Proceedings of the 5th WSEAS Int. Conf. on SIMULATION, MODELING AND OPTIMIZATION, Corfu, Greece, August 17-19, 2005 (pp319-324)

Towards an Agent-based Model of Adaptive Organisations

Oswaldo Terán[1], Juan de Lara[2], Manuel Alfonseca[2]
[1]Centro de Simulación y Modelos (CESIMO)
Universidad de Los Andes
Mérida, Venezuela
[2]Escuela Politécnica Superior
Universidad Autónoma de Madrid
Madrid, Spain

Abstract: - This work describes an agent-based model of an organisation. The model is made of workers, which are assigned tasks that have to be solved by analyzing some information items. Information quality has been modelled by associating to each item a probability of being wrong. Workers can interact with each other to recommend information items. During the simulations, we have induced deep reorganisation by changing the quality of the information items, inducing strong structural change. We have experimented with different information seeking behaviour for the workers and analyzed organisation performance, group formation and structural change in the periods of strong change.

Key-Words: - **Agent-based modelling and simulation, organisational modelling, self-organisation, structural change.**

1 Introduction

Computer simulation is an inexpensive means to perform virtual experiments that would be difficult, non-ethical, expensive or impossible to perform in reality. Many disciplines – such as physics, mathematics, computer science, biology and economics – use simulation techniques as a means for research or decision making. With the emergence of agent-based simulation techniques [10], simulation is also becoming popular in the social sciences. In agent-based modelling and simulation, the fundamental element is the agent [11], which can be defined as a "*computer system, situated in an environment, which is able to perform flexible and autonomous actions to achieve its design objectives*". In this way, the simulation is performed by the interactions of the agents, which give rise to "*macroscopic*" or emergent phenomena [1].

Organisational modelling [6][8][9][12][16][17][18][19] focuses on the decision-making process and behaviour of the individuals and the overall organisation. The former is essential for modelling and analysing organisational systems [7][13]. Recently, agent-based approaches have emerged, such as Carley's models for understanding structural change and learning in organisations. One of her models [4][5] was the starting point of the model we show in this paper.

In this paper, we describe our agent-based model of an organisation. We have modelled two kinds of interactions, inter-individual and individual-knowledge. Thus, conceptually, two different nets appear during the simulation: one relating workers (the *who-who* network), and the other relating workers with knowledge (the *who-what* network). Both networks evolve due to three reasons: by the interest of the workers in obtaining new knowledge, by affinity (workers that use similar knowledge) and by learning, depending on organisational performance. This latter behaviour is one of the differences of our model with Carley's [4][5]. The concept of endorsement [6][14] is used by the workers in order to classify other workers and knowledge items according to past performance.

As the workers learn to differentiate the best knowledge items, they modify both networks. We call this a period of "*soft*" structural change. After some time, the organisation achieves some degree of stability: self-organisation has emerged. At some points in time, when self-organisation is achieved, we induce a strong change in the organisation by modifying the quality of the knowledge items. Our aim is to analyse the two kinds of structural change: soft and strong. For this purpose, we measure performance, dynamic group formation and structural change. Group formation is examined, as in [5], in its simplest forms, groups of three collaborating workers (*triads*).

The rest of the paper is organized as follows. Section 2 presents the basic organization model. Section 3 briefly comments the implementation in the SDML language. Section 4 discusses the experiments performed and the results obtained. Finally, section 5 ends with the conclusions and future work.

2 The Organisation Model

Our organizational model is made of workers, which are modelled as agents. The organisation goal is to classify a problem in one of two different classes. The problem consists of a number of Boolean information items. If the majority of the items is "1", then the problem is classified as "1", and "0" otherwise. The classification work is partitioned, in such a way that each agent in the organisation is assigned a number of items. Thus, agents do not have a complete knowledge of all the items in the problem. Agents count the "1" and "0" and issue an answer. The organisation answer is then calculated as the majority of "1" or "0" answers. This is a classical problem in organisation modelling [4][5], and similar problems of distributed observation are found in control theory[20]. Figure 1 shows a scheme of the model.

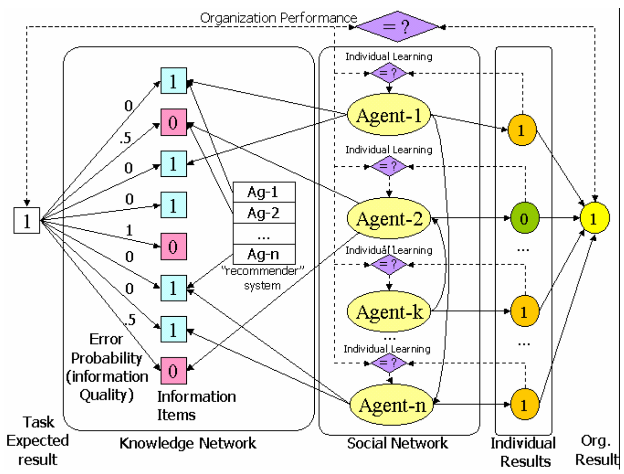

Figure 1: Scheme of the Organisation Model

The simulation is organized in days, and each day 40 problems are presented. A problem is made by first selecting the correct answer (0 or 1), and then deriving a binary vector that represents it. Each element in the vector is assigned a probability (1, 0 or 0.5) of being wrong. Thus, if it is 0, the item is equal to the correct answer. If it is 1, the item is the complement. Finally, if it is 0.5, then the value is equal to the correct answer with a probability of 0.5.

Each agent has a limited memory, where he can store some information items. Agents interact, and in each interaction an agent can take a new information item (and discard an old one). We have modelled three different types of behaviour for the agents. This behaviour governs how an agent interacts with other agents. When solving a problem, the agents choose probabilistically between one of them. The first one is called *passive*. Here we model agents that communicate with agents that are similar to them, in the sense that they use a similar set of information items. The second one is called *active*. In this mode, agents interact with others that have different information. Finally, in the *learning* mode, agents interact with other agents that gave them good items in the past. Agents classify other agents as well as knowledge items using *endorsements* [6][14]. This concept is implemented by the following formula:

$$E(b,\bar{a}) = \sum_{val(\bar{a}_i)\geq 0} b^{val(\bar{a}_i)} - \sum_{val(\bar{a}_i)<0} b^{|val(\bar{a}_i)|}$$

Where **b** is the chosen basis (we use 1), **a** is a vector of attributes of the object to be endorsed and **val** is a function giving the value of that attribute.

In addition, when agents are *active*, there is a small probability for an agent to choose an item directly from the knowledge network, if it is not being used by other agents. This will be useful to "discover" good items when we change the quality of the items. Active agents can also take items from a "recommender system", which contains the best items found so far by each agent.

When an agent gives his answer, he can compare it with the correct answer. Thus, if his answer is correct, it is counted in the first sum as +1, while if it is incorrect, then it is counted in the second sum as – 1. Agents only remember the four last days (but there are 40 problems each day).

3 Implementation

For our implementation, we have used the SDML language (a Strictly Declarative Simulation Language) [15], based on the KD45 modal logic. The language has a number of predefined agent types, which have to be sub-classified to implement the agents for the problem. The language is declarative, agent actions are specified by means of rules. Agents may be nested, and at each level of nesting it is possible to define a number of time levels: the most external one is called "*eternity*". *Initial* and *final* rules can be associated with the different time levels. For instance, in our organisation model, we have defined the *day* time level, and inside it the *problem cycle*. Each day is made of 40 problem cycles, meaning that 40 problems should be solved each day. A complete simulation run lasts for 500 days. Figure 2 shows the database of one agent during one of the simulations.

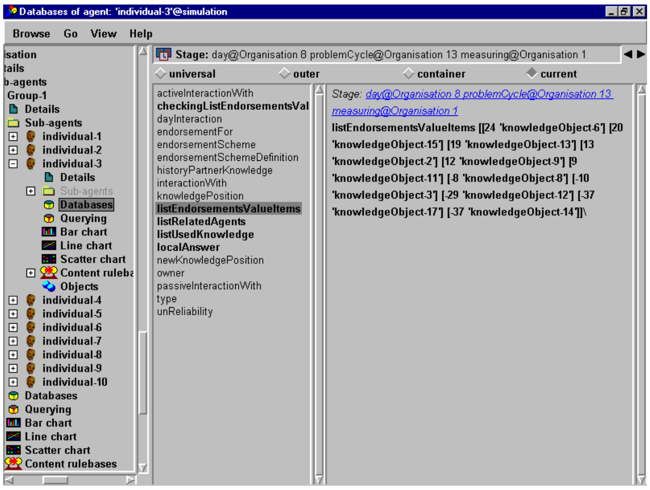

Figure 2: Database of one agent during the simulation

4 Experiments

Here we describe a number of experiments with the organisation model. First, we experimented with the agents' behaviour, by making the probability of choosing one of the three behaviours equally biased, making them always *passive*, always *active* or always *learning*. We also experimented with different memory sizes (3, 7 and 10). The reliability of the input items was changed during the simulation runs (every 100 days), as this paper's objective was to examine the dynamics of the simulated organisation when the quality of the items changes. The problem is composed of 20 items, to 7 of which we assigned a 0 probability of being wrong, to another 7 a probability of 0.5 and to the rest a probability of 1.

Simulations were run for 500 days. The reliability values were randomly changed on simulation days 100, 200, 300 and 400. As explained above, the individuals solve 40 problems every day, one per problem cycle. Running the simulation for one hundred days before modifying the reliability of the information items, guarantees that agent learning, group formation and performance stabilises. For the case of equally biased individuals, the simulation was run 4 times. Thus, for this case, 4 * 5 = 20 simulation periods of soft structural change, and 4 * 4 = 16 transition periods from soft structural change to strong structural change, were analysed. In average, each simulation took about two days to finish in a Pentium IV computer.

In the next subsections we analyse some of the measured properties for the type of experiments performed.

4.1. Organisation Performance

Organisational performance measures the percentage of problem cycles where the organisation gave the correct answer. We describe the seven experiments that were performed:

a) *Base model.* The different behaviour modes are equiprobable. Agents were assigned a memory of 7 items. The results of one experiment are shown in Figure 3. The time to achieve maximal performance is about 20 days. After the reliability of the input items changes (every 100 days), the length of the transitory period reduces to less than half. This is because, at the beginning of the simulation, the agents' mental model is empty – it contains only 2 information items without endorsements. On the other hand, after a reliability change, the agent knows many items, although they may be wrongly endorsed after the change, and the agents need to update them. Just after a strong structural change, the performance decreases to around 0.5, i.e. the organisational performance becomes random. In this model, once the stable period is reached, the randomness of the environment input items has little impact on the performance.

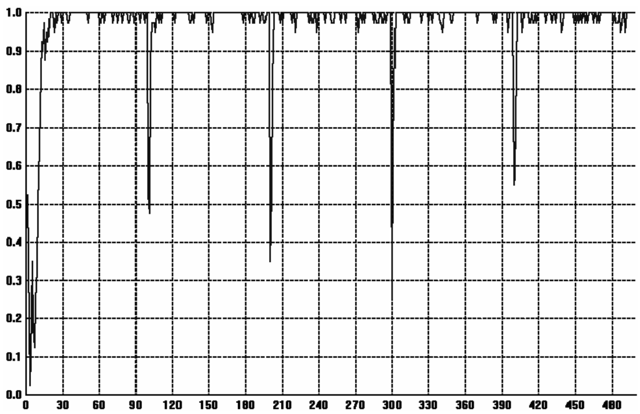

Figure 3: Performance for the Base Model.

b) *Active model.* The agents are always active. Compared with the base case (a), a faster learning is observed at the beginning of the simulation.

c) *Short memory model*. Identical to the base model, but agents have a memory size of 3. Results are quite similar to the base model, but after

Proceedings of the 5th WSEAS Int. Conf. on SIMULATION, MODELING AND OPTIMIZATION, Corfu, Greece, August 17-19, 2005 (pp319-324)

changing the reliability of the information items, the performance of this model is affected more strongly than in the base model.

d) *No recommendation model*. Similar to the base model, but without recommendation system. Results are also quite similar to the base model.

e) *Passive model*. The agents are always passive. In this case, after some time, interaction stops, as the individuals exhaust their possibilities for information exchange. This confirms Carley's results, but in our case learning and information diffusion seem to be much faster. After a strong structural change is induced, good learning is a question of chance, and behaviour becomes good or bad in general, without improving beyond a certain upper bound, different for each period of soft structural change (see Figure 4).

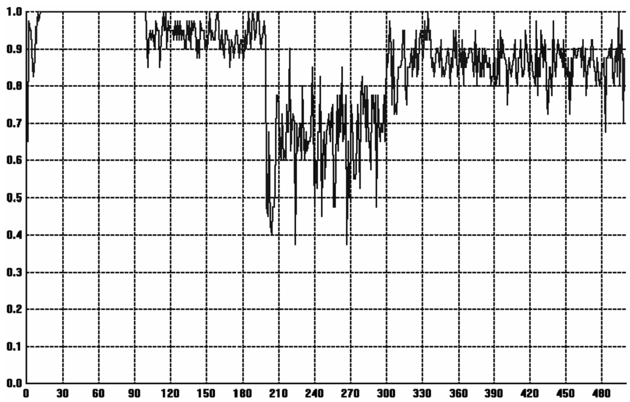

Figure 4: Performance for the Passive Model.

f) *Learning model.* The agents are always learning. Stabilization takes longer and performance oscillates in a wider interval (which only occasionally reaches the upper bound of 1), than to the base model. In this case, the randomness in the environment continues affecting the organization's performance over time.

g) *No-noise model*. Similar to base model, but without noise (information currently unused by other agents). The presence of noise means that the agents can choose items from the environment. In this model, the performance improves until it oscillates in a certain range (see Figure 5), which only occasionally reaches the value of 1. In general, randomness in the environment continues affecting the organisational behaviour.

As a summary, active learning and noise from the environment seem to be the main sources of the differences between the seven experiments. Results from the last three models are quite different from the first four. In the first four experiments, agents are either active or can obtain noise information from the knowledge base (currently unused by other agents). This ability seems essential to recover the performance after a strong structural change. Interestingly, individual learning does not seem to help the organisation performance. Moreover, the recommender system does not seem to improve the performance, as long as individuals are able to interact in an active way.

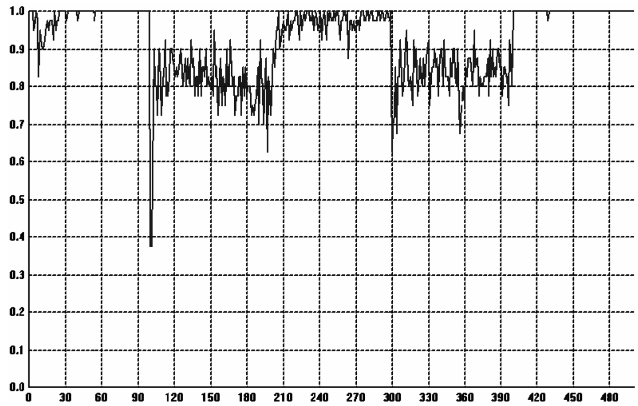

Figure 5: Performance for the No-noise model.

4.2. Dynamic Group Formation: Triads.

Here we present the results of analysing dynamic group formation during the simulation. As in [5], we analyse the simplest form: groups of tree or *triads*. Again we comment the results for the seven experiments:

a) *Base model*. Stabilisation of triads takes longer than performance stabilisation, about 40 days (compared to 20 days for performance). These results agree with Carley's model [5], although in our case both stabilisations are faster (see Figure 6).

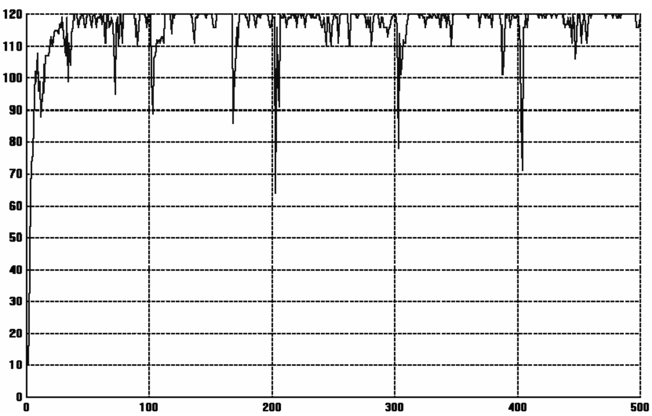

Figure 6: Triads for the base model.

b) *Active model*. Similar to the base model, but the number of triads is much more variable.

c) *Short memory model*. Triads are much more unstable than in the two previous models. They rarely reach the maximum of 120, sometimes going down to values such as 70, during the period of soft structural change, which never happens in the base model. This might happen because individuals choose more dynamically new individuals and forget the old ones, which favours fast learning.

d) *No recommendation model*. Stabilisation takes about 60 days, much longer than the performance stabilisation, which takes about 30 days.

e) *Passive model*. Once the interaction of agents stops, the number of triads achieves a fixed value, which is different for every simulation experiment. In the case shown in Figure 7, this value happened to be 20.

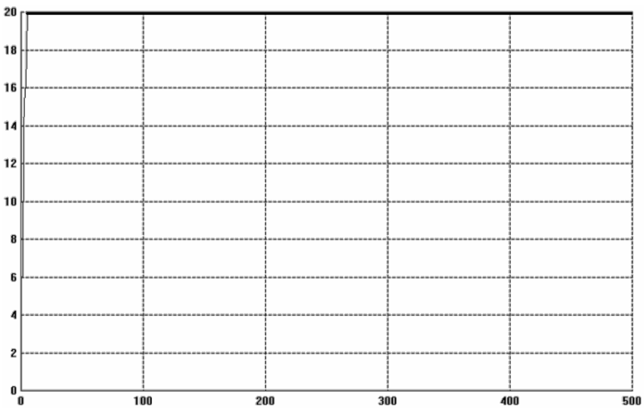

Figure 7: Triads for the passive model

f) *Learning model*. Here the maximal number of triads (120) is never reached, but groups are more stable. This may happen because the individuals keep their endorsements of other individuals for several days, thus slowing down the group change.

g) *No-noise model*. Results are similar to the base model.

As a summary, organisational structure usually takes longer to stabilise than performance. The exceptions are the passive and no-noise model. In addition, learning behaviour in agents seems to produce more stable groups, because of the selection mechanisms of the agents with which the interactions take place.

4.3. Structural Change: Organisational Behaviour and Reorganisation.

Here we study structural change, i.e. organisational behaviour within a period with fixed reliability of the information items, and re-organisation after a variation of the reliability of these items. Structural change is analysed by examining the endorsements assigned to the items, depending on their reliability. When agents learn, they will give higher endorsement value to those with reliability 1, less to those with reliability 0.5, and the lowest to those with reliability 0. Once the reliability of the items changes, reorganisation must happen, as the agents should adapt the endorsement to the new environment (i.e. the *who-what* network should be heavily changed).

We have studied the differences between the total endorsements of items with two different reliabilities. For example, Figure 8 shows this difference for items with reliabilities one and zero. It can be seen how agents quickly discard bad items (as the quantity becomes larger) in the period after the strong structural change.

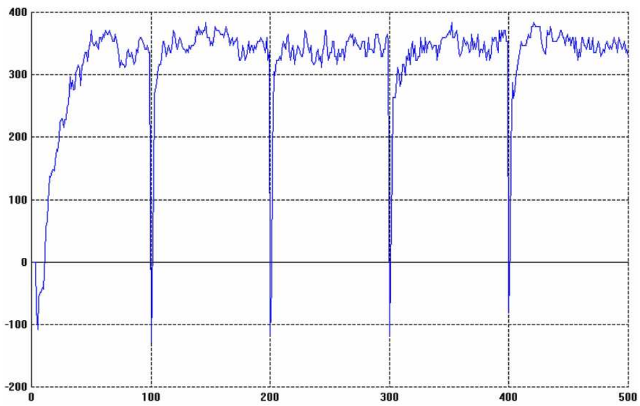

Figure 8: Relation between used items with reliability 1 and 0 for the base model.

However, the agents seem confused with items of reliability 0.5, as expected, even in the case of learning agents. These also have an upper bound for learning (sometimes quite poor). This is partly due to the fact that learning agents do not take items from the environment nor from the recommendation system. Despite these limitations, learning individuals behave much better than passive ones. This suggests that, for modelling learning agents, it is important to consider indirect interactions, such as those provided by a recommendation system, and the interaction with the organisational environment, such as our notion of noise. The simulated random influence from the environment helps the organisation to avoid being trapped in an old and obsolete culture.

5 Related Work

This work is an improvement of Carley's model CONSTRUCT-O [5]. That model is an improvement on previous ones [2][3]. In CONSTRUCT-O, the agent interaction style is restricted to be passive or active, where each agent can be assigned a mixed strategy. No strong changes were induced nor indirect communication modelled.

6 Conclusions and Future Work

In the present paper, we have continued and improved Carley's model by adding a new interaction style (learning) to the agents, using the concept of endorsement, and inducing strong structural changes to the organisation. We have added a probability of being wrong to each problem item (modelling the quality of the information). After the stabilisation of the organisation (soft structural change) we induce a strong change by modifying these probabilities. Our results for the periods of soft change agree with Carley's model, although the stabilisation in our models is much faster.

We have also experimented with other knowledge store and diffusion mechanisms, such as the "recommender system". Some interesting results were found. For example, individual learning (as described here) is not essential for organisation performance, while active behaviour is. Active behaviour leads to learning at a higher level: "structural learning" within an organisation. That is, changes in groups of agent collaborators. Working groups may change without damping performance, and they are more stable with learning agents. Being able to access unused knowledge (noise) becomes essential when drastic changes are made to the problem assumptions. Organisation with passive agents cannot react to this kind of changes.

In the future, we plan to extend the model by considering more complex knowledge structures, similar to the Internet. Other organisation structures and hierarchies, mental models for the agents and more complex problems could also be considered.

References:

[1]Alfonseca, M., de Lara, J. 2002. "*Two level evolution of foraging agent communities*". BioSystems(66), Issues 1-2, pp.: 21-30. Elsevier.

[2] Carley, K 1992. "*Organizational Learning and Personnel Turnover*".Org.Science,3(1): 20-46.

[3] Carley, K. and Svoboda, S. M. 1996. "*Modeling Organizational Adaptation as a Simulated Annealing Process*." Sociological Methods and Research 25(1): 138-168.

[4] Carley, K. 2001. "*Intra-organizational Complexity and Computation*". in Companion to Organizations, Oxford UK: Blackwell.

[5] Carley, K., Hill, V. 2001, "*Structural Change and Learning Within Organizations*." In Dynamics of Organizations: Computational Modeling and Organizational theories. Edited by Alessandro Lomi and Erik R. Larsen, MIT Press/AAAI Press/Live Oak, Ch. 2.pp 63-92.

[6] Cohen, P. 1985. "*Heuristic Reasoning: An Artificial Intelligence Approach*". Pitman Advanced Publishing Program, Boston.

[7] Cohen, M., March, J. and Olsen J. 1972 "*A Garbage Can Model of Organisational Choice*". Administrative Science Quarterly, 17(1), pp. 1-25.

[8] Cyert R., March J. 1963. "*A Behavioural Theory of the Firm*". Englewood Cliffs, Prentice Hall, NJ.

[9] Cyert R., March J. 1959. "*A Behavioural Theory of Organisational Objectives*". In Haire M. (ed.), Modern Organization Theory (pp. 76-90), Wiley New York.

[10] Gilbert, N., Troitzsch, K. 1999. "*Simulation for the Social Scientist*". Open University Press.

[11] Jennings, N. R., Sycara, K., Wooldridge, M. 1998. "*A Roadmap of Agent Research and Development*". Autonomous Agents and Multi-Agent Systems, 1, pp.: 7-38. Kluwer.

[12] March, J., Simon H. 1958, "*Organisations*". New York, Wiley.

[13] Masuch M and LaPotin. 1989. "*Beyond Garbage Cans: An AI Model of Organizational Choice*". Administrative Sciences Quarterly, 34, 38-67.

[14] Moss, S. 1995. "*Control Metaphors in the Modelling of Decision-Making Behavior*", Computational Economics, v.8, pp. 283-301

[15] Moss S., Gaylard H., Wallis S., and Edmonds B. 1998. "*SDML: A Multi-Agent Language for Organizational Modelling*", Comp. Mathematical Organization Theory, 4(1), pp. 43-69.

[16] Newell, A. 1990. "*Unified Theory of Cognition*". Harvard University Press, MA.

[17] Newell A, Simon H. 1976. "*Computer Science as Empirical Enquiry: Symbol and Search*". Communications of the ACM, 19(3).

[18] Simon, H. 1984. "*The Sciences of the Artificial*". MIT Press, Cambridge, Mass., USA.

[19] Simon, H. 1948. "*The Proverbs of Administratio*". Public Administration Review, 6(June), pp. 53-57.

[20] Tripakis, S. 2005. "*Two-phase distributed observation problems*". Proc. IEEE Application of Concurrency to System Design (ACSD). 98-105.