Measuring and Visualizing Organizations by Relating Structure to Performance

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

[1] UMIST, Manchester M60 1QD, England
`j.gabbai@stud.umist.ac.uk, allinson@umist.ac.uk`
[2] BAE SYSTEMS, ATC PO Box 5, FPC 267, Bristol BS34 7QW, England
`andy-src.wright@baesystems.com`

Abstract

Within a Multi Agent System (MAS) environment, principled metrics are developed that encapsulate the structure and performance of organizations. From extensive simulation work, we can explore performance/cost/structure trade-offs; and, by incorporating data visualization techniques, we can observe the emergence of organization classes and begin to identify optimum organizational structures to meet specified constraints and tasks. We illustrate our approach through specific examples and suggest future directions.

1. Background and motivation

An organization be it a government health care system, a distributed (i.e. spatially and/or temporally separated) sensor array system, birds in a flock, or nodes in a telecommunications network, will usually exhibit emergent behavior. This emergent behavior is typically unintended and detrimental [2]. Not all emergent behavior is unfavorable: positive emergence can be found in ant path planning, bird flocking and the Internet [4, 5]. An organization with positive emergence is usually described as "*a whole greater than the sum of its parts*" [6-8]. Organizations that are closely linked to their environment and display adaptability and robustness to change are known as self-organizing systems [5, 6, 9].

In the past, man-made systems and organizations for the creation and maintenance of products have been relatively easy to design, understand and alter. This, however, is changing [2, 6, 7]. We can see this by comparing trends in agent function and use within current organizations and systems, with the subsequent challenge to our understanding of organizational behavior (see Table I, adapted from [4]).

Unfortunately, we have only a limited principled methodology of how to organize complex, interdependent, heterogeneous, semi-autonomous agents – and the infrastructure to support them – into aggregates with predictable, reliable, and stable behavior on a very large scale [4]. We also lack a solid understanding of which types of organizational structures are appropriate to which organizational function; a centralized organization favors complex but static problems, whereas a decentralized system will work well for a dynamic problem when the costs of reconfiguration are low [1], but there are many other organizational types (c.f. discussion about organizational spectrum in [10]) and requirements.

TABLE I
TRENDS AND CHALLENGES TO OUR UNDERSTANDING OF ORGANIZATIONS

Trends in Organizations/Systems	Challenge to out Understanding
Increase in number of agents	As agents are social, they communicate; if every agent is connected to every other agent, they will have a 'communications overload'. Agents may then spend more time processing information rather than acting on it. Building a system that balances information overload with global coordination is a challenge [1].
Increase in agent specialization	It is easier to make agent skills domain specific, rather than build a generic agent. However, efficient and effective agent coordination is required if heterogeneous agent cooperation is required to solve a problem. Specialization may also lead to centralized control.
Decrease in agent capability	Simple agents are easier to maintain. However, the desired behaviors in the organization may have to be engineered to emerge out of the simple agents' interactions, rather than explicitly design them in [1].
Decrease in resources and resource slack	As resources in a system are reduced, the interdependence between the agents increases. This requires good resource allocation and agent coordination [3].

These uncertainties are becoming more apparent to industry; quoting Sir Richard Evans, chairman of the defense company BAE SYSTEMS, "Systems capability has become more important than individual technologies and products. Obviously it's easier to make a single item, however sophisticated, than to integrate it into a large environment of complex devices and understand how it will perform" [11].

This critical issue can be applied not only to products but to entire systems and even organizations that exhibit high levels of interactivity and complexity. To explore these questions, we utilize Multi Agent Systems (MAS).

MAS can be seen as a collection of autonomous entities which interact with each other locally, affecting the global behavior of a system. Therefore agents are an appropriate way of exploring the questions raised above [12]. Unfortunately the MAS field (e.g. load balancing [13]), while advancing research in the architecture for individual agents and agent communications, has placed the exploration of agent society and organization as a peripheral theme, "primarily a specific coordination technique – not really one of the central intellectual issues of the field" [14, 15]. However, by emphasizing the plurality of agents and the organizational structure that binds them, the focus is shifted from designing ((***intelligent agent***) *systems*) to (*intelligent* (***agent systems***)). This may initially seem to be counterintuitive, but as agents get smarter, their functionality in fact reduces [1].

The increasing trends for complexity in industry and the lack of explicit MAS research into organizations outside of coordination techniques is the basis for the following research questions:

- An organization's behavior/performance primarily is a function of its environment, the composition of the individual agents and therefore the way they interact; how agents are connected determines the organizational structure [3, 10, 16]. How do we measure and relate the relationships between structure and performance?
- Which organizational structures or attributes (e.g. centralized, hierarchical) are more suited to which performance requirement, how and why?
- How can you describe the cost/benefit trade off as a function of an organization's structure?

Note that cost is defined as either a financial or biasing (i.e. promote or avoid some specified factor) attribute that is a function of the individual members in the organization, the cost of maintaining the organization and the cost of procuring the organization. For example, homogeneous organizations tend to be cheaper to procure and operate compared to organization comprised of heterogeneous members. Benefit is defined as a performance measure of how well the organization performs, such as resilience to individual failure, changing environmental conditions or achieving certain tasks within time and/or cost constraints. Usage of the term "cost" refers to all of these conditions

To examine these research questions, an extendable simulation that can construct organizations of varying complexity and size to solve simple tasks was developed. The primary goals of the simulation is to provide a test bed to measure organizational structures and relate this to various performance metrics. The agents in the simulation, if communicating, do not do so explicitly, but use stigmergy which leads to emergent behaviors, and this simplifies the structural analysis of the organization.

2. Organizational Metrics Concept Demonstrator

The Java based Organizational Metrics Concept Demonstrator (OMCD) creates a MAS scenario to examine the emergent organizational behavior of homogenous and heterogeneous agents, whose composition determines the organization's structure and therefore performance.

The simulation is based on a two-dimensional grid which has no boundary conditions 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 inside the signal's region will travel up the signal gradient to the source. This is a form on non-explicit communication that relies on stigmergy. Signals remain until the target is removed and can overlap, implying that agents will move up the steepest possible slope. To summarize, an agent, j, will have one or more capabilities, i, defined as c_{ij} where $j \in \mathbb{N} \wedge j \le N_a$, $i = \{\text{search, remove, communicate}\}$ and

TABLE II
LIST OF AGENT CAPABILITIES AND DEFINITIONS

Capability	Description
Search	The distance at which agent can sense a target. Defined as a range where $c_{\{\text{search}\}} \in \mathbb{N}$ and $\left(c_{\{\text{search}\},j} > c_{\{\text{remove}\},j}\right) \vee \left(c_{\{\text{search}\},j} = 0\right)$ for $\forall j$
Remove	The distance at which agent can remove a target. Defined as a range where $c_{\{\text{remove}\}} \in \mathbb{N}$ and $\sum_{j}^{N_a} c_{\{\text{remove}\},j} > 0$ This ensures that the simulation can end.
Communicate	The distance a signal will travel when placed over a target. Defined as a range where $c_{\{\text{communicate}\}} \in \mathbb{N}$ and $c_{\{\text{communicate}\},j} > 0 \Leftrightarrow c_{\{\text{search}\},j} > 0$ for $\forall j$

N_a is the total number of agents in the organization. c_{ij} describes the range of capability i; if no capability is present, c_{ij} is zero. Capability definitions along with mutually exclusive and overriding conditions are listed in Table II. Validation of the model and further details about the simulation are covered in [17].

3. Structural metrics

The previous section introduced a simulated scenario that simply by varying the capabilities of agents, the corresponding global behavior changes due to their interactions. This section discusses how these interactions can be described and recorded. Using non-temporal interaction matrices, the relationship between agents and targets can be recorded and measured. An example of a relationship chart is provided in Fig. 1. Similar approaches on visualizing agent conversations can be found in [18]. Note that the term A_n\{Search\} implies that an agent has a search capability, and by definition (see Table II) a communications capability, but no ability to remove targets.

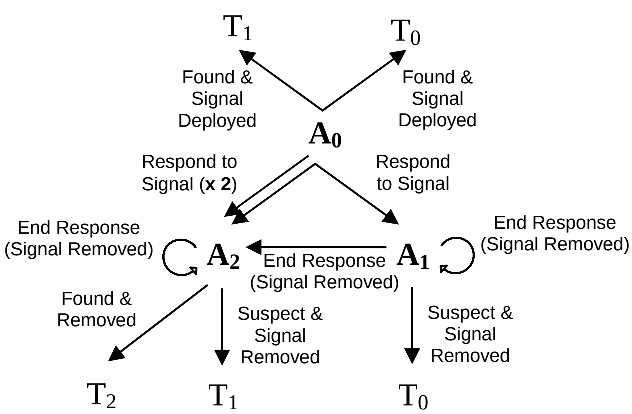

Fig. 1. Events and relationship between agents and targets of potential scenario. Shown are three targets, T_0, T_1 and T_2 and three agents, Agent A_0 \{Search\}, A_1 \{Remove\} and A_2 \{Remove\}. Agent A_0 finds T_0 and T_1 and deploys a separate signal for each target. At the same time, agent A_2 discovers T_2 without the aid of a signal and automatically removes T_2 because signals are not a prerequisite for target removal in this particular scenario. Agents A_1 and A_2 both respond the A_0's signal for T_0, as they are within the signal range and have a remove capability. A_1 arrives before A_2 and removes T_0, leaving it to indirectly inform itself and A_2 to stop moving up the signal gradient by removing the target and signal placed there. Following that, A_2 proceeds to find and remove S_1, following A_0's signal. The ↻ symbol is short-hand for describing agents indirectly instructing themselves. For the sake of diagrammatic clarity, targets appear twice to allow for an event based representation (top to bottom) of the process. Such an approach will differ from a temporal representation.

Based on these network graphs, we can construct an interaction matrix, **M**, that details the direction of interaction and the interaction itself so that it can be further examined, as shown in Table III. As the interactions are non-explicit an agent's causal command can only be linked to an agent (A_0\{Respond to Signal\} → A_1), rather than an agent's sensing function. If the agents were more complex, there would be merit in detailing which agent function affected which other agent's function.

TABLE III
EXAMPLE INTERACTION MATRIX TABULATING THE RELATIONSHIPS IN FIGURE 1

			To					
		Relationship Type	Agent 0	Agent 1	Agent 2	Target 0	Target 1	Target 2
From	Agent 0	Found Target				1	1	0
		Removed Target				0	0	0
		Respond to Signal	0	1	2			
		End Response	0	0	0			
	Agent 1	Found Target				0	0	0
		Removed Target				1	0	0
		Respond to Signal	0	0	0			
		End Response	0	1	1			
	Agent 2	Found Target				0	0	0
		Removed Target				0	1	1
		Respond to Signal	0	0	0			
		End Response	0	0	1			

The interaction matrix, **M**, can be used to measure various structural properties, which are discussed below.

3.1. Centrality of communication

Centralization refers to overall integration or cohesion of a network graph, indicating the extent to which a graph is organized around its most central point [19, 20]. We use this as a measure of the degree of centralization. The degree of a point is defined by the number of arrows efferent or afferent to the point in a network graph [19, 20].

The following equation has been adapted from the concept of degree centrality to cope with multiple links. Rather than give the centrality of the entire organization, it measures the centrality of communication; whilst other approaches take unconnected nodes into account, this method ignores them. The centrality of communication, C_C, is given by:

$$C_C = \frac{1}{n}\sum_{j=1}^{N_a}\left(1-\frac{k_j}{k^*}\right) \text{ where } k_j > 0 \tag{1}$$

The denominator k^* is the largest number of arrows afferent and efferent from a point or points in the organization, and n refers to the global number of connections between agents. Either agents are connected so $n = n+1$, or they are not. This is regardless of the *number* of links. N_a refers to the number of agents in the organization. It is important to note that the centralization of communication metric is time-dependant, as k^* can change. In the implementation considered here, this metric is applied to organizations that are considered stable.

3.2. Degree hierarchy

The notion of degree of hierarchy is based on the idea that all complex systems, including informal organizations, have a certain level of hierarchy [21, 22]. Krackhardt [23] developed a measure of degree of hierarchy that indicates the extent to which relations among the individuals in the organization are “ordered,” and there is little, if any, reciprocity. This measure of degree of hierarchy is defined as follows:

$$H_D = 1 - \left(\frac{V}{V^*}\right) \tag{2}$$

Where V is the number of unordered or reciprocated links in the organization (i.e. A_0 is linked to A_1 and A_1 is linked to A_0), and V^* is the sum of all links in the organization. A graph that is completely hierarchical will have no “reciprocated” or symmetrical links. The degree of hierarchy in a completely hierarchical network graph will be ‘1’, whereas a completely non-hierarchical graph will be indicated by a value of ‘0’.

3.3. Agent specialization

The specialization of agents in an organization is defined as the uniqueness of a single agent relative to other agents, which is useful when comparing different organizations’ capability distributions, particularly where the number of agents for each group is different. We do this by measuring the normalized deviation of each agent’s capability in relation to the group average, given by:

$$S_{c_i} = \frac{1}{2(N_a - 1)} \sum_{j=1}^{N_a} \left(\frac{\left| \overline{c}_i - c_{ij} \right|}{\overline{c}_i} \right) \quad \text{for} \quad 0 < i \le N_c \tag{3}$$

Where N_c refers to the number of capabilities and c_{ij} to the value of capability i where $i \le N_c$ specific to agent j where $j \le N_a$. This can either be true or false, or a quantity, so $c_{ij} = 0 \vee 1$ or $c_{ij} = 0 \vee \mathbb{N}$. The quantity refers either to a physical factor, or a relative level of quality. The above metric allows us to find the uniqueness of a capability across a group of agents. If all agents in an organization have a capability in equal measure, then S_{c_j} is zero. If only one agent has a capability, regardless of the quantity or quality, S_{c_j} will be ‘1’.

3.4. Heterogeneity of capabilities

The heterogeneity of capabilities, S_h, looks at how capabilities are distributed throughout the organization. If the sum of each capability throughout the group is equal, S_h is zero. The greater the difference, the more S_h will tend towards 1. For this comparison to work, $\tilde{c}_{ij} = \{0,1\}$, where 1 is defined as the maximum possible capability value, rather than actual group maximum, if it is known. S_h is defined as:

$$S_h = \frac{1}{N_a} \sqrt{\sum_{i=1}^{N_c} \left(\frac{C}{N_c} - \sum_{j=1}^{N_a} \tilde{c}_{ij} \right)^2} \tag{4}$$

Where

$$C = \sum_{i=1}^{N_c} \sum_{j=1}^{N_a} \tilde{c}_{ij} \tag{5}$$

4. Performance metrics

Performance metrics are more specific to the organization compared to the generally generic structural metrics. However, while the performance metrics discussed here are specific to the OMCD, the concepts can be applied to different domains.

4.1. 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 the 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 [24]. The normalized time taken is defined as:

$$\tilde{\overline{\tau}} = \frac{\overline{\tau}}{D} \left(w_{N_a} \cdot \tilde{N}_a + \sum_{i=1}^{N_c} w_{c_i} \cdot \sum_{j=1}^{N_a} \tilde{c}_{ij} \right) \tag{6}$$

Where D is the grid area and $\overline{\tau}$ is the average time taken. The number of agents and each agent capability is normalized 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, as discussed in Section 1. There are significant drawbacks in using such a weighting system for bias allocation that are well documented, and other alternatives will be considered, but for this analysis, this approach is acceptable.

4.2. Contribution to global utility

The contribution to global utility is simply a score that an agent has, based on the number of agent to agent and agent to target connections afferent and efferent to it, normalized by the number epochs required for scenario convergence. Applying this definition to Table II, the Agent A_0 would be the highest scoring agent. A more developed approach could be implemented where there is a weighting between the removal of targets and the centrality of an agent, but the weighting would be subjective, and so is avoided in this initial analysis.

4.3. Robustness to failure

The robustness metric measures the impact of organizational failure. This can be defined in several ways, such as failure of homogeneous agents, random failure, or failure based on a set criterion. The approach taken here is the latter, and is defined as the difference between normal time taken to remove all targets, τ, and the time taken after a specified agent has been removed, $\overline{\tau}'$, defined as the one with the highest contribution to global utility.

A more detailed discussion of the interaction chart, the interaction matrix, **M**, and how (1 – 5) can be implemented for a given **M**, as well as further normalization techniques for comparison of different OMCD organizations is provided in [25].

5. Cost vs. performance analysis in homogeneous organizations

We have previously demonstrated how we can link one performance attribute to one structural metric, namely robustness to centrality [10]. We now cover multidimensional data analysis of organizations. We begin the analysis using simple homogeneous agent organization scenarios that examine the cost/performance relationship of this model. The benefit of using homogeneous organizations to introduce our analysis and explore the search space is the implication that the structural metrics become uniform throughout the analysis. This implies that $C_C, H_D \to 0$ for $\varepsilon \to \infty$, and that at very low $\overline{\tau}$ conditions, C_C and H_D will appear more volatile given a computationally-viable acceptable error for the standard deviation. Furthermore as all agents are uniform, when checking for robustness, removing any random agent would result in a similar $\overline{\tau}'$.

The dataset comprises of the following range of agent organizations:

```
numberOfAgents = 3 TO 9 STEP 1 {
  agentSearchRange = 0 TO 50 STEP 10 {
    agentRemoveRange = 0 TO 50 STEP 10 {
          agentCommunicationsRange = 0 TO 50 STEP 10 {
                                   … }}}}
```

A configuration is maintained as long as the agent combination meets the validity criteria outlined in Table II. All other environmental conditions are kept constant; the unbounded environment size is 100 unit^2 and the number of targets is three.

Fig. 2. shows the normalized time taken to remove all targets, defined in (6) against robustness; the change in normalized time to remove targets when an agent is removed from an organization. The data points are colored using a standard color gradient to represent a third normalized dimension – structure. This can be changed to represent the other metrics described in Section 3. In this case, we look at the relationship between agent density and performance.

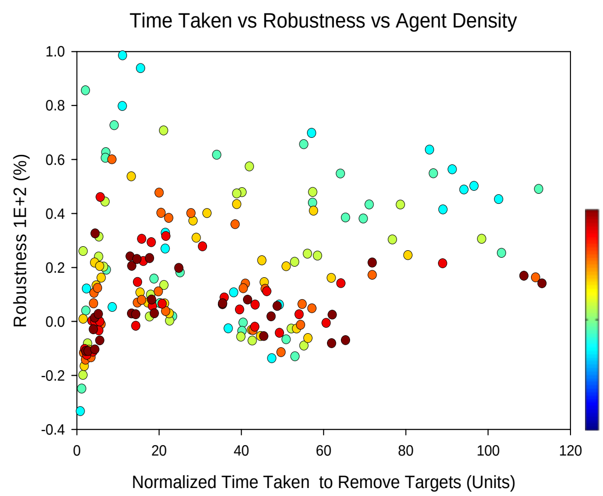

Fig. 2. Normalized time taken to remove all targets for a given organization against the change in normalized time taken after one agent is removed from the homogeneous organization. For the sake of clarity this chart and subsequent charts represent a close up of the data, and do not include all the outlying data points.

The optimal point for an organization where cost efficiency, time to remove targets and impact of failure are all equally important would be the point closest to the origin. However, we find that some points have "negative" robustness. These points represent organizations with significant levels of redundancy. Since the time taken is normalized over the attributes of the agent organization, a "negative" time taken indicates that similar performance was achieved using fewer resources.

Note that the magnitude does not represent the extent of redundancy, but indicates that it is present.

We also observe that organizations which contain many agents have a reduced probability of suffering significant setbacks when an agent is removed. This is to be expected in homogeneous organizations. We also find that the removal attribute is the most important when determining the time taken. Changing the cost weightings to discourage the use of the removal capability will change the distribution of the data points, as shown in Fig. 3.

To implement the cost analysis, capabilities that are outside two standard deviations of the mean are removed. This does not necessarily compromise the analysis as the tail end of the data, which exhibits a Poisson distribution, corresponds to sub optimal solutions for an organization. This approach is also useful when visualizing the data, as values far exceeding the standard deviation will place a significant bias on the color gradient.

We have related organizational composition (number of agents, capability distribution) to performance for homogeneous organizations that have uniform specialization, and negligible centrality and degree hierarchies. As our aim is to relate performance to structure which is a function of the composition, we now consider heterogeneous organizations and appropriate visualization methods of the complex multidimensional relationships between structure and performance.

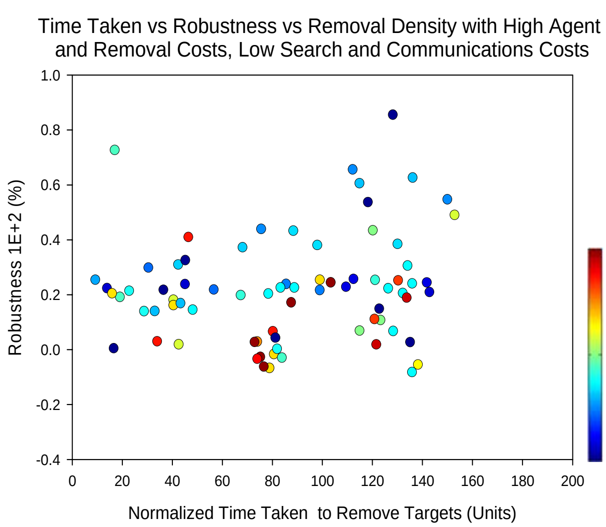

Fig. 3. Normalized time taken to remove all targets for the same organizations in Fig. 2., where the cost of removal and number of agents is increased by an order of magnitude, and the cost of communication and search capabilities are reduced by the same amount. We see that organizations around the optimal point tend to exhibit with low removal densities.

6. Multidimensional data visualization of heterogeneous organizations

To explore the behavior of heterogeneous organizations, we examine a single premise in detail; namely, the materiel cost of the organization is kept constant with three agents and 37 unit^2 of capabilities that are distributed in every possible combination to the three agents (i.e. the bracketed term in (6) is constant although this changes for $\vec{\tau}'$). The environmental conditions are the same as in Section 5. In all, over 6,000 scenarios were run. The resulting dataset is presented in Fig. 4.

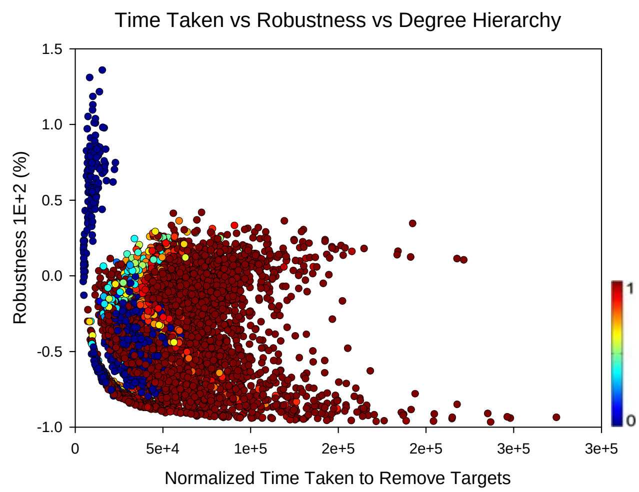

Fig. 4. Normalized time taken to remove all targets for a heterogeneous organization with fixed materiel.

The curve (Pareto frontier [26]) around the axis represents the trade off between time taken and robustness. Depending on requirements, an optimal point along that curve can be determined. We can also appreciate the relationship between degree hierarchy and performance. In organizations where there is no communication or perfectly reciprocal communication, so $H_D = 0$, there tends to be a high impact of failure, although there is a region where $H_D = 0$ organizations cluster around the optimal. To explore further the reasons for this we need to examine the possible correlations between the different structural parameters.

The high-dimensional measurement space that results from the use of numerous metrics requires the application of powerful data visualization tools. One such tool is the Self-Organizing Map (SOM), which SOM permits an ordered two-dimensional representation of much higher dimensional input spaces [27]. In essence, a SOM employs unsupervised clustering techniques that can reveal meaningful groupings of parameters that may form distinct classes of organizations. A SOM for the data generated from this scenario is shown in Fig. 5.

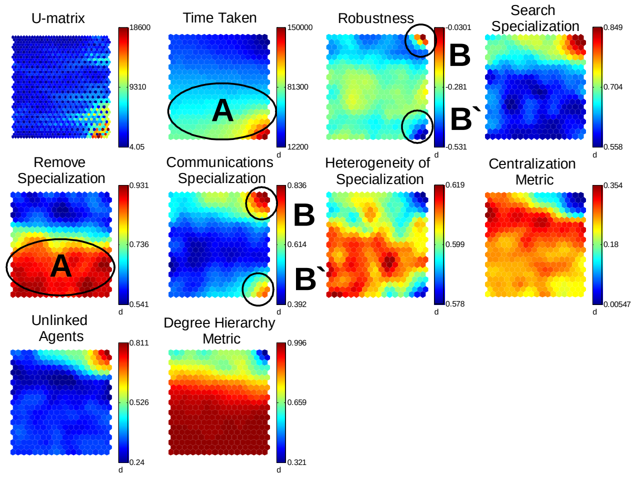

Fig. 5. SOM showing the U-matrix, the two performance outputs (Time taken, and Robustness) and structural metrics that describe the organizations.

Fig. 5 shows patterns in performance and structure that allow us to make inferences concerning organizational types. For example, in region A we see that as the remove specialization increases, the time taken also increases. This is because the removal skill is put to better utilization when distributed amongst all the agents. We also see that two types of high communications specialization lead to both poor robustness performance (region B) and good performance (region B'). Looking at the centralization metric area that corresponds to B, we see that the organization is behaving in a decentralized manner. This implies that the communications capability is not being utilized to its full potential. We also see that as heterogeneity of the organization increases, there is a general trend towards more centralized control. This implies that the organization has some form of coordination, and looking at the degree hierarchy, we note that this coordination tends to be mainly unidirectional, or top-down.

7. Conclusions and future work

Our principal aims are to attain a better understanding of emergent global behavior in organizations and consequently, improve their design, and the design process itself. These aims are addressed through the use of MAS simulations of organizations. Collated data comprises a set of generic structural metrics that are tabulated against a set of domain-specific performance metrics as detailed in Section 4.

In order to examine the resulting search space, various organizational configurations are simulated to build up a database of structure and performance relationships. We use a variety of visualization techniques, such as multi-dimensional plots and SOMs. Such methods, as we have demonstrated, provide a powerful means of visualizing the relationships between multi-dimensional parameters in a low-dimensional manner with the potential discovery of key organizational types.

Our research program will be to extend the simulation and analysis framework so that corresponding organizational models can be incorporated into a feedback learning system with more advanced cost-functions (incorporating procurement costs as well as operational ones) to find and maintain organizations at a required performance setting. Much of our work is directed at the aerospace industry with specific application domains such as Unmanned Air Vehicle (UAV) composition based on mission requirements and available inventory.

8. Acknowledgments

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

9. References

[1] V. Parunak, "Industrial Applications of Agents," presented at 2nd European Agent Summer School, Saarbrucken, 2000.

[2] J. Meek, "Things Falling Apart," in *The Guardian*. London, 2001, pp. 2 - 3.

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

[4] K. Kelly, *Out of control : the new biology of machines, social systems and the economic world*. Reading, Mass. ; Wokingham: Addison-Wesley, 1994.

[5] J. Cohen and I. Stewart, *The collapse of chaos : discovering simplicity in a complex world*. New York: Viking, 1994.

[6] R. Ashby, "Principles of the self-organizing dynamic system," *Journal of General Psychology*, vol. 37, pp. 125 - 128, 1947.

[7] S. A. Kauffman, *The origins of order : self-organization and selection in evolution*. New York ; Oxford: Oxford University Press, 1993.

[8] I. Prigogine and I. Stengers, *Order out of chaos : man's new dialogue with nature*. London: Fontana Paperbacks, 1985.

[9] M. C. Yovits, G. T. Jacobi, and G. D. Goldstein, *Self-organizing systems*. Washington, D.C.: Spartan Books, 1962.

[10] J. M. E. Gabbai, W. A. Wright, and N. M. Allinson, "Relating Organisational Structure to Performance: An Initial Look at Centralisation," presented at NET.Object Day, Erfurt, Germany, 2002.

[11] R. Evans, "Hard Choices," in *Aerospace International*, May 2001.

[12] M. Wooldridge, *Multiagent systems: a modern approach to distributed artificial intelligence*. Cambridge, Mass. ; London: MIT Press, 1999.

[13] M. Heusse, D. Snyers, S. Guérin, and P. Kuntz, "Adaptive Agent-Driven Routing and Load Balancing in Communication Network'," presented at Ants '98, Brussels, 1998.

[14] L. Gasser, "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.

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

[16] R. Lewin, *Complexity: Life on the Edge of Chaos*. London: Phoenix Paperbacks, 1995.

[17] J. M. E. Gabbai, W. A. Wright, and N. M. Allinson, "Centralised and Distributed Organisational Control," in *Intelligent Data Engineering and Automated Learning*, H. Yin, N. M. Allinson, R. Freeman, J. Keane, and S. Hubbard, Eds. Manchester: Springer, 2002, pp. 279-284.

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

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

[20] S. Wasserman, and Faust, K., *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press, 1994.

[21] H. A. Simon, *The New Science of Management Decision*. N.J.: Prentice Hall, 1977.

[22] M. Ahuja and K. Carley, "Network Structure in Virtual Organizations," *Organization Science*, pp. 741-747, 1999.

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

[24] T. F. Chan, G. H. Golub, and R. J. LeVeque, "Algorithms for computing the sample variance: analysis and recommendations," *The American Statistician*, vol. 37, pp. 242-247, 1983.

[25] J. M. E. Gabbai, N. M. Allinson, and W. A. Wright, "Relating Organisational Structure to Performance," in *e-Logistics for a Fail-Safe World: The Comprehensive Handbook of Holonic Enterprises*, M. Ulieru and A. W. Colombo, Eds. Berlin: Springer Verlag Heidelberg, 2003, pp. TBA.

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

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