Visualisation of Multi Agent System Organisations using a Self-Organising Map of Pareto Solutions

J. M. E. Gabbai[1], W. A. Wright[2] and N. M. Allinson[1]

UMIST, Manchester M60 1QD, England[1]
BAE SYSTEMS, ATC PO Box 5, FPC 267, Bristol BS34 7QW, England[2]

Abstract. The structure and performance of organisations – natural or man-made – are intricately linked, and these multifaceted interactions are increasingly being investigated using Multi Agent System concepts. This paper shows how a selection of generic structural metrics for organisations can be explored using a combination of Pareto Frontier exemplars, from extensive simulations of simple goal-orientated Multi Agent Systems, and exposé of organisational types through Self-Organising Map clusters can provide insights into desirable structures for such objectives as robustness and efficiency.

1 Background and Motivation

The increasing trends for complexity in industry and the lack of explicit Multi Agent Systems (MAS) research into organisations outside of coordination techniques [1, 2] is the basis of our research programme. An organisation's behaviour/performance is primarily a function of its environment, the composition of its individual agents and how they interact. In turn, how agents are connected determines an organisation's structure [3]. Our underlying research questions include how to measure and visualise the relationships between structure and performance, and why organisational structures/attributes (e.g. centralised, hierarchical) are more suited specific performance requirements.

To examine these research questions, an extendable simulation has been developed that can construct organisations of varying complexity and size to solve simple tasks. This is discussed in §3, proceeding an overview of our work to develop a method of graphing organisational structure and a set of generic metrics to quantify structural relationships.

2 Relating Organisational Structure to Performance

Much work has been carried out in the MAS and Distributed Artificial Intelligence (DAI) community to formalise individual utterances between agents [4]. These 'conversations' have been retrospectively examined using methods developed for Social

Network and Dooley Graphs, where nodes represent identities of particular agents as well as the state of information transferred [5]. In our simulation, we chart the conversations between agents and the type of information conveyed. The graphs can be described in matrix form, which lends itself to deeper analysis of organisational structure. Using this organisational structure matrix, we can quantify the type of structure based on a set of metrics. Below is an overview of some of these metrics. For a full description of how the network graphs are constructed, matrices developed and details of all metrics see [6].

- Centrality of communication – The overall cohesion of a network graph, indicating the extent to which a graph is organised around its most central point [7]. This metric has been adapted to cope with multiple connections afferent or efferent between two or more points. A measure of '1' indicates a fully centralised network while '0' indicates no communication or fully decentralised network where all communication is equal between nodes.
- Degree hierarchy – Krackhardt [8] developed a measure of degree of hierarchy that indicates the extent to which relations among the nodes in a network are ordered and there is little, if any, reciprocity. A measure of '1' indicates a fully hierarchical network while '0' indicates a flat organisational structure.
- Specialisation – In heterogeneous organisations, capabilities and skills will be unevenly distributed. For each particular capability, we measure the volatility of distribution in agents over the entire organisation. A measure of '1' indicates a fully specialised skill, meaning only one agent has a particular skill. '0' indicates that all agents have the said skill with equal degree.
- Heterogeneity of capabilities – The heterogeneity of capabilities looks at how capabilities are distributed throughout an organisation while '0' indicates that the sum of each capability throughout the group is equal. The greater the difference, the more this measure will tend towards '1'.

Using these generic metrics, we can relate them to organisations and their specific performance metrics in an effort to understand how structure affects performance, and aid our understanding and design of organisations.

3 Organisational Metrics Concept Demonstrator

The Java based Organisational Metrics Concept Demonstrator (OMCD) simulation is based on a two-dimensional grid which has no boundary conditions and where the agents have a simple "find and remove" objective. The agents move around the grid using a random walk searching for one or more 'targets'. When a target appears within an agent's search range, the agent communicates that a potential target has been found by placing a communication 'signal' around the target. The signal is strongest at the source, and tails off to zero at the edges. Agents that can remove targets and are within the signal's region will travel up the signal gradient to the source. The communication is recorded in a relationship matrix outlined in §2. To summarise, an agent, j, will have one or more capabilities, i, defined as c_{ij} where $j \in \mathbb{N} \wedge j \leq N_a$, $i = \{\text{search, remove, communicate}\}$ and N_a is the total number of agents in the organisation. c_{ij} describes the range of capability i; if no capability is present, c_{ij} is zero. Validation of the model and further details about the simulation are covered in [6].

3.1 Time taken to reach termination condition

The termination condition of the simulation is defined as the removal of all targets in the environment. The time taken, τ, to remove all targets is the average time taken from a set of simulations, or epochs ε, based on a single scenario configuration, but with random start positions. The average time taken has an error margin determined through standard statistical methods. The normalised time taken which also includes the 'cost of organisation' dimensional parameter is defined as:

$$\tilde{\tilde{\tau}} = \frac{\overline{\tau}}{D}\left(w_{N_a} \cdot \tilde{N}_a + \sum_{i=1}^{N_c} w_{c_i} \cdot \sum_{j=1}^{N_a} \tilde{c}_{ij} \right) \tag{1}$$

where D is the grid area and $\overline{\tau}$ is the average time taken. The number of agents and each agent capability is normalised over the entire data set so that $\tilde{N}_a$ and $\sum_{i=1}^{N_a} \tilde{c}_{ij}$ for $\forall j = \{0,1\}$. Weights can be assigned to the number of agents, w_{N_a} and the individual capabilities, w_{c_i} to signify relative financial cost or bias of capabilities. There are drawbacks in using such a weighting system for bias allocation which are well documented, but for this analysis such an approach is acceptable.

3.2 Robustness to failure

Resistance and adaptability to failure (or changes to environmental factors) ought to be a major design consideration in large organisations. This is increasingly the case in intricate organisations, where the need for "efficiency" has led to organisations and systems to be on the perpetual brink of collapse. Organisations have various modes of failure. They include single point failure of a member, agent or node through to full homogeneous failure where all members of a network with the same vulnerability fail (for example a virus affecting all unprotected computers running a specific operating system). In this discussion we explore the single point failure of the most influential agent.

We define influence as the degree centrality of an agent (communication afferent and efferent) combined with the contribution to global utility an agent provides via the removal of targets. For every scenario, agents are defined as belonging to an organisational set A such that $\forall j : a_j \in A$. To test the effect of failure, the most influential agent, a'_j, is removed. Formalised, we define this as:

$$a'_j : \max\left(w_d \tilde{c}_{D_j} + w_u \tilde{t}_j \right) \notin A \text{ for } \forall j \tag{2}$$

where $\tilde{c}_{D_j}$ and $\tilde{t}_j$ are normalised measures of agent degree centrality and contribution to global utility. Weights w_d and w_u are also assigned respectively. Once the most influential agent is removed, the scenario is repeated with the reduced organisational set. The new normalised time taken, $\tilde{\bar{\tau}}'$, as defined in (1) is then used to determine the robustness to failure, given by $\left(\tilde{\bar{\tau}}'/\tilde{\bar{\tau}}\right)-1$.

4 Simulation Results and Analysis

To explore the relationship between structure and performance of predominantly heterogeneous organisations, we examine a single premise in detail; namely, the materiel cost of the organisation is kept constant with three agents and 37 arbitrary unit^2 of capabilities that are distributed in every possible combination to the three agents (i.e. the bracketed term in (1) is constant although this changes for $\tilde{\bar{\tau}}'$). w_{N_a} and $\forall i : w_{c_i}$ defined in (1) are set to '1'; w_d and w_u defined in (2) are set to '½'. Environmental conditions are kept constant; the unbounded environment size is 100 unit^2 and the number of targets is three. In all, over 6,000 scenarios were run. The resulting dataset is presented in Fig. 1.

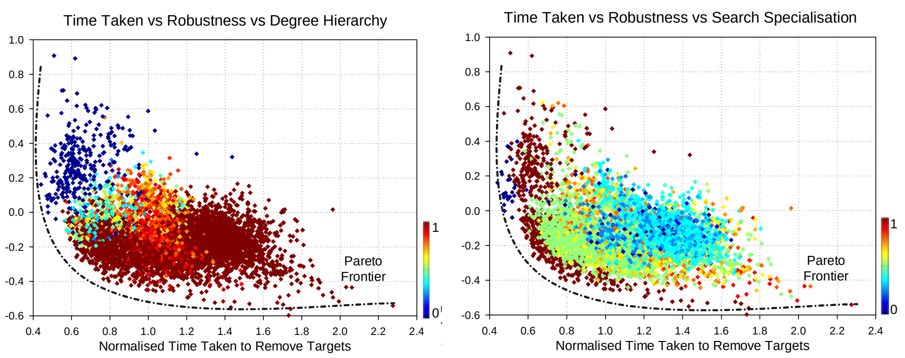

Fig. 1. Normalised time taken to remove all targets for a heterogeneous organization with fixed materiel, with the colour scale showing two different organisational structure metrics.

The surface in the above figures labelled "Pareto Frontier" [9] around the axis represents the trade off between time taken and robustness. Depending on requirements, an optimal point along that curve can be determined. In fact, we have found that when normalised time taken is plotted against robustness to failure for any set of organisational configurations the output always forms a Pareto frontier. The left hand graph shows how the degree of hierarchy affects the performance; flat organisations tend to be more efficient in removing targets but are prone to significant failure whereas hierarchical organisations are less efficient, but more robust. The right hand graph explores the relationship of search specialisation. We see that a homogeneous distribution of the search skill provides the best performance. Interestingly this is followed by a second band indicating that when only one agent has the search skill the organisation performs well, and is less prone to failure. However, it is difficult to pick up organisational types unless they are in the extremities as discussed above.

The high-dimensional measurement space that results from the use of numerous metrics thus requires the application of powerful data visualisation tools. One such tool is the Self-Organising Map (SOM), which permits an ordered two-dimensional representation of much higher dimensional input spaces [10]. In essence, a SOM employs unsupervised clustering techniques that can reveal meaningful groupings of parameters that may form distinct classes of organisations. Using the SOMine software package, a SOM trained output map for the data generated from this scenario is shown in Fig. 2

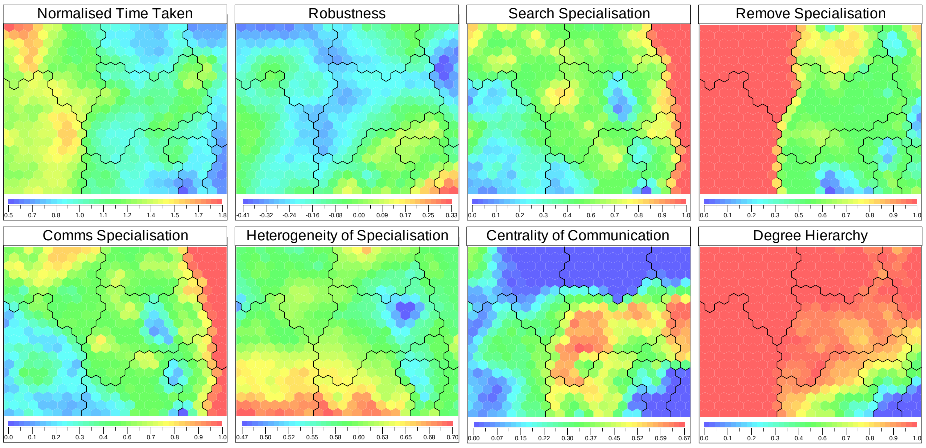

Fig. 2. SOM showing the two performance objectives (Time taken, and Robustness) and structural configurations that describe the organisations. Clusters are marked using boundaries.

Here we see how differing organisational structures are suited to subtly different performance requirements. A small cluster exists where uncoordinated search-heavy agents perform efficiently and are not too prone to failure (bottom left). Another region is where agents are far more specialised and rely a lot more on communication and coordination to be both efficient and robust (top right). We can also see that in general, organisations with a high degree of removal capability tend to perform better, but this performance is enhanced with regard to robustness agents communicate (middle right). Altering the cost weights in (1) will affect this observation. We can quantify these observations by highlight-

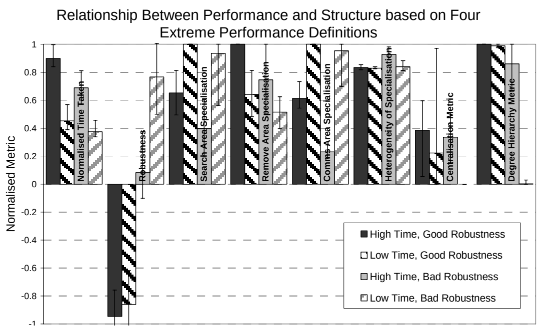

Fig. 3. SOM requirement extremities quantified

ing regions in the SOM and extract statistical data for all the parameters within the region. In this particular case, we are interested in the extremities of our organisational requirements.

Through visual inspection, four data sets for low/high time taken, good/bad robustness combinations were extracted and are shown in Fig. 3. Looking at the graph, we confirm that though organisations with no coordination perform marginally better in finding targets quickly, they are very poor at handling member failure. Organisations that have an element of centralisation and hierarchy, thus implying communication and coordination, tend to compensate for lack of brute force with communication.

4 Future Work

Our research program will be to extend the simulation and analysis framework so that corresponding organisational models can be incorporated into feedback learning systems. Much of our work is directed at the aerospace industry and in particular the coordination and resource allocation of Unmanned Air Vehicles (UAVs).

Acknowledgements: This work is supported by an EPSRC studentship and BAE SYSTEMS, for an Engineering Doctorate (EngD) based at UMIST.

5 References

[1] Odell, J., Parunak, V. D., and Fleischer, M., "The Role of Roles," *Journal of Object Technology*, vol. 2, pp. 39-51, 2003.

[2] Gasser, L., "Perspectives on Organizations in Multi-agent Systems," in *LNAI: Multi-Agent Systems and Applications*, M. Luck, V. Marik, O. Stepankova, and R. Trappl, Eds. Prague: Springer, 2001, pp. 1-16.

[3] Bonabeau, E., Theraulaz, G., J.L, D., Aron, S., and Camazine, S., "Self-organization in social insects," *Trends Ecol. Evoll.*, pp. 188-193, 1997.

[4] Cohen, P. and Levesque, H., "Communicative Actions for Artificial Agents.," presented at First International Conference on Multi-Agent Systems, 1995.

[5] Parunak, H. V. D., "Visualizing agent conversations: Using Enhanced Dooley graphs for agent design and analysis.," presented at 2nd International Conference on Multi-agent Systems, 1996.

[6] Gabbai, J. M. E., Wright, W. A., and Allinson, N. M., "Measuring and Visualizing Organizations by Relating Structure to Performance," presented at IEEE International Conference on Industrial Infomatics (INDIN 2003), Banff, Alberta, Canada, 2003.

[7] Freeman, L. C., "Centrality in Social Networks I: Conceptual Clarification," *Social Networks*, vol. 1, pp. 215-239, 1979.

[8] Krackhardt, D., Blythe, J., and McGrath, C., "KrackPlot 3.0: An improved network drawing program," *Connections*, vol. 2, pp. 53-55, 1994.

[9] Pareto, V., *Manual of Political Economy*, English translation ed. New York: Macmillan, 1971.

[10] Kohonen, T., *Self-Organising Maps*, vol. 30. Berlin: Springer, 1995.