

Exploring Robustness in the Context of Organizational Self-Design*

Sachin Kamboj and Keith S. Decker

Department of Computer and Information Sciences

University of Delaware

Newark, DE 19716, USA

{kamboj, decker}@cis.udel.edu

Abstract

Robustness is the ability of a multiagent system to recover from failures and exceptions. In particular, the system should be able to recover from task and agent failures and the failure of any single agent or group of agents should allow the graceful, predictable degradation of performance. Hence, designing robust systems is a critical challenge facing many multiagent designers.

Organizational Self-Design (OSD) has been proposed as an approach to constructing suitable organizations at runtime in which the agents are responsible for constructing their own organizational structures. OSD has also been shown to be especially suited for environments that are dynamic and semi-dynamic. However, the problem of making these self-designed organizations robust is still an open research problem that has not been studied to any considerable extent. In this paper, we focus on developing and evaluating robustness mechanisms that can be used by the agents in conjunction with OSD.

Introduction

Robustness may be defined as the ability of a multiagent system to recover from failures and exceptions. An exception may be defined as a departure from an “ideal” system behavior (Dellarocas & Klein 2000). Recovery would then involve the execution of some corrective measures to reinstate the ideal system behavior.

Achieving robustness is particularly challenging in dynamic and semi-dynamic environments, since the problem characteristics, available resources or agent capabilities may change over time. Multiagent organizations for such environments must include two components — the first component is responsible for monitoring the performance of the organization and for discerning whether or not the measured performance falls within the design parameters of the organization. The second component is responsible for explicitly changing the organization if it fails to meet its design goals.

This monitoring and design may be done by an entity external to the organization (i.e. by the multiagent designer) or by the component agents themselves. In an extreme form of the latter approach, the agents come up with a new, implicit, one-off organization for each problem instance. This is what happens in the contract net protocol (CNP) (Smith 1988).

*This paper is a work in progress.

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This latter approach is referred to in the literature as organizational self-design (OSD) (Ishida, Gasser, & Yokoo 1992; Kamboj & Decker 2007) and in this approach the agents are responsible for designing their own organizational structure at run-time. OSD is especially suited in situations where the environment is semi-dynamic as the agents can adapt to changes in the task structures and environmental conditions, while still being able to generate relatively stable organizational structures that exploit the common characteristics across problem instances. However, this amortization of organizational costs across problem instances comes at a price — such generated organizations can be particularly susceptible to agent failures. Since the constituent agents are responsible for enacting particular goals within the organization, the failure of a single agent, performing a critical role, could bring down the whole organization. Hence, such organizations might be *less* robust than the one-off organizational schemes (such as CNP).

In this paper we would like to study various approaches that can be used to increase the robustness of organizations generated through the use of OSD. Furthermore, we primarily focus on agent failures¹ in worth-oriented domains. Our approach is geared towards real-world applications in grid/volunteer/cloud computing where a large pool of agents is available though the availability of any single agent cannot be guaranteed.

A primary challenge to incorporating robustness in the OSD process is the continuously changing set of agents and the roles that they enact. This leads to a continuously changing (and distributed) organizational knowledge that must be preserved across agent failures. We feel that robustness has not been studied in the context of OSD to any significant extent. Whereas we (Kamboj & Decker 2007) alluded to this problem in our previous paper, we did not present any algorithms and did not discuss the problem in any depth. The algorithm presented in (DeLoach, Oyenan, & Matson 2008) would respond to a single agent failure by performing a complete reorganization — an extremely expensive process.

¹Another aspect to robustness in multiagent systems is task failures. We won't be concerned with task failures because — (a) agent failure is significantly harder than task failure because the failure of an agent results in a loss in both its organizational knowledge and its contextual problem-solving state; and (b) our approach can easily handle task failures through rescheduling (Wagner, Raja, & Lesser 2006).

(Sims, Goldman, & Lesser 2003) also address robustness by reorganizing, however their approach is specific to the distributed sensor network and would have to be adapted for general purpose worth oriented domains. Other approaches to OSD (Ishida, Gasser, & Yokoo 1992; So & Durfee 1996; Frederic & Jacqueline 1995) tend to completely skim over the problem.

This paper will present and evaluate algorithms for both of the two commonly used approaches to robustness:

1. the *Survivalist Approach* (Marin *et al.* 2001), which involves replicating domain agents in order to allow the replicas to take over should the original agents fail; and
2. the *Citizen Approach* (Dellarocas & Klein 2000; Klein, Rodriguez-Aguilar, & Dellarocas 2003), which involves the use of special monitoring agents (called *Sentinel Agents*) in order to detect agent failure and dynamically startup new agents in lieu of the failed ones.

Our goal is to allow the organization, at its best, to function without any performance degradation in the face of failures. At its worse, the organization should degrade gracefully in proportion to the number of failures.

Note that we do *not* present any new approach to OSD in this paper. Instead we add robustness to the OSD approach presented in (Kamboj & Decker 2007). Also, we are *not* trying to develop any new and general approaches to robustness. Instead we were trying to address the robustness issues that arise when using OSD.

The organization of the rest of this paper is as follows. In the next section we discuss the task model that we use for worth-oriented domains. This is followed by a discussion of our approach to OSD. Finally, we evaluate the presented algorithms.

Task Model

We use TÆMS as the underlying representation for our tasks (problem instances). TÆMS (Lesser *et al.* 2004) (Task Analysis, Environment Modeling and Simulation) is a computational framework for representing and reasoning about complex task environments in which tasks (problems) are represented using extended hierarchical task structures (Chen & Decker 2005). The root node of the task structure represents the high-level goal that the agent is trying to achieve. The sub-nodes of a node represent the sub-tasks and methods that make up the high-level task. The leaf nodes are at the lowest level of abstraction and represent executable methods – the primitive actions that the agents can perform. The executable methods, themselves, may have multiple outcomes, with different probabilities and different characteristics such as quality, cost and duration. TÆMS also allows various mechanisms for specifying subtask variations and alternatives, i.e. each node in TÆMS is labeled with a characteristic accumulation function that describes how many or which subgoals or sets of subgoals need to be achieved in order to achieve a particular higher-level goal. TÆMS has been used to model many different problem-solving environments including distributed sensor networks, information gathering, hospital scheduling, EMS,

and military planning. (Decker 1995; Decker & Li 2000; Chen & Decker 2005; Zimmerman *et al.* 2007).

Formally, the primary input to the multi-agent system (MAS) is an ordered set of problem solving requests or task instances, $\langle P_1, P_2, P_3, \dots, P_n \rangle$, where each problem solving request, P_i , can be represented using the tuple $\langle t_i, a_i, d_i \rangle$. In this scheme, t_i is the underlying TÆMS task structure, $a_i \in \mathbf{N}^+$ is the arrival time and $d_i \in \mathbf{N}^+$ is the deadline of the i^{th} task instance². The task t_i is not “seen” by the MAS before the time a_i , i.e., the MAS has no prior knowledge about the task t_i before the arrival time, a_i . In order for the MAS to accrue quality, the task t_i must be completed before the deadline, d_i .

Furthermore, every underlying task structure, t_i , can be represented using the tuple $\langle T, \tau, M, Q, E, R, \rho, C \rangle$, where: T is the set of tasks; $\tau \in T$, is the root of the task structure; Q is the set of quality/characteristic accumulation functions (CAFs); E is the set of (non-local) effects (NLEs) and R is the set of resources. For a more detailed formal definition of each of these, refer to (Kamboj & Decker 2007; Decker 1995)

Organizational Self Design

In our approach, the organizational design is directly contingent on the task structure of the problems being solved and the environmental conditions under which the problems need to be solved. Here, the environmental conditions refer to such attributes as the task arrival rate, the task deadlines and the available resources. We assume that all problems have the same underlying task structure, henceforth called the global task structure.

To participate in the organization, each agent must maintain some organizational knowledge. This knowledge is also represented using TÆMS task structures, called the local task structures. These local task structures are obtained by rewriting the global task structure and represent the local task view of the agent vis-a-vis its role in the organization and its relationship to other agents. Hence, all reorganization involves rewriting of the global task structure. However, note that the global task structure is *NOT* stored in any one agent, i.e. no single agent has a global view of the complete organization. Instead each agent’s organizational knowledge is limited to the tasks that it must perform and the other agents that it must coordinate with — it is this information that is represented using the local task structures.

To allow for a change in an agent’s organizational knowledge, we define three rewriting operations on a task structure:

Breakup: Breakup involves dividing the local task structure of an agent so that it can be allocated to another agent. The algorithm for breakup is shown in Algo. 2 and an example is shown in Figure 1. In this example, the task structure represented by Root-1 is broken twice — once at node D to generate Root 2 and a third time at Node E to generate Root 3.

² \mathbf{N} is the set of natural numbers including zero and \mathbf{N}^+ is the set of positive natural numbers excluding zero.

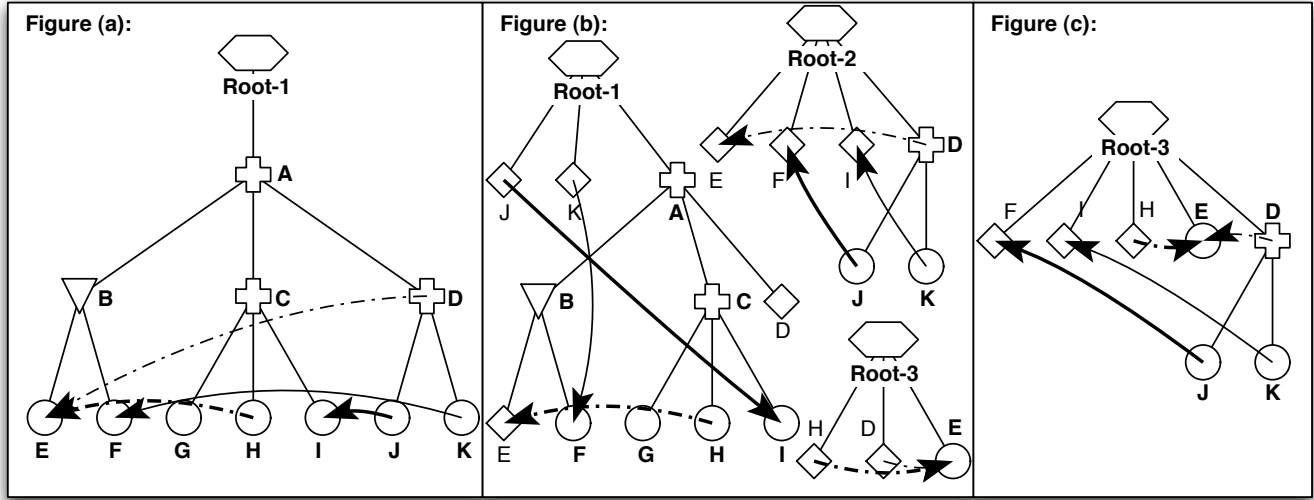


Figure 1: **Task rewriting:** *Fig (a)* shows the global TÆMS task structure: The polygons (labelled A – D) represent tasks and the circles (labelled E – K) represent executable methods. The + iconography indicates a *SUM CAF* while ∇ represents a *MIN CAF*. The arrows represent NLEs — the thick arrows represent hard constraints such as *Enables* (represented by a solid arrow from J to I) and *Disables* (represented by a broken arrow from H to E). The thin arrows show soft constraints such as *Facilitates* (solid arrow from K to F) and *Hinders* (broken arrow from D to E). Method characteristics and other details are omitted. *Fig (b)* shows the breakup of Root-1 at nodes D and E. The diamonds represent non-local nodes, that are the responsibility of some other agent. *Fig (c)* shows the merging of Root-3 and Root 2.

Merging: Merging involves combining two different local task structures from two different agents to form one local task structure. The algorithm for merging is shown in Algo. 3 and an example is shown in Figure 1.

Cloning: Cloning involves creating a complete copy of a substructure, so that it can be allocated to another agent. Cloning serves two purposes: (a) it can be used for load balancing similar to the work of (Shehory *et al.* 1998); and (b) it can be used to increase the *robustness capacity* of an agent by having multiple agents work on the same task simultaneously. The algorithm for cloning is given in Algo. 4 and an example is shown in Figure 2.

Our approach to OSD involves starting with a single agent responsible for the global task structure (i.e. the local task structure is equal to the global task structure). Each agent in the organization follows the algorithm presented in 1, if an agent is overloaded, it either breaks or clones its local task view and then spawns a new agent. If the agent is underloaded, on the other hand, it composes with another agent, merging the local task structures of the two organizations.

Algorithm 1 CHANGEORGANIZATION

- 1: **if** ISAGENTOVERLOADED() **then**
 - 2: SPAWNAGENT()
 - 3: **else if** ISAGENTUNDERLOADED() **then**
 - 4: COMPOSEAGENT()
-

Algorithm 2 BREAKUP (τ, v)

- 1: $\bar{\tau} \leftarrow \text{DESCENDENTS}(\tau) - \text{DESCENDENTS}(v)$
 - 2: $\bar{v} \leftarrow \text{DESCENDENTS}(v)$
 - 3: **for all** $\{ N \mid N \in \text{NLES}(\tau) \}$ **do**
 - 4: **if** (SOURCE(N) $\in \bar{\tau}$ **and** SINK(N) $\in \bar{v}$) **or**
 (SOURCE(N) $\in \bar{v}$ **and** SINK(N) $\in \bar{\tau}$) **then**
 - 5: $x \leftarrow \text{GETNONLOCALNODE}(\text{SOURCE}(N))$
 - 6: $y \leftarrow \text{GETNONLOCALNODE}(\text{SINK}(N))$
 - 7: $M \leftarrow \text{COPYNLE}(N)$
 - 8: REPLACENODE(N , SOURCE(N), x)
 - 9: REPLACENODE(M , SINK(N), y)
 - 10: $x \leftarrow \text{GETNONLOCALNODE}(v)$
 - 11: REPLACENODE(τ, v, x)
 - 12: **return** CREATEROOTNODE(v)
-

Robustness Mechanisms

Both of our robustness mechanisms involve three parts: (a) monitoring for agent failure; (b) maintaining state information for all the agents; and (c) restarting failed agents.

Furthermore, the underlying mechanism for monitoring and restarting is the same across the robustness mechanisms. Monitoring is achieved by sending out periodic *Are-You-Alive* messages to the set of monitored agents and waiting for *Alive* reply messages. If a reply is not received within a certain interval, we assume that the agent is dead and send a restart message to the environment. The individual mechanisms, however, differ in *who* is responsible for the monitoring and *which* set of agents are monitored.

State information is needed to restart a failed agent. At

Algorithm 3 MERGE (τ, v)

```
1: for all {  $x \mid x \in \text{DESCENDENTS}(v)$  } do
2:    $y \leftarrow \text{FINDNODE}(\tau, x)$ 
3:   if NULL( $y$ ) then
4:     DELETENODE( $v, x$ )
5:     ADDNODE( $\tau, x$ )
6:   else if TYPE( $\tau, y$ ) = NonLocal  $\wedge$  TYPE( $v, x$ )
       = NonLocal then
7:     MERGENODES( $\tau, y, x$ )
8:   else if TYPE( $\tau, y$ ) = Local  $\wedge$  TYPE( $v, x$ )
       = NonLocal then
9:     DELETENODE( $v, x$ )
10:  else if TYPE( $\tau, y$ ) = NonLocal  $\wedge$  TYPE( $v, x$ )
       = Local then
11:    REPLACENODE( $\tau, x, y$ )
12:  return  $\tau$ 
```

Algorithm 4 CLONE (τ, v)

```
1:  $\bar{\tau} \leftarrow \text{DESCENDENTS}(\tau)$ 
2:  $\bar{v} \leftarrow \text{DESCENDENTS}(v)$ 
3:  $\phi \leftarrow \text{CREATECLONECONTAINER}(v)$ 
4: for all {  $x \mid x \in \bar{v}$  } do
5:    $y \leftarrow \text{COPYNODE}(x)$ 
6:   ADDNODE( $\phi, y$ )
7: for all {  $N \mid N \in \text{NLES}(v)$  } do
8:   if SOURCE( $N$ )  $\in \bar{\tau}$  then
9:      $x \leftarrow \text{CREATEINHERITINGNODE}()$ 
10:    ADDNODE( $\tau, x$ )
11:     $L \leftarrow \text{COPYNLE}(N)$ 
12:     $M \leftarrow \text{COPYNLE}(N)$ 
13:    REPLACENODE( $N, \text{SINK}(N), x$ )
14:    REPLACENODE( $L, \text{SOURCE}(L), x$ )
15:    REPLACENODE( $M, \text{SOURCE}(M), x$ )
16:     $y \leftarrow \text{FINDNODE}(\phi, \text{SINK}(M))$ 
17:    REPLACENODE( $M, \text{SINK}(M), y$ )
18:   else if SINK( $N$ )  $\in \bar{\tau}$  then
19:     {Similar to the source}
20:  return  $\phi$ 
```

a minimum, this state information should contain the *organizational state* (i.e. the local task structure) of the agent being restarted. However, the local task information is not sufficient for restarting an agent in a complex domain. The restarted agent will still need information about the *execution context*, i.e. information about the outstanding task instances, information about the methods of a task instance that have already been executed (so that the agent does not try to re-execute them) and information about coordination commitments (because the subtasks have non-local effects and are interdependent on each other). The coordination mechanisms also differ in how they keep a track of this execution context.

Citizens Approach

The citizens approach involves creating a special monitoring agent (called a *sentinel agent*), which is responsible for all the robustness related responsibilities of the organization.




Figure 2: **Task Rewriting (cont.):** Figure showing the cloning of Node C in Root 4. The \sqcup node (Node C(C)) is used to represent a clone container, a node created for storing the “clones” of a node; the \perp node (Node S(C)) is used to select which clone to “execute” — for the purposes of robustness, all the clones will be executed; finally the \times nodes represent the NLE-inheriting-methods — these nodes form a conduit for NLEs to the non-clone parts of the task structure.

This approach is the simplest to execute — the sentinel agent is the sole monitor that is responsible for monitoring all the agents in the organization. However, to maintain state information, the sentinel needs to listen to *all* the messages exchanged by *all* the other agents, since it needs to store both the set of spawning/composition messages (in order to track the changes to the local task structures of the agents) and the set of execution/coordination messages between the agents (in order to keep a track of the execution context).

Hence, not only does the sentinel effectively becomes a conduit for all the messages, it also has global knowledge about the complete organization — a problem we were trying to avoid by using the OSD approach in the first place. Furthermore, the sentinel can (a) quickly become overwhelmed by all the messages that it needs to track and (b) become a central point of failure³. The solution might be to add multiple sentinel agents — we will now need to create an organization for the sentinels (for which we could, again, use OSD) and a way of monitoring the monitors.

Hence, we focus on developing algorithms for the survivalist approach and use the citizens approach for comparison.

Survivalist Approach

In the survivalist approach, there are no special agents responsible for monitoring and restarting failed agents. Instead the domain agents divide the monitoring responsibilities amongst themselves. Furthermore, some/all domain agents may be replicated in order to (a) increase the robustness capacity of the organization; (b) decrease the response time to a failure, and (c) process task instances in parallel, thus helping to balance the load.

³It's unreasonable to assume that the other agents might fail, but the sentinel will never fail

The obvious advantage of the survivalist approach is that no one agent is overburdened with the monitoring responsibilities. Also there is no central point of failure and no agent with global knowledge of the organization. Furthermore, the survivalist approach can take into account the interplay between a satisfying organizational structure and probability of failure. For example, one way of achieving a higher level of robustness in the survivalist approach, given a large number of agent failures, would be to relax the task deadlines. However, such a relaxation would result in the system using fewer agents in order to conserve resources, which in turn would have a detrimental effect on the robustness. These advantages come at a cost of increased complexity of the monitoring mechanism.

Creating a monitoring set of agents The monitoring set of an agent, A, is defined as the set of agents that are responsible for monitoring A for failures. We assume that the minimum cardinality of this set, N is an input to the organization⁴. Also in our approach, all monitoring is mutual, i.e. if A is in the monitoring set of B (i.e. if A is responsible for monitoring B), then B is in the monitoring set of A. This is by design, because Agent A on receiving an *are-you-alive* request from Agent B, already knows that Agent B is alive and does not need to send it a separate request.

Each agent is responsible for determining its monitoring set. At the time the agent is first created, it runs the following algorithm:

1. Determine my *related set*. The related set of an agent, A is the set of agents that have a coordination relationship with this agent. (This coordination relationship would exist because of interdependent tasks and NLEs in the task structures of the agents).
2. If the number of related agents is greater than N , send a message to each of the related agents requesting the cardinality of their monitoring set, and goto Step 3. If this number is less than N , add all of these agents to my monitoring set and goto Step 4.
3. Pick N agents with the lowest monitoring-agent-set cardinalities to be my monitoring set.
4. Send messages to other agents requesting their monitoring set cardinalities. (This can be done using a single broadcast). Pick agents according to their cardinalities until either we have N agents in our set or until all the agents have been exhausted (i.e. there are less than N agents in the whole organization).

Finally, once an agent has its monitoring set, it can send a message to each agent in its set requesting them to monitor its health. In addition to being distributed, other advantages of this algorithm are: (a) Steps 1–3, can be piggy-backed onto the coordination-mechanism negotiation messages exchanged with the agents in the related set and (b) This scheme will reduce the frequency of *are-you-alive* mes-

⁴It should be possible to develop an algorithm for learning the optimal value of N given the environment conditions — i.e. the probability of failure. We plan to incorporate this into our future work.

sages transmitted since the agents will be communicating in-band as a part of their normal tasks processing.

Augmenting the robustness capacity of an organization

The robustness capacity of an organization is defined as the number of agent failures that an organization can withstand. The robustness capacity is equal to the kill count minus 1, where the kill count is the *minimum* number of agent that need to be killed in order to kill the organization (i.e. ensure that the remaining agents cannot complete tasks without having to restart more agents).

The robustness capacity of an organization is dependent on (a) the underlying global task structure and (b) the way it has been divided amongst the agents. The CAFs of the global task structure, especially the root CAF, determines the number of alternatives available for achieving a task. For example, a one-level deep task structure with a *MAX* CAF and three subtasks would have three alternative ways of achieving the task. If each of these three alternatives was divided amongst three agents, the resultant organization would have a kill count of 3 and a robustness capacity of 2.

For example, consider the organization shown in Figure 1(b). This organization has a kill count of two since if the agent responsible for Root-2 (Task D) is killed, the agent responsible for Root-1 would still be able to complete the task and vice-versa. Note that the kill count is 2, even though there are three agents in the organization.

Augmenting the robustness capacity of an organization is the process of adding agents to the organization so as to increase its kill count. Again, we are assuming that the desired kill count, K is an input to the organization. A trivial way to do this would be to replicate each agent $K-1$ times. However, this would be inefficient as it does not take into account the existing kill-count of the organization.

The first step towards increasing the robustness capacity of an organization would be to compute the existing kill-count, and then to “add” agents by breaking up the global task structure and spawning agents in a way that increments this kill-count. Unfortunately, the bad news is that computing the kill-count of an organization based on an underlying TÆMS task structure is NP-hard. An informal proof follows:

This proof is based on the reduction of a minimum set covering problem to a TÆMS based organization, where the kill-count of that organization would be the solution to this problem. Assume a ground set M consisting of m elements, $\{e_1, e_2, \dots, e_m\}$ and n subsets $\{s_1, s_2, \dots, s_n\}$. Create a TÆMS task structure, with a *MAX* CAF as the root and the subsets, $\{s_1, s_2, \dots, s_n\}$ as its subtask nodes. Finally replace each node s_i with a *MIN* CAF task, the subtasks of which will be the methods, $\{m_1, m_2, \dots, m_j\}$, where each method corresponds to an element of s_i . Finally, assign m agents to the organization, where each method corresponding to e_i is assigned to agent a_i . This reduction will provably take polynomial time.

Since, (a) the problem of computing the kill count/robustness capacity of a problem is NP-hard and (b) augmenting the organization by spawning agents at specific places will interfere with other desirable character-

istics such as balancing the execution time and maximizing quality, we chose an alternative approach to augmenting the robustness capacity.

In our approach, the initial root node of the global task view is cloned N times and each clone is allocated to a separate agent. These agents are responsible for individually forming their own independent organizations and spawning and composing with agents independently.

Frequency of *are-you-alive* messages Ideally, we want each agent in the monitoring set of an agent A to send an *are-you-alive* request at a different time. To achieve this, we initialize each agent with a random seed. The *next-poll-time* is initialized to the *poll-interval* plus this random seed. Also the next-poll-time is recalculated on receiving *any* message from the monitored agent.

Evaluation

To evaluate the two robustness mechanisms, we ran a series of experiments that simulated the operation of the OSD organization when those mechanisms were employed. We tested the performance of the survivalist approach against the citizens approach with the following (per agent/per cycle) probabilities of agent failures: 0.000, 0.002, 0.006 and 0.010. Here a probability of 0.006 means that on every clock cycle, each agent has a 0.6 % chance of failing.

We used a randomly generated TÆMS task structure with a maximum depth of 4, branching factor of 3, and NLE count of 10 to seed the experiments. We were careful to use the same task structure, task arrival times, task deadlines and random numbers for each of the (robustness-mechanism, failure probability) pairs. Each experiment was repeated 15 times using a different randomly generated task structure. The experiments were run for 2500 clock cycles. The average results of various measured parameters are shown in Figure 3.

We also tested the statistical significance of the obtained results using the *Wilcoxon Matched-Pair Signed-Rank* tests with $p < 0.05$. *Matched-Pair* signifies that we are comparing the performance of each robustness approach on precisely the same randomized task set, environmental conditions and failure probabilities within each separate experiment. Some interesting observations are:

- We were pleasantly surprised by the overall performance of these approaches. In particular there was no statistical difference between the number of tasks completed in the absence of agent failure and all the situations in which the citizens approach was employed. Only the survivalist approaches at probabilities of failure of 0.006 and 0.010 performed statistically worse than the no failure case.
- As expected the fewest number of agents were used when the probability of failure was 0. Surprisingly, not only did the survivalist approach use fewer agents than the citizens approach, this result was *statistically significant* for failure probabilities of 0.006 and 0.010. We are not really sure why the survivalist approach uses fewer agents than the citizens approach — we expected the survivalist approach to use more agents since the survivalist ap-




Figure 3: Graph showing the various measured parameters for the different robustness mechanisms. The numbers below the mechanisms indicate the probability of agent failure per agent per cycle.

proach pro-actively replicates agents to increase the robustness capacity of the organization. More research will be needed to ascertain the exact cause of this result.

This result does, however, show that the survivalist approach is a credible distributed alternative to the citizens approach.

- The citizens approach used a significantly larger number of messages to achieve similar levels of robustness. This is primarily because of the monitor agent has to record all the messages exchanged by the agents.
- Finally, the turnaround time (the time difference between task completion and generation) is statistically significantly lower for the survivalist approaches with probabilities of 0.002 and 0.006 than all the other approaches including the no failure approach. This is probably the result of extra agents better balancing the load of the organization.

Conclusion

This paper was primarily concerned with the robustness of organizations, generated at run-time through the use of organizational-self design (OSD), in the presence of agent failures. We have incorporated the two commonly used approaches to robustness, that is, the citizens approach and the

survivalist approach into our OSD system. The citizens approach is simpler and more effective than the survivalist approach but suffers due to the use of a single centralized and omniscient monitoring agent to achieve its robustness. The survivalist approach, on the other hand, is truly distributed and we have shown it to be a credible alternative to the citizens approach since it uses fewer agents and fewer (communication) resources to achieve similar levels of robustness.

In our future work, we would like to develop a more fine-grained method for augmenting the robustness capacity of an organization in the survivalist approach — one that will use some heuristic to compute or *underestimate* the current robustness capacity of the organization and then augment the robustness capacity by cloning specific agents. We would like to see if such a fine-grained approach will use even fewer agents than the currently presented survivalist approach.

References

- Chen, W., and Decker, K. S. 2005. The analysis of coordination in an information system application - emergency medical services. In *Lecture Notes in Computer Science (LNCS)*, number 3508, 36–51. Springer-Verlag.
- Decker, K. S., and Li, J. 2000. Coordinating mutually exclusive resources using gpgp. *Autonomous Agents and Multi-Agent Systems* 3(2):133–157.
- Decker, K. S. 1995. Environment centered analysis and design of coordination mechanisms. *Ph.D. Thesis, Department of Computer Science, University of Massachusetts, Amherst*.
- Dellarocas, C., and Klein, M. 2000. An experimental evaluation of domain-independent fault handling services in open multi-agent systems. *Proceedings of the International Conference on Multi-Agent Systems (ICMAS-2000)*.
- DeLoach, S.; Oyen, W.; and Matson, E. 2008. A capabilities-based model for adaptive organizations. *Autonomous Agents and Multi-Agent Systems* 16(1):13–56.
- Frederic, G., and Jacqueline, A. 1995. *Logical reorganization of DAI systems*. Lecture Notes in Computer Science. Springer Berlin / Heidelberg.
- Ishida, T.; Gasser, L.; and Yokoo, M. 1992. Organization self-design of distributed production systems. *IEEE Transactions on Knowledge and Data Engineering* 4(2):123–134.
- Kamboj, S., and Decker, K. S. 2007. Organizational self-design in semi-dynamic environments. *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* 1220–1227.
- Klein, M.; Rodriguez-Aguilar, J.-A.; and Dellarocas, C. 2003. Using domain-independent exception handling services to enable robust open multi-agent systems: The case of agent death. *Journal for Autonomous Agents and Multi-Agent Systems* 7(1 - 2):179 – 189.
- Lesser, V. R.; Decker, K.; Wagner, T.; et. al. 2004. Evolution of the gpgp/tæms domain-independent coordination framework. *Autonomous Agents and Multi-Agent Systems* 9(1-2):87–143.
- Marin, O.; Sens, P.; Briot, J.; and Guessoum, Z. 2001. Towards adaptive fault tolerance for distributed multi-agent systems. *Proceedings of European Research Seminar on Advances in Distributed Systems (ERSADS 2001)*.
- Shehory, O.; Sycara, K.; Chalasani, P.; and Jha, S. 1998. Agent cloning: an approach to agent mobility and resource allocation. *IEEE Communications Magazine* 36(7):58–67.
- Sims, M.; Goldman, C. V.; and Lesser, V. 2003. Self-organization through bottom-up coalition formation. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 867–874. New York, NY, USA: ACM Press.
- Smith, R. G. 1988. The contract net protocol: High-level communication and control in a distributed problem solver. In *Distributed Artificial Intelligence*, 357–366. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- So, Y., and Durfee, E. H. 1996. Designing tree-structured organizations for computational agents. *Computational and Mathematical Organization Theory* 2(3):219–245.
- Wagner, T.; Raja, A.; and Lesser, V. 2006. Modeling uncertainty and its implications to sophisticated control in taems agents. *Autonomous Agents and Multi-Agent Systems* 13(3):235–292.
- Zimmerman, T. L.; Smith, S.; Gallagher, A. T.; Barbulescu, L.; and Rubinstein, Z. 2007. Distributed management of flexible times schedules. In *Sixth International conference on Autonomous Agents and Multiagent Systems (AAMAS)*.