

Toward an Understanding of the Impact of Software Personal Assistants on Human Organizations

Steven Okamoto
Carnegie Mellon University
sokamoto@cs.cmu.edu

Paul Scerri
Carnegie Mellon University
pscerri@cs.cmu.edu

Katia Sycara
Carnegie Mellon University
katia@cs.cmu.edu

ABSTRACT

Intelligent software personal assistants for human organizations are an active research area within the multiagent community. However, while many capabilities for these software personal assistants are imagined or already developed, there has been no quantification of how an organization's performance is improved by software personal assistants. Moreover, while intuitively organizations will adapt to take advantage of the new technology, there has been no work looking at how organizations should or will change in response to the new technology. This paper presents a first step toward addressing this oversight. Specifically, a computational model of the working of an organization and how software personal assistants will affect that organization is developed that allows effects of software personal assistants to be modeled. By varying the potential capabilities of the software personal assistants and the structure of the organization, we can explore the impact of the technology. Our results show that managing task contingencies can greatly improve organizational performance by as much as 45%.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*multiagent systems*

Keywords

Organizational Design, Personal Software Assistants

1. INTRODUCTION

A variety of recent agent research has focused on developing techniques for building intelligent software personal assistants (SPAs) for humans in complex organizations. Each human in the organization will have his or her own SPA which is anticipated to have a wide variety of capabilities from scheduling joint activities[19, 23, 11, 21], monitoring and reminding individuals of key timepoints[7, 4], sharing key information[23, 20], assisting in negotiation decision

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

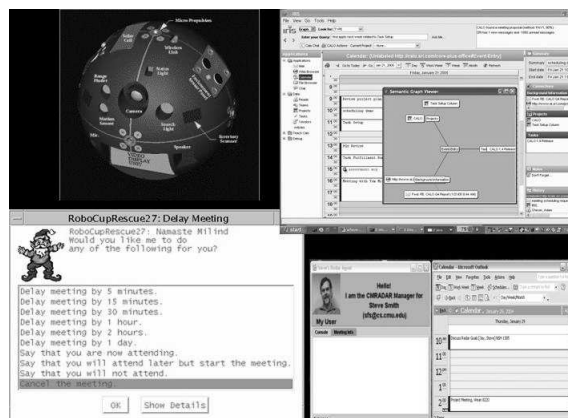


Figure 1: A sampling of SPA technology. Clockwise from upper right: CALO, RADAR, Electric Elves, PSA.

support[14], and even ordering lunch[7]. Just as previous technological innovations have revolutionized human organizations, SPAs are likely to revolutionize the way organizations ranging from first responders to the military to commercial enterprises organize their joint activities. However, despite much research into enabling technologies, there have been no previous efforts to understand how human organizations might utilize the technology. Moreover, it is not even well understood which SPA capabilities are most useful for improving the effectiveness and efficiency of an organization.

Outside of the multiagent community, organizations and their adaption to technological change has been extensively studied. Some research, particularly within economics, uses abstract analytical models to determine optimal features of an organization given the task for which it is designed[12, 10, 3]. Other work, often in sociology, attempts to develop analytic models to explain observed phenomena[5, 17]. Unfortunately, such analytic models are typically too abstract to model the subtle effects of SPA technology. One sub-field of research focuses on design and adaption of human organizations[5], but has not looked specifically at the impact of new technology on those organizations. Within the multiagent community, researchers such as Dignum[9], Horling[13] and Dastani[8], have developed interesting techniques for modeling organizations. As these languages mature they will play an important, complementary role to this work. Hence, while organizations are extensively studied, no work specifically allows quantification of the impact of SPAs on an organization.

In this paper, we present our initial efforts at modeling the impact of SPAs on human organizations. The aim is to provide SPA developers with input on which capabilities are most likely to be useful and to provide SPA adopters with input on how to best apply the technology. Specifically, we are interested in organizations configured for a particular large scale task, e.g., response to a major disaster or a military operation, where SPAs such as the DARPA COORDINATORS assistants[23] may be deployed. Since SPA capabilities typically affect detailed work activities and their broader influence is unclear, we must model the organization at a fine grained level of detail, capturing, for example, message paths, decision making, times, etc, in order to see their effects. An abstract simulation environment takes proposed organizations and tasks to be performed by that organization and computes key properties of task execution, including how well and how quickly the organization performed the task and how robustly it handled individual failures. The simulation captures important aspects of the operation of the organization, such as non-determinism and individual cognitive limits, but abstracts away domain level details, making it feasible to evaluate many instances.

In our computational model, SPAs are represented as changing and relaxing constraints on human cognition. With reference to ongoing projects and previously published literature about SPA technology, we identified four key capabilities that are being developed and inferred how these would alter how a person operates in an organization. For example, several key technologies[23, 4] will have the effect of increasing the speed with which an individual in an organization will be able to make decisions. The model of the organization allows the impacts of capabilities to be considered individually or in combination.

To quantify the impact of SPAs on an organization, we generated a large number of different organizational structures and evaluated the performance of each organizational structure with and without combinations of SPA capabilities. This general approach was used for two reasons. First, it enabled us to quantify the impact of SPAs on the organization by measuring the difference in performance between organizations with and without SPAs. Specifically, we found that managing contingencies to prevent actor overload improved performance most significantly, allowing a 45% performance improvement over an organization without SPAs. Second it allowed us to compare the performance of different organizational structures to determine what changes may be needed to leverage SPA technology. We found that the types of hierarchies currently used by many organizations remains effective in the presence of SPAs.

2. RELATED WORK

The effect of information technology on human organizations has been a focus of considerable research, especially in business and organization theory[5, 17, 1]. While it is generally assumed that the deployment of such technology in an organization reduces coordination costs and enhances performance[16], other research[6] indicates that the adoption of information technology can have unexpected and deleterious effects. It is thus essential to consider the interactions between proposed technology and the dynamics of the organization deploying that technology.

Even in cases where the adoption of information technology improves organizational efficiency, the underlying orga-

nizational structure must often change in order to realize the full benefits. Two notable and well-studied such technologies are electronic marketplaces in supply chain management[17] and e-mail for communication and coordination[1]. While the organizational adaptations needed to leverage those technologies are now understood, there is to date no research on the organizational changes necessary to most fully exploit the newer technology of SPAs. Research on SPAs has instead focused on issues of actually building and deploying such complex software systems, in domains ranging from military applications[23] to disaster response[20] to office scheduling and collaboration[7, 19] to space environments[4]. Studies involving agents actually assisting human teams perform tasks have been performed[22, 7], but such an approach is impractical for studying the effects of agents on large organizations, and also require the SPA to already be built.

Organizational research in the multiagent community has tended to focus on modeling organizations in order to guarantee logical properties during execution[9]. While important for a complete understanding of organizational behavior and the successful integration of agents into human organizations, this approach does not specifically address the impact of SPAs, and also neglects quantitative evaluation of organizational performance. Recent work[13] addresses the issue of quantitative evaluation, but is focused primarily on organizational modeling as opposed to design, and still does not specifically take into account SPAs.

3. PROBLEM

We represent an organization of human actors that perform tasks and communicate in order to make decisions. Successful execution of sensing tasks yields information that is used as input to making decisions. This abstraction describes a wide variety of real-life organizations, from military units that detect and eliminate enemy threats, to businesses that identify customer needs and reallocate investments.

A number of assumptions govern the behavior of actors in our model:

- An actor that attempts to perform a task may *fail*. The actor cannot try to perform the task again.
- Tasks have fixed execution durations during which time they must be performed, uninterrupted.
- Actors can perform only a limited number of tasks simultaneously. In particular, an actor can only perform a single sensing task at a time, or 3 decision-making tasks at a time. Actors cannot simultaneously perform sensing and decision-making tasks. An actor who attempts to exceed these limits becomes overloaded and cannot perform *any* tasks.
- Actors can only transmit information along predetermined communication links.
- Actors are bounded in the number of communication links they can maintain. Each actor can have at most max_{in} incoming communication links and max_{out} outgoing communication links.
- Communication takes a fixed amount of time.
- Actors can only perform decision tasks if they have all the required information.

We do not explicitly represent SPAs in the organization, but instead model their effects as changes in organizational constraints and parameters. For example, the presence of SPAs may be modeled by increasing the communication speed from a baseline for an organization without SPAs.

Formally, an organization O is a tuple

$$O = \langle A, T, I, Prov, Req, Cap, Fail, G, speed_a, speed_c, M, d \rangle$$

with the following components:

- $A = \{a_1, a_2, \dots, a_m\}$: the set of human actors in the organization
- T : the set of tasks to be performed by the actors. Tasks are partitioned in two orthogonal ways:
 - $T = T_P \cup T_C$
 - * $T_P = \{t_1, t_2, \dots, t_n\}$: the set of *primary tasks* to be performed.
 - * $T_C = \bigcup_{t_i \in T_P} C_i$: the set of *contingencies* to be performed in case of task failures.
 - $T = T_D \cup T_S$
 - * T_D : the set of *decision-making tasks* that require information to be performed.
 - * T_S : the set of *sensing tasks* that provide information when successfully performed.
- $I = \{I_1, I_2, \dots, I_l\}$: the set of information produced by sensing tasks and used by decision tasks.
- $Prov : T \rightarrow \mathcal{P}(I)$: the information provision function mapping (sensing) tasks to the set of information provided by successfully completing the (sensing) task.
- $Req : T \rightarrow \mathcal{P}(I)$: the information requirement function mapping (decision-making) tasks to the set of information required to perform the (decision-making) task.
- $Cap : A \times T \rightarrow [0, 1]$: the capability function indicating how well actors perform tasks.
- $Fail : A \times T \rightarrow [0, 1]$: the failure function indicating the probability actors will fail at performing tasks when attempted.
- G : the directed graph with actors as vertices and directed edges indicating possible communication paths between actors.
- $speed_a$: the rate at which an actor can make a decision.
- $speed_c$: the speed of communication.
- M : the assignment matrix of primary tasks and contingencies to actors. Each primary task and contingency must be assigned to exactly one actor.
- d : the deadline by which time all decisions must be made in order for the organization to succeed.

For each primary task $t_i \in T_P$, there exist 0 or more contingencies, $C_i = \{c_i^1, c_i^2, \dots, c_i^{k_i}\}$. The contingencies are pre-determined backup tasks that are invoked in case of task failure, and are ordered by preference. Hence contingency c_i^1 must be performed if primary task t_i fails, and contingency c_i^2 must be performed if contingency c_i^1 fails, and so on. The set of all contingency tasks is written $T_C = \bigcup_{t_i \in T_P} C_i$.

3.1 An Organization Model

The directed graph G in the specification of the organization describes the structural communication links that exist between members of the organization. The vertices of G are the actors in A . A directed edge from $a_i \in A$ to $a_j \in A$ indicates that a_i can transmit information to a_j . A path through G must exist from one actor to another when the first supplies information to the second, i.e., a path must exist from two actors a_i and a_j if $M(a_i, t) = 1$ and $M(a_j, t') = 1$ and $Prov(t) \cap Req(t') \neq \emptyset$. Communication will actually occur on the edges of G if it falls on a path required for delivering information.

Notice that tasks in T_C may not actually be performed, hence edges in G that only transmit information for these tasks may not actually be used. Moreover, even those edges that are on paths required for tasks in T_P may not be used if the actor performing the task fails, since a contingency would then be invoked. Hence we distinguish between two graphs, G and \bar{G} . We refer to G as the potential structure graph, since it must contain all the edges that could be required for any possible combination of tasks and contingencies. \bar{G} , on the other hand, is called the *instantiated structure* and only contains those edges that are actually being used given the current set of tasks and contingencies.

Human cognitive limitations restrict the number of contacts a person can effectively handle. We model this as constraints on the potential structure graph G . In particular, we bound the in- and out-degrees of the actors in G : $\forall a \in A, indegree(A) \leq max_{in}$ and $outdegree(A) \leq max_{out}$.

3.2 Execution

The actors attempt to perform primary sensing tasks assigned to them in M . For any primary task failure, contingencies are invoked as described in section 3. If an actor a succeeds at a sensing task t_s (whether it is a primary task or contingency), the actor gains information $Prov(t_s) \subseteq I$. This information has an associated quality which we denote by $Qual : I \rightarrow [0, 1]$. For all $I_j \in Prov(t_s)$, we let

$$Qual(I_j) = Cap(a, t_s) \in [0, 1] \quad (1)$$

The actor then transmits the newly learned information to any neighbors as defined in G . These neighbors store the information in their personal knowledge bases and also in turn propagate the information to their neighbors.

When an actor assigned a decision-making task gains all the information required to make the decision, either from other actors or by performing the sensing tasks itself, it attempts to perform the decision (which can fail). If the actor succeeds, the decision is made with a quality that depends both on the capability of the decision-making actor and the quality of the provided information, and the organization gets reward for making the decision. This reward is denoted $Reward : T_P \cap T_D \rightarrow [0, 1]$. For a decision-making task t_d successfully performed by actor a , if $t_d = t_i \in T_P$ or $\exists t_i \in T_D$ with $t_d = c_i^r \in T_C$, reward is given by

$$Reward(t_i) = Cap(a, t_d) \left(\sum_{I_j \in Req(t_d)} Qual(I_j) \right) \quad (2)$$

where $Qual(I_j)$ is defined above in equation 1. The total reward earned by the organization will be one measure of organizational performance, as described next.

3.3 Objective Function

The “performance” of an organization is complex and multi-attributed. In this work, we focus on three specific measures.

- *Success Rate* This is a measure of how well the organization handles individual failure and is calculated as the fraction of the time the organization makes all of the decisions represented in T_D before the deadline d .
- *Reward* This is a measure of how well the organization makes decisions when it succeeds, and is calculated:

$$\sum_{t_d \in T_D \cap T_P} \text{Reward}(t_d)$$

where $\text{Reward}(t_d)$ is given by equation 2 in section 3.

- *Speed* This is a measure of the organization’s efficiency and is calculated as the difference between the deadline time d and the time at which the last task finishes execution.

The overall objective function used for performance is a weighted sum of these three metrics.

4. EXPERIMENTAL SETUP

We evaluated a large number of organizations with a variety of organizational structures. The basic experimental setup was as follows:

- 20 human actors.
- 44 total primary tasks and contingencies.
 - 4 decision tasks, each requiring information provided by 3 distinct sensing tasks.
 - Each decision task had 1 contingency.
 - Each sensing task had 2 contingencies.
- Performance function: each metric (success rate, reward, and speed) was normalized prior to weighting. The weights were:
 - Success rate weight = 5
 - Reward weight = 1
 - Speed weight = 1

We tested 9 types of organizational structures. These 9 types of structures can be grouped into 4 broad classes:

- Ring. Actors form a singly linked cycle. If actors can link to more than one actor (i.e., $max_{out} > 1$), actors form at regular intervals across the ring to minimize the average path length to every actor.
- Regular hierarchy. Actors form a full n -ary rooted tree, where n is the span of control. We tested hierarchies with three spans of control:
 - RH3: regular hierarchy with span of control of 3.
 - RH4: regular hierarchy with span of control of 4.
 - RH5: regular hierarchy with span of control of 5.

Every actor in a regular hierarchy has one outgoing link and a fixed number of incoming links, so this structure is invariant to changes in max_{out} and max_{in} . Note, however, that max_{in} must be greater than or equal to the span of control for a regular hierarchy to be allowed.

- Scale free network. These were generated according to [2]. A small initial set of actors were chosen and the remaining actors were added sequentially, with each adding max_{out} directed links to the actors already in the graph. Link formation was done preferentially to actors with higher connectivity. No actor was allowed to have more than max_{in} incoming links.
 - SF: linking preference was proportional to the in-degree of the target actor.
 - SFpreR1: each actor in the initial set added an edge at uniformly at random to another actor in the initial set before generation began. Preferential linking was by in-degree.
 - SFB: linking preference was proportional to the total degree (in-degree + out-degree) of the target actor
 - SFBpreR1: each actor in the initial set added an edge at uniformly at random to another actor in the initial set before generation began. Preferential linking was by total degree.
- Random graph (Rand). max_{out} links were added randomly from each actor to other actors. No actor was allowed to have more than max_{in} incoming links.

4.1 Baseline Performance

In the baseline configuration, $max_{out} = 1$ and $max_{in} = 4$. Note that this meant that actors in ring structures did not have any links across the ring. In addition RH5 could not be used in the baseline configuration since actors in RH5 must have 5 incoming links, which exceeds the constraint on max_{in} .

Figures 2 – 4 show organizational performance as described in section 3.3. Figure 2 shows the success rate component, Figure 3 shows the reward component, and Figure 4 shows the speed component. Baseline performance is depicted by black columns labeled “No SPAs”.

Observe that the regular hierarchies show the highest success rates in Figure 2. This is because there is a high congruence between hierarchical structure and the decision making problem we are considering: actors lower in the hierarchy perform sensing tasks, and they transmit information upward to decision-making actors. The structure of the organization naturally constrains the flow of information to the actors needing that information. The root actor will eventually have all information gained by actors in the organization, and is thus well placed to make decisions or to perform contingencies for decision-makers lower in the hierarchy.

In contrast, ring structures show the lowest success rate. This is because of the lack of congruence between the ring structure and the problem. Because the ring structure is strongly connected, all actors eventually gain any information gained by any actor. However, there is no structural force to encourage the specialization of actor, and so many actors overload by attempting to simultaneously perform sensing contingencies and make decisions.

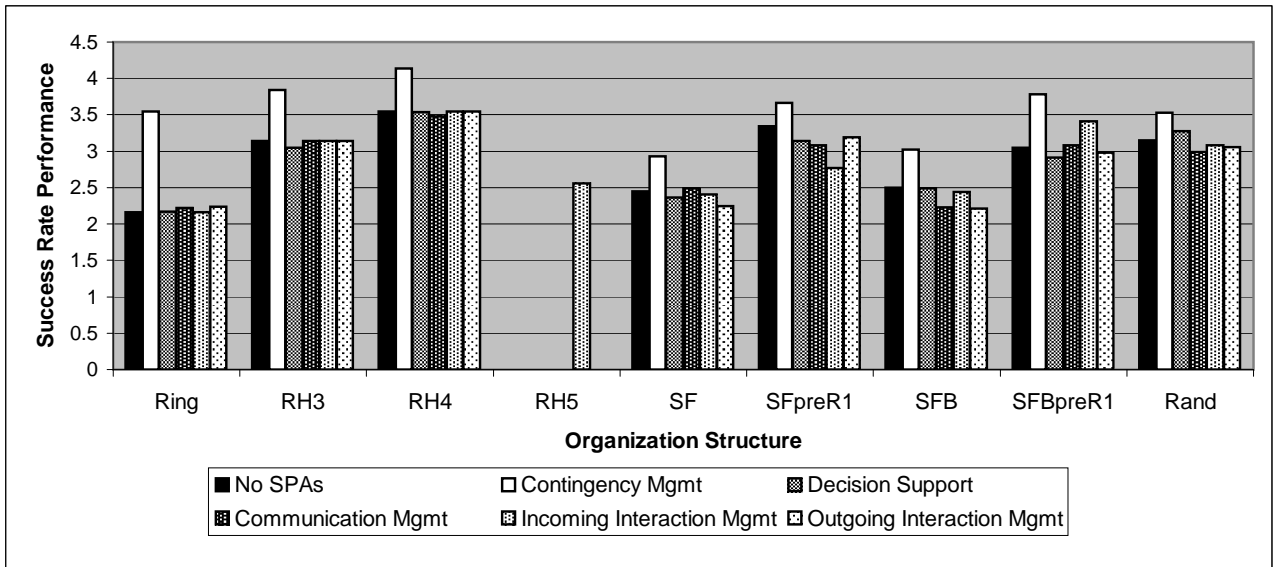


Figure 2: Success rate contribution to performance.

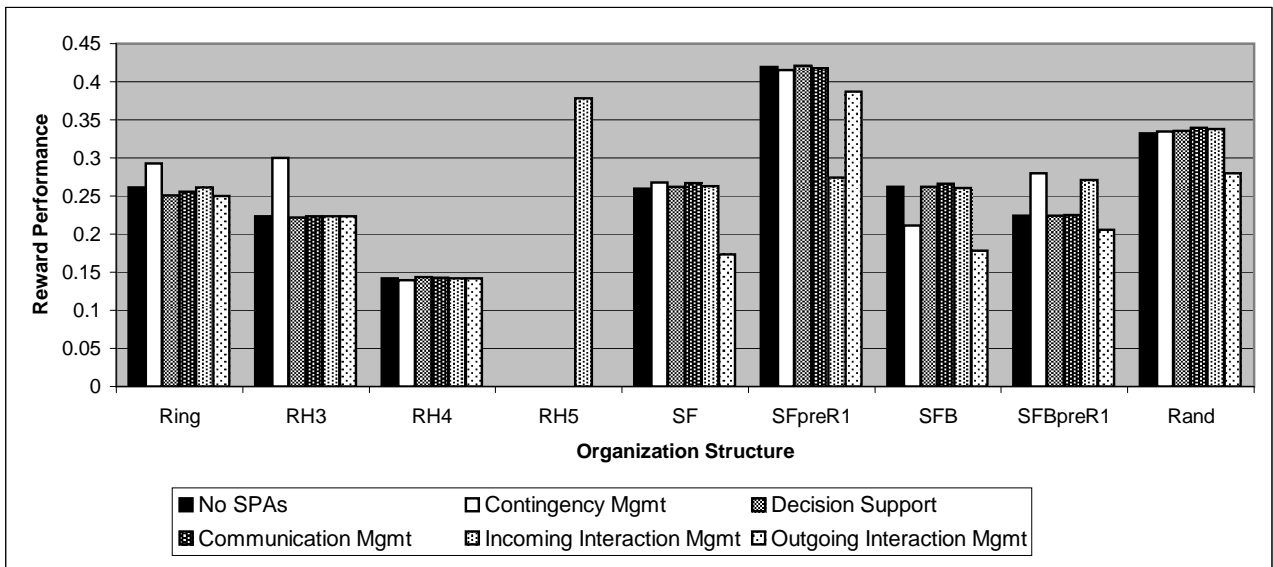


Figure 3: Reward contribution to performance.

Ring structures also show the lowest speed in Figure 4. This is because actors are only allowed a single outgoing link. Thus there are no edges crossing the cycle. As a result, the diameter of the ring structure is linear in the number of actors, while it is logarithm in the number of actors for the other structure types.

Results presented in the next section will detail how performance changes in the presence of SPAs.

5. MODELING PERSONAL ASSISTANTS AND RESULTS

The capabilities of SPAs are generally being designed to overcome perceived cognitive limitations of humans that are thought to limit the performance of an organization[22, 23, 4]. In this section, we describe four such limitations and the models of SPA capabilities — contingency management, decision support, communication management, and inter-actor interaction management for information sharing — that might overcome these limitations. We also present experimental results showing the effect of these capabilities on organizational performance in figures 2 – 4.

Many other SPA capabilities (and corresponding cognitive limits) might have been modeled. We focused on this initial set because there was active research in the field on developing these capabilities. Moreover, we selected SPA capabilities that appeared most suited to the special purpose organizations that are the focus of this paper. Future work will expand the scope of SPA capabilities, including to those where there is not yet active research.

5.1 Managing Task Contingencies

An SPA may manage contingencies in order to prevent its actor from being overloaded. Instead of invoking contingencies only when a domain-level failure is suffered, the SPA may monitor the actor’s current status and workload and automatically invoke contingencies for any tasks that would overload the actor. Similar approaches have been developed or proposed for SPAs in other environments[7, 15]. Intuitively, these SPA capabilities should make the organization more robust to individual failure. Moreover, we hypothesize that it will be possible for the organization to arrange itself such that more capable team members will be involved in more contingencies, since their task performance provides the most reward to the whole organization. Managing contingencies or assigning roles has been a key area of research, but this is the first work that attempts to quantify the benefit of such a capability on an organization and understand how the organization may change to leverage this capability.

The results of adding contingency management capabilities to the simulated organizations are shown in figures 2 – 4 in the white columns labeled “Contingency Mgmt”. From Figure 2, it is clear that contingency management greatly improved the success rate for all tested organization structures. It was particularly effective in improving the success rate of the ring structure, which is especially prone to actor overload: the success rate of the ring structure improved by 64%, while the success rate of RH4 improved by only 16%. This large gain translated into an overall gain in performance of 45% for the ring structure, against an average performance gain of 17% for all structures.

As shown in Figure 3, the addition of SPA-managed contingencies also improved the reward obtained by some of

the organizations, in accordance with our hypothesis. For example, reward increased by 24% for SFBpreR1 and 34% for RH3. The only structure that lost reward was SFB, which saw reward decrease by 20%.

5.2 Decision Support

As introduced above, a key cognitive constraint on humans in an organization is the number of decisions that can be effectively made in a limited amount of time. SPAs could relax this constraint in a variety of ways, e.g., taking over routine tasks to provide more time for “decisions”[18, 20] or collating and presenting information in a way that speeds up the actual decision-making process[15]. Without reference to specifically how it is done, in the computational model of the organization, we simply allow a person to effectively make more decisions in the same period of time, if they have access to SPA technology, by doubling $speed_d$, the speed of making decisions.

The performance of organizations with this capability is shown in figures 2 – 4 in the columns labeled “Decision Support”. We expected that this capability would impact the success rate and the speed of execution. It would affect the success rate because each decision would have execute more quickly, thereby reducing the amount of time that the actor may be overloaded. It would also increase the speed of execution directly. Both these factors turned out not to have noticeable impacts of performance because the decision duration was already quite low, especially in comparison to the deadline (which was used to normalize the speed metric). When modeling other problems this capability may yet prove to be significant and merits further study.

5.3 Communication Management

Wired and wireless networks upon which much communication in an organization travels is amazingly fast in contrast to human communication speeds. However, in most cases, people are still required to input information and subsequently process received information. The input and output time is subject to human limitations and is fast becoming (if it is not already) the bottleneck in intra-organization communication, and could be alleviated by SPAs.

We modeled this capability by doubling the communication speed in the organization. The effects on performance can be seen in figures 2 – 4 in the columns labeled “Communication Mgmt”. This capability had little effect on either success rate or reward. As shown in Figure 4, it also had only small but noticeable effect on the speed of most organizations, increasing average speed by 8%, because they already displayed efficient communication. As expected, however, this capability had a substantial impact on ring-structured organizations, with a speedup of 29%.

5.4 Managing Interactions for Information Sharing

The rate at which decision-makers in an organization can effectively make decisions is limited by two key factors: (i) their ability to get the appropriate information to make required decisions and (ii) their innate processing speed for making those decisions (see section 5.2). The cognitive load of managing incoming information from others in the organization and directing outgoing information to others is a constraint that may be relaxed by SPA technology[15, 17]. Specifically, a decision-maker may be able to handle input




Figure 4: Speed contribution to performance.

from a greater number of members of an organization and hence, either get more direct access to that information or make more decisions. For example, the CEO of an organization needs information from all over parts of the organization to make key strategic decisions. However, it is cognitively and organizationally infeasible for a person to directly receive input from a large number of people[5]. Hence, in the case of the CEO, that information must pass through a small number of department heads, both delaying the information and potentially distorting it. Conversely, suppliers of information are limited in how many others they can provide information to, requiring organizational structures that channel information and introduce delays and distortion. SPA technology could allow information suppliers to more directly communicate with the decision-makers requiring their information.

The effect of managing interactions was modeled by changing the values of max_{in} and max_{out} . max_{in} was increased from 1 to 2, and max_{out} was increased from 4 to 5. We considered the two effects separately, with the resulting performance shown in figures 2 – 4 by the columns marked “Incoming Interaction Mgmt” and “Outgoing Interaction Mgmt”, respectively.

It is important to note that unlike the other capabilities tested, interaction management related directly to the organizational structure by allowing actors to form more links. Some structural classes were completely unaffected by such changes. For example, in a regular hierarchy, each actor has a single outgoing link, and a fixed number of incoming links (bounded above by max_{out}), and so changing max_{out} had no effect on regular hierarchy performance. Similarly, increasing max_{in} had no effect on the ring structure performance. However, increasing max_{out} allowed actors in the ring structure to add links crossing the cycle, which reduced the diameter of the graph by half. This translated into impressive gains in speed of 36%, as shown in Figure 4.

The increase of max_{in} from 4 to 5 allowed us to test a new regular hierarchy, RH5, in which actors had a span of control of 5. This structure had a poor success rate, lower

than that of both RH3 and RH4, as shown in Figure 2. This is because the hierarchy of RH5 was too shallow; there were too few interior actors to make all the decisions and so they are being overloaded.

5.5 Multiple Capabilities

We also tested the organizations with different combinations of SPA capabilities. Because there are four capabilities (those described in subsections 5.1 – 5.4), with information sharing including two distinct effects, there are a total of 32 distinct combinations of capabilities. Figure 5 shows the performance breakdowns of the SFpreR1 scale free structure for all combinations of our SPA capabilities; the number of capabilities in the combination increases along the x-axis, from no capabilities on the left to all capabilities on the right. It is clear that not all combinations of capabilities have equal effects on performance, and many (including the combination of all capabilities) even negatively impact performance. This emphasizes the need to carefully select SPA capabilities to fit the organization.

Several of the effects seen in the individual capability experiments were evident in this experiment as well. For instance, managing contingencies to prevent actor overload provided substantial gains in organizational performance. Also, managing interactions frequently decreased performance by encouraging flatter, bushier structures. The resulting decrease in levels of redundancy that can be exploited through the assignment of contingencies leads to higher chances of actor overload.

Managing contingencies can offset this increase, as shown by the best performing organization using contingency management, communication management, and incoming interaction management. In this case, the increased number of incoming links allow a more capable actors to make decisions, yielding higher rewards. Speed is also increased both directly by communication management, and indirectly because of the bushier structure. This sort of interplay between capabilities is important to take into account when evaluating and designing SPAs.

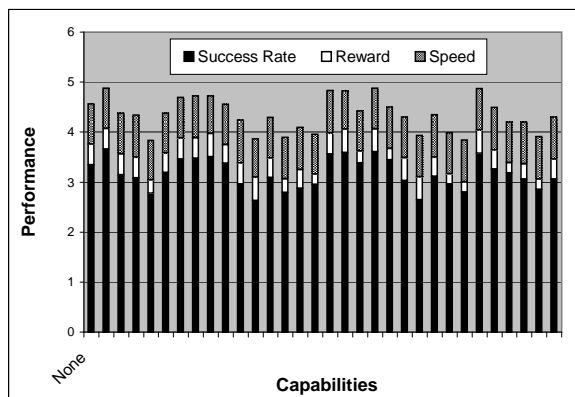


Figure 5: Performance breakdown of organization with scale free structure for all combinations of SPA capabilities.

6. CONCLUSIONS

We have taken an initial step toward identifying the ways in which intelligent personal assistants will impact the performance of human organizations. We have hypothesized a number of possible capabilities of SPAs, evaluated their effect on performance, and shown that the structure of the underlying organization has a substantial effect on the effectiveness of the SPA in improving performance. Our results indicate that SPAs can substantially improve organizational performance by flexibly managing pre-planned contingencies in order to reduce actor overload. However, the effects of the SPAs are far from dramatic and do not impact all structures evenly. Significantly, more work will be required to understand precisely when SPAs are useful. More importantly more analysis should be performed to determine what SPA capabilities might make a dramatic impact.

7. REFERENCES

- [1] Manju Ahuja and Kathleen M. Carley. Network structure in virtual organizations. *J. Computer-Mediated Communication*, 3(4), 1998.
- [2] A.-L. Barabasi, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173–187, 1999.
- [3] A. W. Beggs. Queues and hierarchies. *Review of Economic Studies*, 68:297–222, 2001.
- [4] J.M. Bradshaw, M. Sierhuis, Y. Gawdiak, R. Jeffers, N. Suri, and M. Greaves. Adjustable autonomy and teamwork for the personal satellite assistant. In *IJCAI-01 Workshop on Autonomy, Delegation, and Control: Interacting with Autonomous Agents*, 2001.
- [5] K. Carley. Computational and mathematical organization theory: perspective and directions. *Comput. Math. Organ. Theory*, 1(1):39–56, 1995.
- [6] K. Carley. Computational organization science: A new frontier. *Proceedings of the National Academy of Sciences*, 99:7257–7262, 2002.
- [7] H. Chalupsky, Y. Gil, C. Knoblock, K. Lerman, J. Oh, D. Pynadath, T. Russ, and M. Tambe. Electric Elves: Applying agent technology to support human organizations. In *International Conference on Innovative Applications of AI*, pages 51–58, 2001.
- [8] Mehdi Dastani, M. Birna van Riemsdijk, and John-Jules Ch. Meyer. *Multi-agent Programming: Languages, Platforms and Applications*, chapter Programming Multi-Agent Systems in 3APL. Springer Science+Business Media, Inc., 2005.
- [9] Virginia Dignum, Javier Vázquez-Salceda, and Frank Dignum. Omni: Introducing social structure, norms and ontologies into agent organizations. In *PROMAS*, pages 181–198, 2004.
- [10] Luis Garicano. Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108:874–904, 2000.
- [11] L. Garrido and K. Sycara. Multi-agent meeting scheduling: Preliminary experimental results. In *ICMAS-96*, 1996.
- [12] John Geanakoplos and Paul Milgrom. A theory of hierarchies based on limited managerial attention. *Journal of the Japanese and International Economies*, 5:202–225, 1991.
- [13] B. Horling and V. Lesser. Quantitative organizational models for large-scale agent teams. In *Proceedings of the International Workshop on Massively Multi-Agent Systems*, pages 297–312, December 2004.
- [14] C. Li, J. Giampapa, and K. Sycara. Bilateral contract negotiation decisions with uncertain dynamic outside options. *IEEE Systems, Man and Cybernetics, Part C.: Special Issue on Game Theoretic Analysis and Stochastic Simulation of Negotiation Agents*, 36(1), 2006.
- [15] Pattie Maes. Modeling adaptive autonomous agents. *Artificial Life Journal*, 1(1 & 2):135–162, 1994.
- [16] Thomas W. Malone and Kevin Crowston. The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1):87–119, 1994.
- [17] Thomas W. Malone, Joanne Yates, and Robert I. Benjamin. Electronic markets and electronic hierarchies. *Communications of the ACM*, 30(6):484–497, 1987.
- [18] Tom Mitchell, Rich Caruana, Dayne Freitag, John McDermott, and David Zabowski. Experience with a learning personal assistant. *Communications of the ACM*, 37(7):81–91, July 1994.
- [19] Pragnesh Jay Modi, Manuela Veloso, Stephen F. Smith, and Jean Oh. CM Radar: A personal assistant agent for calendar management. In *LNCS*, volume 3508, pages 169 – 181. Jan 2005.
- [20] N. Schurr, J. Marecki, M. Tambe, P. Scerri, J.P. Levis, and N. Kasinadhuni. The future of disaster response: Humans working with multiagent teams using DEFACTO. In *AAAI Spring Symposium on Homeland Security*, 2005.
- [21] T. Shintani, I. Takayuki, and K. Sycara. Multiple negotiations among agents for a distributed meeting scheduler. In *ICMAS-00*, 2000.
- [22] K. Sycara and M. Lewis. *Team Cognition*, chapter Integrating Agents into Human Teams, pages 203 – 233. Erlbaum Publishers, 2004.
- [23] T. Wagner, J. Phelps, V. Guralnik, and Ryan VanRiper. COORDINATORS: Coordination managers for first responders. In *AAMAS’04*, 2004.