions is essentially one of communication. In the first condition, the agents are completely unaware of each other’s actions other than their locations. In the second and third, the grooming agent transmits to the groomee its role as a target for grooming, and the groomee holds still. In the



**Fig. 1.** The architecture of a BOD agent. Behaviors ( $b_1 \dots$ ) generate actions ( $a_1 \dots$ ) based on their own perception (derived from *sensing*, the eye icon, and their internal variable state). Actions which affect state external to their behavior (*expressed actions*, the hand icon), may be subject to arbitration by *action selection* (AS) if they are mutually exclusive (e.g. sitting and walking). In this diagram, three actions 1, 3 and 5 are taking place concurrently, while 2 and 4 have been inhibited by action selection.



**Fig. 2.** Results showing the impact of adding simple communication to facilitate social grooming. In the first condition there is no communication, in the second recipients are informed of (and tolerate) active grooming, in the third they these behaviors apply also during the approach to groom.



second condition, this happens only when the groomer is actively grooming the agent. In this case, the communication is analogous to *feeling the action* of grooming. In the third condition, this transmission takes place during the groomee's approach. Here the communication is analogous to *seeing or hearing the intent* to groom from a distance.

It is interesting that the communication behavior in condition 3 does not significantly reduce the time the groomer spends approaching the agent over condition 2, but it does have significant impact on how much time the agents have to dedicate to their more primal drives, as witness the decrease in grooming time. Thus if the simulated agents were a valid model of some animal species, these results would indicate a selective pressure for the behavior patterns in condition 2, but not for those in condition 3.

### 7.3 Complexity in the IMA supporting Simple Communication

At what engineering cost was the cooperation shown in Figure 2 achieved? First, consider the design of the original agents in condition 1 (Figure 3). The behaviors are simplified, but the reactive-plan is precisely represented. The plan determines the agent's priorities, but both their assessment and their achievement are carried out by the behaviors. Items on the left of the plan are more abstract; solid vertical lines indicate sub-plans. Higher items on a sub-plan have higher priority. Parenthetical expressions are perceptual preconditions. Boxes surround primitive actions.

The farthest left sub-plan is the root of the plan hierarchy, referred to as the *drive collection* (D). Other sub-plans are called *competences* (C). The program cycle of the action-selection module has two phases. During the first phase, the highest priority drive which is currently triggered by its preconditions passes action attention to whichever of its children was most recently active. In the second phase, the child performs the next pending step in its own execution. The first phase provides for alertness or reactivity, the second for persistence within a particular plan. When a sub-plan terminates, attention is returned to the root of that particular drive. For full details of the plan representation (including quantitative experimental evaluation) see Bryson [8].

The exact mechanisms of BOD are not significant to the point of this chapter; they are drawn here partly to show the sorts of coordination developed in the IMA community — this system is similar to many other hybrid architectures (see Section 5 above) — but more to illustrate an implementation of very simple communication in an IMA MAS. Contrast Figure 3 with Figure 4. The additions are in italics. Communication is handled by the new action *notify*, which signals the target of the groomer to set its *groomed-when* variable to the current time. This allows a new perceptual primitive, *being-groomed?* to tolerate the grooming (that is, to hold still.) Notice that *notify* has been added into an action sequence with 'groom'. The only change in condition 3 (not shown) is that *notify* is added into the 'align' and 'approach' boxes as well.

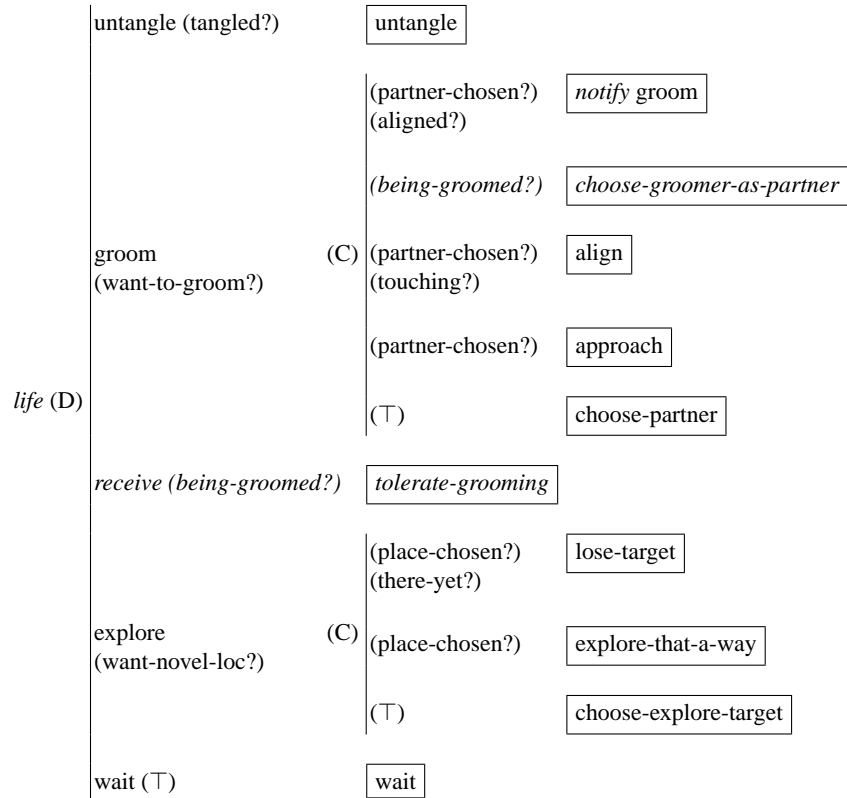
## 8 Summary

This chapter has argued that the agent-decomposition problem should be considered continuous with the behavior-decomposition problem. Sometimes a module within an intelligent system should be an agent, and sometimes it should just be part of an agent.



	<b>Navigate</b>	<b>Groom</b>	<b>Explore</b>
<i>state</i>	x, y, size, name focus-of-attn	drive-level, partner <i>groomed-when, being-groomed?</i>	drive-level direction-of-interest
<i>actions</i>	approach wait, align untangle	groom, choose-partner partner-chosen? <i>tolerate, notify</i>	choose-new-location lose-target, explore want-novel-loc?

(a) Behaviors



(b) Plan

**Fig. 4.** Behaviors and action-selection plans for IMA in condition 2.

I have suggested that the heuristic for determining which condition holds has to do with communication: it is better to embed complexity within a single agent than within the communication process for two reasons: because communication is unreliable and because existing AI and software engineering techniques can reduce the complexity of within-agent coordination.

In Section 3 I motivated my claim by using natural intelligence as an example; I then illustrated it by modeling a natural system. I now close by returning to the hypothetical reasons why mammals communicate so simply in order to answer to the question: When is a module an agent?

- *The environment demands attention.* If communication cannot take place in time, then MAS and communication are not an option. This constraint clearly doesn't apply to agents which are situated in entirely stable, predictable environments with no real time constraints. However, do such agents exist and, if so, do they need to communicate? Communication indicates mutual interdependence, and therefore some form of time constraint. A program that can compute a solution quickly enough that time doesn't matter and return its result reliably doesn't need to be anything as complicated as an agent — it could be a subroutine.
- *Detailed direction requires significant two-way information exchange.* This might not hold for an artificial entity — it might know exactly what another entity needs to know, and be able to feed that information directly. In this case, again, there is no reason for these entities to be agents. They could use pipes instead of KQML.
- *Perception (communication) is unreliable.* This is obviously more true of satellite arrays than modules on the same CPU, but in general it is solved in MAS by protocols. That is, communication reliability issues have been replaced by increasing costs in time of transmission and in the complexity of designing two-way exchanges<sup>2</sup>. However, although modules in IMA may run in parallel, their action selection is often sequential. Consequently, if timing or reliability in communication are issues, MAS coordination strategies may be better than IMA ones.
- *Conspecifics share a great deal of intelligence.* Section 6 argued that only within homogeneous agents does it make sense to coordinate through negotiation, because between heterogeneous agents the amount of work creating a common negotiation framework might as well be put into creating a reactive-plan framework instead. When agents do share a great deal of common code, very sparse, simple communication should be possible.

These are obviously broad generalizations and claims. These issues should probably be studied more formally, and the procedures for agent / behavior module decomposition tested and refined. Further, their may be hybrid cases when *parts* of a MAS may want to be *temporarily* unified in an IMA-like structures. Gajos [21] describes a system that does this within an intelligent environment — IMA-like coordination is used to coordinate care for both individual users and individual rooms. Bryson et al. [12] propose

---

<sup>2</sup> A significant but frequently-overlooked contribution of the Subsumption Architecture [5] was the specification that messages are sent with no hand-shaking and no guarantee of reception. Consequently, no single packet could be crucial, but rather the entire communication stream could be viewed as homologous with continuous sensor information. Section 7 uses this model.

a similar solution, where an individual user agent's action-selection may temporarily absorb as a module agent-like services discovered on the Web.

Nevertheless, my conclusion is that the agent abstraction makes the most sense when transmitting information between components of a system takes some time. In that case, communication should be kept as simple as possible. Where modularity can help simplify the system further, it is generally best coordinated as an IMA, not a MAS.

## Acknowledgments

Jessica Flack, Will Lowe and Lynn Andrea Stein contributed to the research in this chapter. Will Lowe, Krzysztof Gajos and Marc Hauser helped refine earlier drafts.

This work was begun in the MIT AI Laboratory and moved to Franklin W. Olin College of Engineering. At both locations it was conducted under the sponsorship of the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Material Command, USAF, under agreement number F30602-01-2-0512<sup>3</sup>. The work in this chapter was completed at Harvard University in the Primate Cognitive Neuroscience Laboratory where it was sponsored by National Science Foundation (NSF) grant EIA-0132707.

## References

- [1] Philip E. Agre and David Chapman. What are plans for? In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 3–15. MIT Press, Cambridge, MA, 1990.
- [2] Sivaram Balasubramanian and Douglas H. Norrie. A multi-agent architecture for concurrent design, process planning, routing and scheduling. *International Journal of Concurrent Engineering: Research and Applications*, 4(1):7–16, March 1996. Special Issue on the Application of Multi-Agent Systems to Concurrent Engineering, (eds. Brown, Lander and Petrie).
- [3] John A. Bargh, Peter M. Gollwitzer, Annette Lee-Chai, Kimberly Barndollar, and Roman Trtschel. The automated will: Nonconscious activation and pursuit of behavioral goals. *The Journal of Personality and Social Psychology*, 81(6):1014–27, Dec 2001.
- [4] Frederick P. Brooks, Jr. *The Mythical Man-month: Essays on Software Engineering*. Addison-Wesley Publishing Company, Reading, MA, 20th anniversary edition, 1995.
- [5] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2:14–23, April 1986.

---

<sup>3</sup> The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, the Air Force Research Laboratory, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

- [6] Joanna J. Bryson. The reactive accompanist: Adaptation and behavior decomposition in a music system. In Luc Steels, editor, *The Biology and Technology of Intelligent Autonomous Agents*. Springer-Verlag, 1995.
- [7] Joanna J. Bryson. Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, 12(2):165–190, 2000.
- [8] Joanna J. Bryson. Hierarchy and sequence vs. full parallelism in reactive action selection architectures. In *From Animals to Animats 6 (SAB00)*, pages 147–156, Cambridge, MA, 2000. MIT Press.
- [9] Joanna J. Bryson. *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*. PhD thesis, MIT, Department of EECS, Cambridge, MA, June 2001. AI Technical Report 2001-003.
- [10] Joanna J. Bryson. Modularity and specialized learning: Reexamining behavior-based artificial intelligence. In Martin V. Butz, Pierre Gérard, and Olivier Sigaud, editors, *Adaptive Behavior in Anticipatory Learning Systems*, Edinburgh, August 2002. Springer. *forthcoming*.
- [11] Joanna J. Bryson and Jessica C. Flack. Action selection for an artificial life model of social behavior in non-human primates. In Charlotte Hemelrijk, editor, *Proceedings of the International Workshop on Self-Organization and Evolution of Social Behaviour*, Monte Verita, Switzerland, September 2002. *forthcoming*.
- [12] Joanna J. Bryson, David Martin, Sheila I. McIlraith, and Lynn Andrea Stein. Agent-based composite services in daml-s: The behavior-oriented design of an intelligent semantic web. In Ning Zhong, Jiming Liu, and Yiyu Yao, editors, *Web Intelligence*. Springer, 2002. *forthcoming*.
- [13] Joanna J. Bryson and Lynn Andrea Stein. Architectures and idioms: Making progress in agent design. In C. Castelfranchi and Y. Lespérance, editors, *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*. Springer, 2001.
- [14] Joanna J. Bryson and Lynn Andrea Stein. Modularity and design in reactive intelligence. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 1115–1120, Seattle, August 2001. Morgan Kaufmann.
- [15] Peter Coad, David North, and Mark Mayfield. *Object Models: Strategies, Patterns and Applications*. Prentice Hall, 2nd edition, 1997.
- [16] Luis Correia and A. Steiger-Garção. A useful autonomous vehicle with a hierarchical behavior control. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in Artificial Life (Third European Conference on Artificial Life)*, pages 625–639, Berlin, 1995. Springer.
- [17] Nelson Cowan. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Brain and Behavioral Sciences*, 24(1):87–114, 2001.
- [18] D. Davidson. *Inquiries into Truth and Interpretation*. Clarendon Press, Oxford, 1985.
- [19] Frans B. M. de Waal. Primates—a natural heritage of conflict resolution. *Science*, 289:586–590, 2000.
- [20] W. Tecumseh Fitch. The evolution of speech: A comparative review. *Trends in Cognitive Sciences*, 4(7):258–267, 2000.

- [21] Krzysztof Gajos. Rascal - a resource manager for multi agent systems in smart spaces. In *Proceedings of CEEMAS 2001*, 2001.
- [22] Erann Gat. Three-layer architectures. In David Kortenkamp, R. Peter Bonasso, and Robin Murphy, editors, *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, pages 195–210. MIT Press, Cambridge, MA, 1998.
- [23] M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pages 677–682, Seattle, WA, 1987.
- [24] Alexander H. Harcourt. Coalitions and alliances: Are primates more complex than non-primates? In Alexander H. Harcourt and Frans B. M. de Waal, editors, *Coalitions and Alliances in Humans and Other Animals*, chapter 16, pages 445–472. Oxford, 1992.
- [25] Marc D. Hauser. *The Evolution of Communication*. MIT Press, Cambridge, MA, 1996.
- [26] Henry Hexmoor, Ian Horswill, and David Kortenkamp. Special issue: Software architectures for hardware agents. *Journal of Experimental & Theoretical Artificial Intelligence*, 9(2/3), 1997.
- [27] Henry H. Hexmoor. *Representing and Learning Routine Activities*. PhD thesis, State University of New York at Buffalo, December 1995.
- [28] Mark Humphrys. *Action Selection methods using Reinforcement Learning*. PhD thesis, University of Cambridge, June 1997.
- [29] Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall, Englewood Cliffs, New Jersey, 2000. ISBN 0130950696.
- [30] Chris Knight, Michael Studdert-Kennedy, and James R. Hurford, editors. *The Evolutionary Emergence of Language: Social function and the origins of linguistic form*. Cambridge University Press, 2000.
- [31] Kurt Konolige and Karen Myers. The Saphira architecture for autonomous mobile robots. In David Kortenkamp, R. Peter Bonasso, and Robin Murphy, editors, *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, chapter 9, pages 211–242. MIT Press, Cambridge, MA, 1998.
- [32] David Kortenkamp, R. Peter Bonasso, and Robin Murphy, editors. *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*. MIT Press, Cambridge, MA, 1998.
- [33] John E. Laird and Paul S. Rosenbloom. The evolution of the Soar cognitive architecture. In D.M. Steier and T.M. Mitchell, editors, *Mind Matters*. Erlbaum, 1996.
- [34] Victor R. Lesser. Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture. *Journal of Autonomous Agents and Multi-Agent Systems*, 1(1):89–111, 1998.
- [35] Pattie Maes. The agent network architecture (ANA). *SIGART Bulletin*, 2(4):115–120, 1991.
- [36] David L. Martin, Adam J. Cheyer, and Douglas B. Moran. The Open Agent Architecture: A framework for building distributed software systems. *Applied Artificial Intelligence*, 13(1–2):91–128, 1999.

- [37] Marvin Minsky. *The Society of Mind*. Simon and Schuster Inc., New York, NY, 1985.
- [38] Hans P. Moravec. The Stanford Cart and the CMU Rover. In I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 407–419. Springer, 1990.
- [39] Nils J. Nilsson. Shakey the robot. Technical note 323, SRI International, Menlo Park, California, April 1984.
- [40] David Lorge Parnas, Paul C. Clements, and David M. Weiss. The modular structure of complex systems. *IEEE Transactions on Software Engineering*, SE-11(3): 259–266, March 1985.
- [41] E. Pöppel. Temporal mechanisms in perception. *International Review of Neurobiology*, 37:185–202, 1994.
- [42] Ronald A. Rensink. The dynamic representation of scenes. *Visual Cognition*, 7: 17–42, 2000.
- [43] Phoebe Sengers. Do the thing right: An architecture for action expression. In Katia P Sycara and Michael Wooldridge, editors, *Proceedings of the Second International Conference on Autonomous Agents*, pages 24–31. ACM Press, 1998.
- [44] Carles Sierra, Ramon López de Mántaras, and Dídac Busquets. Multiagent bidding mechanisms for robot qualitative navigation. In C. Castelfranchi and Y. Lespérance, editors, *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*. Springer, 2001.
- [45] Katia Sycara, Keith Decker, Anandee Pannu, Mike Williamson, and Dajun Zeng. Distributed intelligent agents. *IEEE Expert*, pages 36–45, December 1996.
- [46] Toby Tyrrell. *Computational Mechanisms for Action Selection*. PhD thesis, University of Edinburgh, 1993. Centre for Cognitive Science.
- [47] Mary J. Waller. The timing of adaptive group responses to nonroutine events. *Academy of Management Journal*, 42:127–137, 1999.
- [48] Gerhard Weiß, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, Cambridge, MA, 1999.