A Classification of Communicating Agents

Dominic Tristram, Myra Wilson
Department of Computer Science
University of Wales, Aberystwyth
dbt93@aber.ac.uk, mxw@aber.ac.uk

26.3.99

Abstract

There are many mechanisms with which agents in multi-agent systems can communicate. These mechanisms can be described in various ways, depending on the aims of those running the system concerned. This paper proposes a classification system for such communication methods so that they can be described in a standard way and be more easily compared to other such systems. Part of the classification consists of a coded definition, facilitating automisation of some development.

1 Introduction

In almost every multi-agent system, some sort of communication is used. There are numerous different forms this communication can take, varying both in the physical transmission medium used, the actual data transmitted, its context in the current task, and so on. Many researchers provide this information distributed around the descriptions of their work rather than stating the communication information separately. This is understandable because the communication is nearly always simply a means to an end rather than the subject of the work. It is not always clear, however, what form the communication takes, especially on the physical side. Not all of the information is always given.

The field of multi-agent systems could therefore benefit from a uniform and well-known standard for classifying such communication systems. This would help researchers describe their methods as they could give some simple representation that would be immediately understood, rather than having to use plain English which could be misinterpreted. Those reading the work would receive similar benefits for the same reasons - the information would be immediately available and unambiguous. It would also be consistent across all multi-agent systems, which would make comparisons between them much more straightforward. If the classification system refers to hardware as well as the actual communication, we also achieve a method with which one can specify hardware required for the agents more simply and unambiguously.

It is useful to be able to compare communication systems for several reasons. The most obvious advantage in being able to do so is when testing the level of communication required for a given task. Instead of a long description of the communications used to obtain a set of results, each different communication configuration can be classified. This would mean that when the level of communication to be used for the task is finally chosen, the designers can simply state that communication of a certain class can be used, leaving the actual physical communication methods for later consideration.

This abstraction of the actual physical hardware is not only useful within ones own experiments however. Comparison of communication systems can be used when purchasing agents for a system, or when a range of such systems is available. Instead of stating all of the actual communication capabilities of an agent, it would be simpler to be able to say it can communicate to a defined level. This communication level is more important than the physical actualisation of the communications hardware because it is relatively simple to convert an interface of one physical type to another, but much more difficult to change the complexity of communication an agent is capable of, especially when the control system is simple and most of the communication processing is performed in dedicated hardware.

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

A classification system has been designed that fulfills the above requirements. It is designed to be easy to use and understand, yet be simple and complete. The following sections describe how the information from any given multi-agent communicating system is 'encoded' using the classification, and how a given code is read. This is followed by some examples of real-world systems and how the classification translates to the mechanisms given within them.

2 Related Work and Background

In previous work containing multiple communicating agents there is no uniform classification system. In most cases there is no attempt to classify communication, but when an attempt is made the author usually chooses his or her own previous work, plus the communication used in similar projects, for comparison. It is also worth remembering that because multi-agent systems can be expensive to physically implement in hardware, most work on them is done in simulation. This means that designers often to do not consider the form that communication will take in a finished system as communication can be assumed to be perfect. This is often done by simply using shared memory to simulate the communication interfaces between agents, as in (Arkin *et al.*, 1993).

(Dudek *et al.*, 1993) go some way to classifying communication systems in their taxonomy for swarm robots. The system includes information on the configuration of swarm robots and other swarm-specific aspects and so is not directly comparable to this work, but is nevertheless interesting for its breakdown of communication. The communication side is broken down into range, topology and bandwidth - range has three possible values while the others have four. There is no method of combining these three different aspects of communication to form one overall classification, but the divisions within each aspect of communication themselves are interesting. Range is divided into none, near and infinite. Topology is divided into broadcast, address, tree and graph. Bandwidth is divided into high, motion, low and zero (this is judged on a cost basis). The point to notice here is that some grouping is necessary in every case - it is very difficult to cover every case explicitly.

Other classifications, such as those in (Ashiru & Czarnecki, 1998) are too abstract to be particularly useful. This is usually the result of work where the physical implementation of the communication is not so important as the actual protocol used, as opposed to the above cases where communication is either not really considered, or is considered in detail but with the whole of the rest of the system.

(Balch & Arkin, 1994) comes quite close to defining a useful classification of inter-agent communication. Here, communicating systems are divided by factors such as implicit/explicit communication and the type of data sent. Each class of communication was tested with a set of simple tasks to measure their effectiveness (forage, graze and consume). The classification used is closer to the method described here, but still has too much of the actual protocol included. It is also not general enough as it concentrates on the tasks described.

Matarić (Matarić, Autumn 1997) uses a similar system. She splits communication into two forms, *direct* and *indirect*. Direct communication is a communicative act with the sole purpose of transmitting information. Indirect communication is based on observed behaviour. Further to this, Matarić defines *explicit cooperation* - a set of interactions which involve exchanging information or performing actions to benefit another agent, and *implicit cooperation* which consists of actions an agent takes achieving its goal that indirectly help others.

The lack of a large number of physical communication details, and the fact that there are even fewer detailed discussions on the merits of communications at the physical level, shows that it is still usually the case that work on these systems is normally performed at an abstracted level.

3 Classification Method

The classification of any system is represented by a classification number and an optional alphanumeric sequence called the 'communication code' detailing the mechanism more fully. The number is provided as a broad generalisation of the systems' abilities while the optional code allows the author to specify exactly how the system is composed. This allows comparison between agents and systems using the classification number, while details can also be given in the longer code to describe the components of systems.

It must be noted that the system does not make an attempt to include analysis of any communication protocols or languages used. Including such information would not only make the whole process more complicated, but it is not necessary for the purpose of this work. This system is looking primarily at the complexity and capability of communications between agents in a system, not how these communications capabilities have been used. As such, it is more useful when applied to physical robots, but it still applies in simulation.

3.1 Classification Number

In general the classification number can be considered to rise in line with the complexity of the communication system it is describing, with zero representing a system with no communication and a higher number representing a system that uses fairly complex communication. Using this, one can see at a glance if one system uses more complex communication than another.

The classification numbers are shown in table 1. It can be seen that in general the number rises with complexity, from zero to seven. This range should cover any current and possible forms of communication, and while each classification category may be quite broad, the scale serves its purpose by providing information on general complexity. Any additional detail should be provided in the communication code.

A system with a zero rating performs no form of communication at all. In this class of system the agents may not even be aware of other agents as anything other than objects. They also ignore any form of environmental change effected by another agent, as otherwise indirect communication would be taking place. (Brooks *et al.*, 1990) describes such a system.

A system with a rating of 1 uses 'implicit' communication. This is where the action of agents on the environment can be used as data by other agents without the 'transmitting' agent taking any special action. A good example of this sort of communication is a lawn-mowing/grazing task where the action of an agent when it removes the grass (or other substance being 'grazed') results in subsequent agents knowing that an agent has already passed that way. What they then do with that information is irrelevant in this case. A discussion of this implicit form of communication appears in (Balch & Arkin, 1994). A slightly less clear level 1 system would be one similar to that described in (Hackwood & Beni, 1992). Such a system involves 'signpost robots' that can change the state of other robots that come near. A system such as this might be rated as 1 if the signpost robots were constant, regardless of any events or the identity of the robot. However, should the transmission be changed or switched on or off depending on the identity of the passing robot, the classification number would change to 2 as explicit communication will have occurred.

When a system has a rating of 2 it is performing basic broadcast communication. This is explicit communication at its most simple. For example, it could simply involve an agent turning a light on when it has grasped an object and off when it drops it. This communication is being broadcast to any other agents within range equipped with the correct sensors. It is explicit rather than implicit because the robot deliberately switches on the light for the sole purpose of communicating its actions. The communication would be implicit if the light was triggered as a result of grasping the object because of the way the gripping device worked. If this was the case, the system would be classified as 1 on the scale.

In a system rated 3, communication occurs in both directions. For example, if the agents in the example above were homogeneous, a robot receiving the signal from the first agent indicating it had grasped an object could return a signal informing the first agent that it had received the signal. This form of communication is much more flexible because information can flow in two ways. Even if it does not, the receiving agents can send acknowledgments to the transmitting agent, making the system more robust. (Matarić, Autumn 1997) is a good example of this, with the agents broadcasting their sensory information to others to eliminate sensor blind spots. This sort of communication can also be used to formulate a simple language, as in (Yanco, 1994).

In complex systems, it is often desirable to enable agents to communicate with another specific agent rather than all at once. This is referred to as *directed* communication, but it could also be thought of as 'addressed' or 'agent specific'. In a system classified 4 on the scale this is the form of communication that takes place. An example of this sort of system is a mine-clearing task where one planning agent controls numerous smaller agents individually to direct them to suspected mines. It could not broadcast this information without some form of addressing as it would have no way of differentiating between agents.

Sometimes it is useful for the agents in a system to not only be addressable, but be able to address others individually too. Such a system is quite a lot more complicated that just receiving directed messages as each agent must have knowledge of the other agents, rather than just its own ‘name’. Systems such as these are classified as category 5.

In some systems it may be necessary to transmit a large amount of data. For example, a system where an agent is acting as a remote sensor and returning a constant stream of information. As robots in these systems are not often hard wired to each other, using two-state (binary) communication can be slow when the amount of data to send is high. Greater bandwidth can be achieved by using more individual states. The more states a transmission method can use, the more ‘words’ the agent language can contain, and so the fewer unique words need to be transmitted to get the data across. Such a system would be classified as 6.

As inter-agent communication becomes more complex the extra bandwidth cannot always be found by using more transmission states or increasing the rate at which data is sent. It is sometimes necessary to break data up into smaller ‘packets’ which are then reconstituted into messages upon reception. Along with this the agents must also include error-checking and collision detection if the signal is to get through. This level of communication is classified as 7.

3.2 Communication Code

As well as the numeric classification, it is also desirable to present a more detailed description of the communication system as a whole. Instead of hiding this information behind a general rating it is better to present it as a condensed ‘code’ to show exactly how the system is composed. To do this we must consider several elements of the system:

- Transmission medium such as audio, radio, ultrasonic, etc
- Transmitted states, as in binary, quad-state, etc
- Transmission rate, such as ‘14.4k/S’
- Protocol, including details like handshaking, flow control, etc

To achieve a useful description of the system we must use all of the listed information as well as everything else that is relevant, then represent it in a uniform way. Each feature can be given a unique code, and examples of some of these codes are listed in table 2.

Codes for any particular system are built up from these basic two-letter sequences, with numbers added as appropriate. The basic syntax for any given system is *(system attributes)(interface1)(interface2)(...)*. Thus it it possible to specify the complete system as well as each communication interface, although most systems are unlikely to include more than one.

The syntax can best be demonstrated with an example:

```
(pkdiwl16)(ausrwd4rt5)
```

The above line describes several aspects of the system. Everything in the first set of brackets describes system-wide attributes common over all physical communication interfaces (such as word length). Each physical interface is then described in separate sets of brackets. In this case there is only one physical communications interface, but additional interfaces would simply be added after the first. Note that the code describes a whole system, which could include several heterogeneous agents. It is not the individual agents being defined here, but the *communication interfaces* in the system as a whole.

From the example it can be seen that the system communicates using packets (pk). Facts can not be assumed about the system simply from this however. For example, simply because packets are used it does not mean that they are addressed - it could simply be that the word length of the protocol is larger than that of the transmission device. Equally, we can not assume that the word length is greater - packets may simply be used so that communication between agents does not block other agents communications. It is clear that other information also needs to be provided.

The next code is ‘di’. This shows that the system uses direct addressing. That is to say, each agent will have a unique address and will act upon signals (or packets, in this case) marked with that address. This is not always implicit when broadcasting is not used (the target could be the last actively communicating agent, for example). The same is true when packets are used - this does not imply addressing. A broadcasting system might use packets if the word length is greater than the width of the transmission medium, for example.

The remaining information in the first set of data reveals details of the word length (wl). This is decided by the communication protocol used rather than the physical interface and so belongs in the first set of brackets. The word length is the actual length of the data, as opposed to the actual transmitted data length, which may be larger because of checksums and so on. If these are present they can be shown using the packet length (pl) attribute.

The example system only has one communications interface. This is specified within the second pair of brackets as an audio interface (au) that both sends and receives (sr). The word length (wl) of four shows that information can be sent or received 4 bits at a time. As this is the only interface in the system, the information could have been incorporated into the first set of brackets. This would, however, be less clear, so the separate physical interface definition is used.

The final information in the interface definition shows the baud rate of the device. In this case it is 5, but any number can be used here, including the (K), (M) and (G) modifiers. If the protocol or interface can vary the rate dynamically during operation the maximum rate is given.

When the system contains more than one communication interface common elements still go in the general definition, while differences go into the relevant interface definitions. For example:

```
(wl16)(snrawd16)(rvauwd4pk)
```

This sequence defines a system with two interfaces. The general characteristics show that the system uses a word length of 16 for all communication. The first interface defined is a send-only radio device with a width of 16 bits. The second interface is for receiving signals using audio. This interface is only 4 bits wide and uses packets to split the words up.

It is possible to define very complex systems featuring a number of interfaces using these codes. The well-known symbols also allow 'slotting in' of agents into a system. For example, a system could have a number of interfaces but two only ever communicate with each other in one direction. This could be written like so:

```
(wl16)(snrawd16)(<1>srrawd4pk|2)(srrawd4pk)
(<2>|rvrawd4pk)
```

The above definition shows a system with four interfaces. The first and third do not have other interfaces specified in their definitions, and because they both send and receive we can assume they communicate with other interfaces of the same physical type (radio, in this case). The second interface has a pipe to the fourth. This means that everything it sends gets sent to the fourth interface only (if the : symbol is used it may send signals to other interfaces as well). The definition of the fourth interface starts with the pipe symbol, so we can see that all of its input is from the interface on the other end of the pipe. Note that pipes may have multiple ends, so such an interface may in fact receive communications from more than one other interface, but they will all feature the '|2' or ':2' characters in their definitions.

When combined with a common communications protocol it is easy to see at a glance whether any particular hardware will work in the defined system. This is possible because, as well as actual communications systems, sets of hardware can also be defined in the same way. For example, if the system has been designed and tested in simulation, its creator may wish to see which of his physical systems are most suitable to run it. In this case he would go through a list of agents in his possession and fit their interface definitions to the system interface definitions. A research lab may have five robots with different communications capabilities, perhaps something like table 3. By referring to this list it is easy to see if any of the agents in their present configuration fulfill the interface requirements of the system.

Just as the communication code can make selection of physical hardware easier, it can help in simulation too. In a simulator that is aware of the syntax, the communication code can be entered and an agent automatically generated. This is very useful in systems where generic agents are taken to be 'perfect' and it is the communication system under investigation. For example, if a developer is running experiments to test whether communication improves efficiency in a particular task, designing all of the communication interface simulation is not a useful way to spend a lot of time. By supplying the simulator with the appropriate codes, it can do all of this work automatically, freeing the researcher to spend time on more important things. If the system requires a mobile agent equipped with infra red communications capabilities, such an agent can be constructed easily and supplied with a simple interface to the simulated communications hardware. As this hardware simulation is automatic it is the same for all agents in the simulator, thus providing a uniform 'look' to the simulated agents control system. This allows the developer more time to concentrate on the ideas behind the experiment with the knowledge that the actual physical communication simulation is taken care of. It also allows the designer to easily swap different interfaces (such as infra-red and audio) and let the simulator worry about their properties. This allows the designer to test every possible form of interface easily and choose the best available to build into the final physical design. This sort of functionality is suggested as a positive step in robotics in (Matarić, 1998) and other articles on the continuation of robotics research.

4 Summary and Conclusions

This classification system achieves two goals. Firstly it enables a simple numerical representation of communication complexity which is useful for comparing different communication systems. Secondly, the optional hardware representation allows a designer to easily specify the communication hardware requirements of the agent for use in communication. This provides a level of abstraction that has not been possible before. It also enables easy simulation and rapid prototyping of communicating multi-agent systems as the designer can simulate a generic mobile robot but automatically have the use of the correct communication hardware.

It is unlikely that any system such as this will become universally accepted in the field in the near future. However, for internal use within organisations and cooperating groups the system could save time and money. The standardised specification process and ease of taking a system from the simulation to the physical stage would reduce development times, as well as making it easier to use the same tools for different systems. It is hoped that software currently under development will show how effective the mechanism can be.

References

Arkin, R C, Balch, T, & Nitz, E. 1993. Communication of Behavioural State in Multi Agent Retreival Tasks. *Pages 588–594 of: Proceedings, 1993 IEEE Internationsal Conference on Robotics and Automation*, vol. 3.

Ashiru, I, & Czarnecki, C A. 1998 (May). Experiments in Evolving Communicating Controllers for Teams of Robots. *Pages 3498–3503 of: Proceedings, 1998 IEEE International Conference on Robotics and Automation*, vol. 4.

Balch, T, & Arkin, R C. 1994. Communication in Reactive Multiagent Robotic Systems. *Autonomous Robots*, **1**(1), 27–52.

Brooks, R, Maes, P, Matarić, M, & More, G. 1990 (July). Lunar Base Construction Robots. *Pages 389–392 of: IEEE International Workshop on Intelligent Robots and Systems (IROS '90).*

Dudek, G, Jenkin, M, Milios, M, & Wilkes, D. 1993. A Taxonomy for Swarm Robots. *Pages 441–447 of: Proceedings, 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).*

Hackwood, S, & Beni, G. 1992. Self-Organisation of Sensors for Swarm Intelligence. *Pages 819–829 of: Proceedings, International Conference on Robotics and Automation.*

Matarić, Maja J. 1998. Integration: Challenges and Solutions. *Pages 44–46 of: AAAI 1998 Spring Symposium Series: Integrating Robotics Reseach: Taking the Next Leap.*

Matarić, Maja J. Autumn 1997. Using Communication to Reduce Locality in Distributed Multi-Agent Learning. *Journal of Experimental and Theoretical Artificial Intelligence, special issue on Learning in Distributed Artificial Intelligence Systems.*

Yanco, Holly. 1994. *Robot Communication: Issues and Implementations.* M.Phil. thesis, MIT. Also available as MIT-AI-TR 1478.

Table 1: Classification numbers

Number	Description
0	No communication
1	Implicit communication
2	Broadcast
3	Broadcast bidirectional
4	Directed
5	Directed bidirectional
6	Directed > 2 state
7	Explicit directed with packets

Table 2: Data Codes

Code	Feature
Interfaces	
ir	Infra-red
au	Audio
us	Ultrasonic
ra	Radio
li	Light (visible)
sn	Send
rv	Receive
sr	Send and receive
wd	Transmission width
pl	Packet length
Characteristics	
di	Direct addressing
br	Broadcast
pk	Packet
wl	Word length
rr	Receive rate
tr	Transmit rate
rt	Send/receive rate
Symbols	
:	Communicates with
$\mid$	Piped to
(K)	x 1,000
(M)	x 1,000,000
(G)	x 1,000,000,000
<label>	Text label

Table 3: Agent Interface Examples

Agent	Code
Non-communicating	()
2-way radio	(rasrrt14(K)wd16)
IR	(irsrrt48wd1)
Receive only IR	(irrvrt32wd1)