Exploring the Effectiveness of Agent Organizations

Daniel D. Corkill, Daniel Garant, and Victor R. Lesser

University of Massachusetts Amherst, Amherst, MA 01003, USA

Abstract. Organization is an important mechanism for improving performance in complex multiagent systems. Yet, little consideration has been given to the performance gain that organization can provide across a broad range of conditions. Intuitively, when agents are mostly idle, organization offers little benefit. In such settings, almost any organization—appropriate, inappropriate, or absent—leads to agents accomplishing the needed work. Conversely, when every agent is severely overloaded, no choice of agent activities achieves system objectives. Only as the overall workload approaches the limit of agents' capabilities is effective organization crucial to success.

We explored this organizational "sweet spot" intuition by examining the effectiveness of two previously published implementations of organized software agents when they are operated under a wide range of conditions: 1) call-center agents extinguishing RoboCup Rescue fires and 2) agents learning network task-distribution policies that optimize service time. In both cases, organizational effect diminished significantly outside the sweet spot. Detailed measures taken of coordination and cooperation amounts, lost work opportunities, and exceeded span-of-control limits account for this behavior. Such measures can be used to assess the potential benefit of organization in a specific setting and whether the organization design must be a highly effective one.

1 Introduction

Organization is an important mechanism for improving performance in complex multiagent systems [1–6]. Designed agent organizations provide agents with organizational directives that, when followed, reduce the complexity and uncertainty of each agent's activity decisions, lower the cost of distributed resource allocation and agent coordination, help limit inappropriate agent behavior, and reduce unnecessary communication and agent activities [7–9].

When agents are mostly idle, agents can acomplish needed work whether or not they are well organized. This does not mean that effective organization does not affect how efficiently the agents work together, only that unorganized and even misorganized agents have sufficient time and resources to accomplish system objectives when lightly loaded. Conversely, when every agent is severely overloaded, no choice of agent activities achieves system objectives. In this situation, effective organization can help agents be more efficient while failing to achieve objectives fully, but whether they are well organized or not, the system is unable to perform acceptably. Only as the overall workload approaches the limit of agents' capabilities does organization play a significant role in system performance.

2 Organizational "Sweet Spot"

We first explored this organizational-impact conjecture empirically using an previously implemented and described system of organizationally adept BDI agents [10–13] operating in a well-instrumented and highly parametrized experimental platform adapted from the fire-extinguishing portion of RoboCup Rescue [14]. Organizationally adept *call center* agents direct *fire brigade* resources under their control to extinguish fires in important buildings as quickly as possible. There are no fire-brigade bases in the adapted RoboCup Rescue environment, and brigades typically move directly from fire to fire, remaining deployed if they become briefly idle. The objective is to minimize the total importance-weighted damage to buildings. A call center can use its fire brigades to execute plans to achieve its own goals of extinguishing building fires, and it can request temporary use of fire brigades from other call centers when necessary.

Our goal was to learn how the relative performance of previously evaluated agent organizations in this multiagent system changed when operating in environments well outside the conditions typically studied. Whether the existing agents and organization designs in this system were the best possible was not a concern, as better candidates would affect only the magnitude of the relative performances and not their qualitative characteristics. Some observations were intuitive, but there were also surprises, and we believe this to be the first systematic study of organizational impact in a multiagent system over such a broad range of conditions. We ran and analyzed thousands of controlled and repeatable simulation experiments involving dynamic environments in which new fires occur at various city locations throughout the entire duration of an experimental scenario. In such settings, call-center agents have an ongoing (but potentially changing) firefighting workload in which following organizational guidance offers potential advantages over unguided, reactive local decision-making.

Observation 1: Sweet-spot behavior ⇒ Figure 1 shows, as the firefighting workload increases, the performance benefit provided by call center agents that have been given an effective organization design that specifies a responsibility region for each call center (Org) relative to call-center agents operating without any responsibility-region directives (No Org). Call centers give priority to fighting fires in their responsibility regions when such regions are provided. Each of the four call centers controlled six fire-brigade resources. Performance attained in each of the 320 simulation runs is a raw score of the inverse importance-weighted fire damage in the city. We observed that the performance benefit achieved by organization (the raw score improvement) was greatest when the average firefighting workload on brigades was near their capacity to fight important fires (approximately 2.2 fires per timestep).[1]

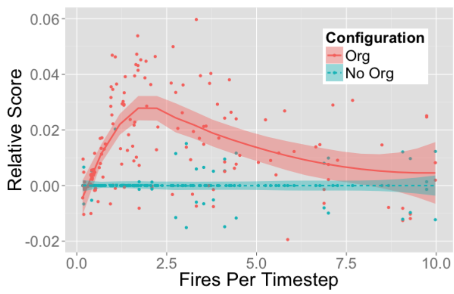

Fig. 1. Relative Score Achieved by Organization

Attenuation of organization benefit outside the sweet spot is a form of phase transition behavior. The transition occurs as the workload approaches the limit of agents' capabilities. The effect of phase boundaries has proved important in satisfiability problems [15–17] as well as to understanding problem difficulty in constraint satisfaction, number partitioning, and traveling salesmen tasks.[2] With multiagent organizations, it is important to determine where on the control complexity scale a system is operating (how important using an effective organization is to system performance) and more generally, when complex multiagent systems are operating within their organizational sweet spot. One may argue that organizations (multiagent or otherwise) will tend to be inevitably operated within the sweet spot region due to real-world economics that limit capabilities and resources to the minimum required to operate effectively.

Upon observing organizational sweet-spot behavior, we took a more detailed look into what was occurring as workload changed that accounted for the benefit attenuation.

3 Performance Factors

Why do we create agent organizations? One reason is that complex agent behavior becomes more structured and understandable through the definition of roles, behavioral expectations, and authority relationships [18]. Additionally, organizational concepts can be used to help design and build agent based systems (organization-based multiagent system engineering). There is also a line of research that addresses *organizational membership* in open agent societies (incentives for organizational recruitment and retention and for the replacement of

[1] All figures illustrate trends as workload (e.g., ignition frequency) is varied. Trend lines are fit using a local linear model, with shaded regions representing a 95% confidence level in the fit. For example, each trend line in the firefighting experiments fits 320 separate simulation runs (drawn as individual dots).

[2] Consider Figure 4 from the classic Kirkpatrick and Selman SAT phase-change paper [16]. If that figure is redrawn as relative difference curves from the k=6/N=40 values, it reveals wide "sweet spot" curves similar to the curves shown in this paper. Relative plots highlight the span and magnitude of performance differences near the phase change, and we consider them more informative in highlighting sweet-spot regions than raw-value plots.

agents that leave the organization).[3] Aligning agents' individual goals and objectives with those of the organization are among the issues addressed in that context. Our focus here is on *organizational control*; specifically, the organizational performance of the members ("how they do their jobs"), rather than on attracting participants from an open pool of agents ("obtaining members for the enterprise") or designing the agent system ("defining what the jobs (roles) are"). We assume here that we have the agents we need, that they all share the organizational objectives (e.g., saving the most important buildings in the city),[4] and that they are competent in their ability to perform tasks necessary to attain that objective.[5]

We distinguish between *operational decision making*, the detailed moment-to-moment behavior decisions made by agents, and *organizational control*, an organization design expressed to agents through directives ("job descriptions") that limit and inform the range of operational decisions made by each agent in the organization. These directives contain general, long-term guidelines, in the form of parametrized role assignments and priorities (e.g., prefer extinguishing fires in region A over fires in region B), that are subject to ongoing elaboration into precise, moment-to-moment activity decisions by the agents [22, 4]. Ideally, following organizational directives should be beneficial when agent directives can be designed that perform well over a range of potential long-term environment and agent characteristics.

3.1 Operational challenges

Without organizational directives, a call center must coordinate with other centers to avoid sending redundant fire brigades to the same fire[6] (using a highest estimated utility protocol to resolve conflicts). Coordination and *retractions* consume valuable time, delaying extinguishing operations. The designed organization only requires coordination if a call center wants to fight a fire outside its responsibility region. When region responsibilities are inappropriate and do not match workloads, fire-brigade *borrowing* requests from overloaded centers increase, again with a loss in performance. When the design is appropriate, retractions are diminished at the risk of more borrowing (as we will demonstrate when we discuss Figure 6) . Call centers must consider all borrowing and loaning

[3] Recent work in open and sociotechnical settings [19, 20] has this emphasis.

[4] There are no non-cooperative agents trying to burn things down. Nevertheless, the cooperative agents sometimes do work at cross-purposes in attaining those objectives (such as all wanting to fight an important fire). This can occur whether the agents are organized or not, because agents have a limited local view of the situation. If unorganized agents did not have the same shared objective as when organized, then some performance gained through organization could stem from the changed objectives. Our assumptions eliminate such a cooperative-objective bonus.

[5] E.g., there is no need to decide if an agent is able to play some role in the organization [21].

[6] In the simulator, every call center receives all fire reports.

options in the context of estimated opportunity costs that are based on potential new fires and uncertainty in the duration of fighting current fires. These are challenging decisions even when agents are well organized.

The call-center agents are highly competent and can make skillful operational decisions to extinguish fires without organizational guidance.[7] Appropriately organized call-center agents, when operating in the sweet spot, should function better than unorganized centers, which must perform consideration of all potential activities and explicitly coordinate them. The organizational complexity in the firefighting system is quite simple. Each call center can perform only two roles: 1) extinguishing fires by directing fire brigades to fight them and 2) loaning fire brigades to another call center. Perhaps counter intuitively, organizational design and control of split roles in homogeneous multiagent systems is more challenging than assigning discrete functional roles to specialized agents in heterogeneous multiagent systems because specialization reduces the space of reasonable choices [8]. The organizational "simplicity" in the firefighting setting means that observed organizational performance differences stem from a relatively small set of organizationally-biased behaviors and are not obscured by complex role and agent interactions.

3.2 Factors affecting organizational performance

We analyzed a number of general factors that influence organizational performance. As these factors change, a designed organization may become highly effective or less effective. In the discussion that follows, we provide an intuitive description of each factor, why it is important, and how it can affect organizational performance. We adjusted each factor individually while holding other environmental settings constant in order to observe its effect on organizational performance independent of the other factors. In total, we conducted a broad analysis that included over 5000 simulation runs with over ten terabytes of simulator output to determine how the general factors of coordination requirements, cooperation benefits, lost opportunity, workload imbalance, and span of control impact the effectiveness of organization. We begin with coordination.

Coordination Requirements Typically, complex tasks performed by multiple agents require coordination, and often a well-coordinated system will perform much better than a system where agents work at cross purposes from only their local, selfish perspectives. In firefighting, coordination is necessary to ensure that call-center agents share responsibility for extinguishing a building only when necessary, and otherwise fight important fires independently (i.e., they do not blindly work on the same fire when more utility could be gained by working on separate fires).

[7] Norms, functions, protocols, etc., are implicitly represented in the plan templates used by these call-center agents. Centers follow these norms (organized or not) and know how to work together to fight fires and share fire-brigade resources.

Coordination is not without associated costs, often involving delays while beliefs, desires, and intentions are communicated. The time required for agents to communicate this information and reconcile it with information from other agents can be significant, especially in cases where agents control resources which must be held in reserve while an agent decides whether it wishes to pursue some goal. Even more significantly, when agents take uncoordinated actions that involve operating in the world, they must deal with the consequences of physically moving resources and then withdrawing them (or having wasted them if they are consumables) once they discover their actions are in conflict with those of another agent. In our analyses, this has been the largest contributor to coordination "cost."

The amount of coordination required is not organization-independent. Organizational directives influence agents to assume specific roles and responsibilities pertaining to certain goals, and assume less responsibility for other goals. The best-case organization for a specific situation would be a perfect partitioning of responsibility regions so that agents select the fires for which they are responsible over those that are the responsibility of others. This ideal situation results in minimal *goal conflicts*, where two agents needlessly pursue the same goal (e.g., extinguish the building at 5$^{\text{th}}$ and Madison). It is important to note that even this organization is not coordination-free, but when each goal is managed and committed to by the agent with the highest expected utility, the committing agent is best suited for reaching out for assistance if necessary. In the context of firefighting, this assistance comes in the form of lending and borrowing fire brigades, an effective remedy for temporal workload imbalances. However, as we will note shortly, excessive resource borrowing leads to inefficiencies in resource provisioning and is often a sign of a more permanent resource imbalance. The worst-case organization (in terms of coordination complexity) would influence every agent to select the same goals (No Org configuration). We analyzed many organization configurations to explore the full spectrum between these two extremes, where organization sometimes cannot prevent agents from selecting the same goals, and at other times, is effective in preventing a goal conflict (which we will also discuss later in conjunction with Figure 6).

This coordination phenomena occurs in firefighting because call centers need to negotiate with other call centers about which fires to fight. In order to come to a resolution for a contested goal, call centers need to compute and share their expected utility with peers. The call center with the highest expected utility will then be responsible for managing fighting the fire, and for borrowing fire-brigade resources from peers if necessary. To investigate the effect of adjusting this coordination cost, we adjusted the *resolution period*, during which call centers reserve resources to fight a fire while waiting for and considering bids from other call centers intent on fighting the same fire. Only after the resolution period has elapsed will the call center with the highest utility commit to fighting the fire. By increasing the resolution period, we increase the cost of coordination while simultaneously making centers more "globally aware" of the utility expected by other agents. By lowering the resolution period, we lower the cost of coordination but make call centers more selfish in that they are less open to considering bids from other centers. Figures 2 and 3, to be discussed shortly, show the effects of "Low-Cost" (short) resolution and "High-Cost" (long) resolution times.

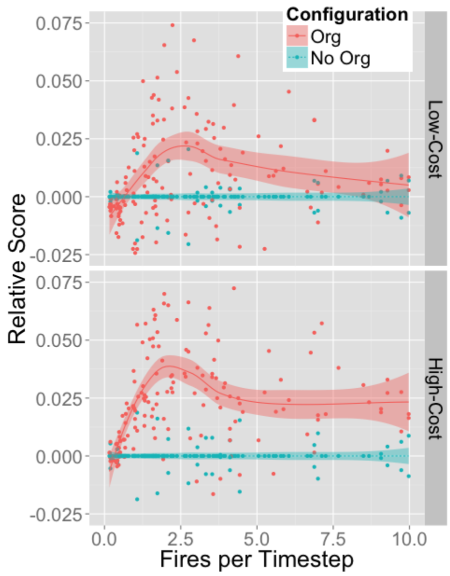

Fig. 2. Varying Coordination Requirements: Score Relative to No Org

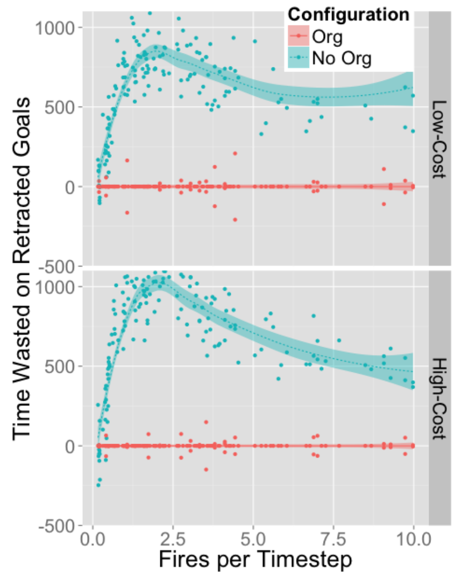

Fig. 3. Varying Coordination Requirements: Cost Effects Relative to No Org

Observation 2: The performance separation of effective organization increases with coordination requirements, without shifting the sweet spot laterally ⇒ We analyzed several organizational designs: 1) a specific responsibility region for each call center (Org) and 2) all centers are responsibible for the entire city (No Org). It seems reasonable to believe that when fires are uniformly distributed, Org would perform best, minimizing goal conflicts while still providing each agent with sufficient beneficial opportunities in its responsibility region. In practice, this is generally true, however, we have found that in cases where, when the conflict resolution period is very short (corresponding to low coordination cost and more selfish agents), the directives supplied to the organized agents do not improve on the No Org baseline. As coordination cost grows, the performance of the organized agents (which need to coordinate less frequently) improves increasingly on the No Org configuration (see Figures 2 and 3).[8]

[8] Figure 3 shows the total retraction time relative to No-Org, which has the most retractions. In both Figures 2 and 3, the 0- and 10-time-steps resolution period results are relative to comparable 0- and 10-time-steps resolution No Org baselines.

Note that with low coordination cost (0-timestep resolution), the difference in performance between the Org and the No Org configuration is only statistically significant within a small window, centered at about 2–2.5 fires per timestep. Correspondingly, the scenario with high coordination cost (10-timestep resolution) achieves a prominent global maximum centered at this time window. From this analysis, it can be seen that when coordination does not incur significant costs, organization is not nearly as beneficial as in cases where coordination (or the absence of needed coordination) is costly. At moderate workload levels, the performance gains afforded by organization reach the maximum. When the simplicity of the scenario does not require coordination, the performance of the Org configuration and the No Org configuration are statistically indistinguishable. Extremely overloaded work scenarios are marked by either statistically indistinguishable performance differences or diminished returns.

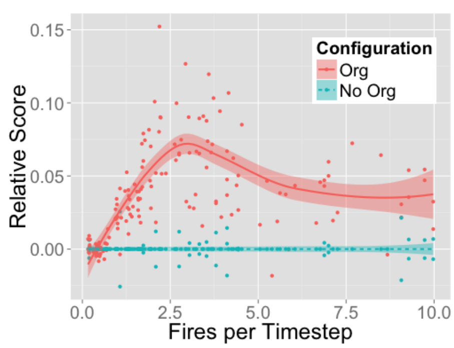

Fig. 4. Relative Score with Twice as many Fire Brigades

Observation 3: Increasing call-center capabilities by adding resources results in a lateral shift and widening of the sweet spot ⇒ The width and position of the sweet-spot window is not fixed, as it depends on the agent's capabilities in servicing goals at either end of the workload range. Call centers become more capable when they have more fire-brigade resources. Figure 4 shows the result of doubling the number of fire brigades controlled by each call center from six to twelve. Now, the organizational sweet spot occurs at a higher workload level: at approximately 2.7 fires per timestep. In addition, the sweet spot is wider as call centers can handle greater task loads before the situation becomes hopeless.

By holding the conflict resolution period constant and varying the number of call centers in the system, we see that coordination complexity is also a function of how "well partitioned" the centers' responsibilities are. In experiments with four call centers, we can see that fewer goal conflicts arise in the Org case than the No Org case. However, if we increase the number of call-center agents to twelve, each with two rather than six fire brigades and responsibility regions that overlap with two other centers, the environmental responsibilities are too precisely partitioned to handle temporal responsibility differences even if, on average over the course of the run, each center's responsibilities are roughly uniform. In Figures 5 and 6, this behavior is reflected in the fact that the number of goal conflicts in the organized, 12-call-center configuration approach the number of conflicts without organization. Correspondingly, the differences in performance between the two configurations are significant. Any advantages to organization under the

4-call-center scenario are lost with the increase in coordination complexity in the 12-call-center scenario. This observation is consistent with the notion that there is an "ideal" number of call centers given the centers' capabilities and the environmental conditions. We do not know for certain that a 4-center organization is the best choice for the environmental conditions that we simulated, but it is certainly better than a 12-center organization, as the 4-center organization provides a better balance between the partitioning of responsibility regions and coordination complexity [23].

Fig. 5. Varying the Number of Call Centers: Relative Score

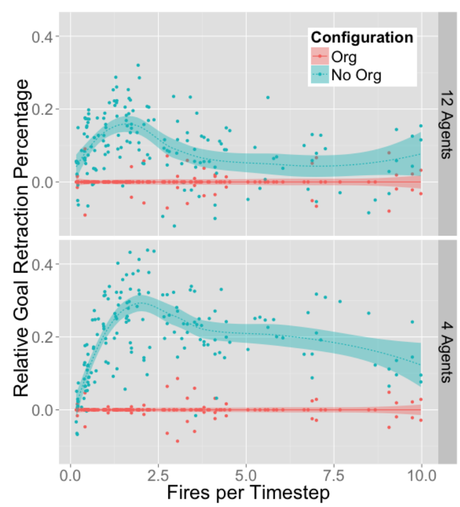

Fig. 6. Varying the Number of Call Centers: Relative Goal-Conflict Rate

Workload Imbalance Organizational directives influence agents to assume responsibility over particular goals and tasks. This reduces the amount of coordination involved in meeting these demands, as there is some expectation of which agent will perform or manage a task. In order for this organizational influence to improve performance, the per-agent workload that is suggested by the organizational directives must be consistent with the distribution of tasks in the environment. Otherwise, some agents have too little work and others have too much. As such, highly beneficial tasks may go without consideration by underloaded agents while overloaded agents struggle to complete all of the tasks they are responsible for. Workload imbalance occurs in firefighting when the distribution of fires throughout the city is not consistent with the size of each of the centers' responsibility region. For instance, if 60% of fires occur in the northwest corner of the city, a partitioning of the city into four equally-sized quadrants would result in a significant average workload imbalance, with the call center in

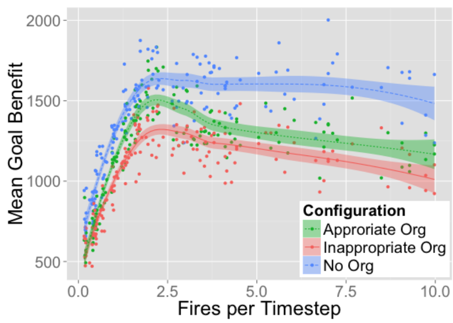

Fig. 7. Varying Workload Balance: Mean Goal Benefit

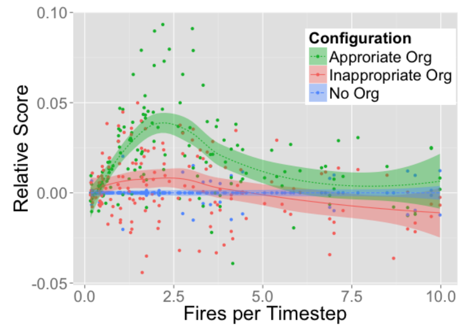

Fig. 8. Varying Workload Balance: Relative Score

the northwest corner of the city having almost six times the workload of other centers.

Observation 4: The performance separation of effective organization increases with increased workload imbalance ⇒ When workloads are imbalanced in this way, call-center agents are not necessarily idle, but instead they work on less beneficial goals. Thus, the penalty occurred by providing these call centers with an inappropriate organization comes in the form of "lost opportunity," where the agent could have performed much more beneficial tasks if it had not been discouraged from doing so by organizational directives. Correspondingly, Figure 7 shows that, as the organizational influences becomes less appropriate, the mean benefit of selected goals becomes lower. A surprising observation shown in Figure 7 is that the No Org case has the highest mean goal benefit of all of the configurations. This is due to No Org agents' preference to selfishly commit to attractive goals which other agents may already be working on, introducing additional goal conflicts and coordination cost.

Observation 5: Extreme workload imbalance, high or low, causes organizationally guided performance to converge to non-organized performance ⇒ On the other end of the spectrum, both Appropriate and Inappropriate Org's less beneficial goals result in a direct lowering of overall score. Figure 8 indicates that this behavior essentially lowers the Appropriate Org curve onto the No Org curve, while still maintaining a window in the workload spectrum where organization is especially advantageous.

Span of Control An important factor in determining if and how agents should be organized is span of control. Simply adding resources (or performers) to a task does not result in constant gain per added resource, and can even result in a net loss of utility. This phenomena is found in many real-world settings [23] where organizations attempt to scale the number of performers without correspondingly scaling management capacity (e.g., hundreds of construction workers cannot be managed by a single foreman). In the firefighting simulator, per-resource effectiveness is diminished above a parameterized call center span-of-control limit.

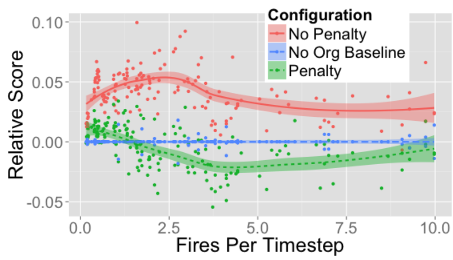

Fig. 9. Span of Control Analysis

Observation 6: Increasing the number of call-center agents beyond what is necessary given their span-of-control capabilities adds coordination requirements (to keep them out of each other's way), decreasing the organizational benefit separation compared to a suitable number of centers ⇒ Span-of-control limits are both important and ubiquitous, since centralization is not generally tractable or realistic. When exceeded in RoboCup Rescue firefighting, performance per brigade is attenuated, counteracting coordination reductions from centralization. Otherwise, one center could handle all brigades. We explored span of control using a configuration where a single call center agent is responsible for managing all 24 fire-brigade resources in the system, but with a span-of-control limit imposed after 6 utilized brigades. Then, we increased the span-of-control capability of the center to 24 to understand how the single call-center agent would perform with no span-of-control limit. We compared these two cases with the baseline configuration where the fire brigades are distributed evenly across four call centers, each controlling 6 of them. Because no call center coordination is needed when there is a single center, in cases where fewer than 6 brigades are needed to execute all of the tasks in the environment, both of the single-agent configurations outperform the multiagent configuration (Figure 9). At a workload level of one fire per timestep, the limited resource effectiveness incurred by the span-of-control penalty becomes more significant than the coordination cost in the multiagent case. Further, since the single-agent case incurs no coordination complexity, there is a noticeable peak in the single-agent configuration without a span-of-control penalty, corresponding to the coordination-cost peak discussed previously.

Observation 7: Coordination requirements that exceed an agent's span-of-control capabilities inverts the performance curve ⇒ Figure 9 shows that the sweet spot obtained when running under the best case scenario of a single call center with no coordination requirements is inverted into a "sour spot" when span of control is considered. Intuitively, the sweet spot is inverted because this is the point in the workload spectrum where it is most important that fire-brigade resources be managed effectively. With span-of-control limits imposed, fire-brigade effectiveness is diminished.

4 MARL Organizations

We next looked for sweet-spot behavior using another previously described and implemented system involving agent organizations. This second system operates in a very different setting: organizing agents that are learning task-assignment policies that optimize service time for tasks arriving in a network [24]. Each vertex in this network is a node, and each edge identifies a potential action. Tasks arrive according to a Poisson distribution, and have variable difficulty (measured as time units) governed by an exponential distribution. Every task is spawned at some vertex v, augmenting agent v's *routing queue* with the new task. Agent v can then decide whether to work on a task locally, adding that task to v's *work queue*, or to forward the task to one of v's neighbors. At every timestep, task at the head of v's work queue is decremented, indicating that it is one timestep closer to being completed. Once the number of remaining timesteps has reached 0, the task is removed from the queue and agent v may proceed to complete the next task in the work queue. Agents receive the inverse of task service time as a reward when a task is completed. To operate effectively in this setting, agents must construct estimates of task service time given locally observable state information such as the size of neighbor's work queues and historical completion times when forwarding tasks.

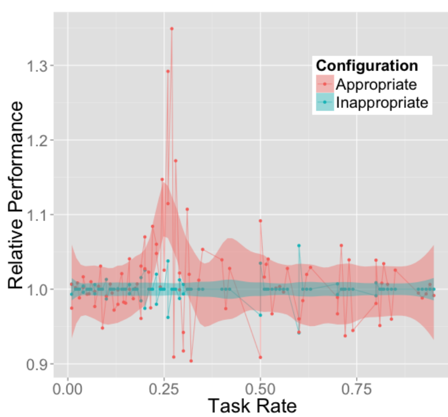

Fig. 10. Relative Performance of MARL Organizations

In this domain, each agent is either a subordinate or a supervisor. Supervisors are responsible for transferring experiences between subordinates that are experiencing similar environmental conditions. Appropriate organizations in this task allocation domain are those that arrange supervisors in a way that exploits similarities between agents. If a group of subordinates frequently experience the same environmental conditions, a great deal of transfer learning can take place. If subordinate groups experience vastly different environmental conditions, transfer learning can occur less frequently, thus not taking advantage of the benefits that organization provides. As in firefighting, an organizational arrangement of supervisors that is appropriate given a particular task distribution may be inappropriate under a different task distribution, so the organization is only effective if the actual distribution is consistent with the expectations assumed in the designed supervisor arrangement.

We used a completely different agent implementation and environment simulator in exploring sweet-spot behavior in multiagent reinforcement learning

(MARL) organizations. For our experiments, we used a 100-agent lattice network and considered two agent organizations. The first organization arranges 4 supervisors such that agents are assigned to supervisors based on their distance from the border of the lattice. The second organization arranges 4 supervisors according to quadrants of the lattice. Tasks are then distributed on the lattice originating from the boundary. Under this model, the former organization is considered "Appropriate" since it partitions agents in a manner that maximizes the similarity of agents in supervisory groups. The latter organization is considered "Inappropriate" since it arranges agents in a way that prohibits effective experience sharing. Given this setup, we experimentally varied problem difficulty by increasing the mean of the Poisson distribution governing task distribution. Evaluation was performed in terms of area under a learning curve (AUC), modeled as an exponential moving average of system-wide task service time. When the system converges more quickly to an optimal policy, the area under this curve will be smaller. To characterize relative performance differences across a wide array of problem difficulties, AUC was normalized relative to the Inappropriate Org configuration.

Observation 8: The MARL system also has a sweet spot $\Rightarrow$ Figure 10 shows more performance variability than occurred with firefighting, but a statistically significant sweet spot arises around a per-agent task rate of 0.25 tasks per timestep. At this workload, the Appropriate Org's performance dominates the Inappropriate Org's. Elsewhere, the two are statistically indistinguishable. The results in the MARL domain are particularly clear. When tasks arrive so frequently that agents cannot compute meaningful policies and the learning process diverges, a supervisor structure that is highly effective in the sweet spot does not help in transferring reasonable policies. On the opposite end of the workload spectrum, when tasks arrive so infrequently that agents do not need to act intelligently in order to service the requests in a timely manner, policy transfer is not important. It is clear from this analysis that even with a completely different set of system dynamics and agent behaviors, an organizational sweet spot exists.

5 Closing Thoughts

Although we have measured and analyzed agent-organization performance under widely varying conditions using only two previously implemented and studied systems (each operating in a different problem domain), we believe that the qualitative behaviors we observed are general and apply to multiagent organizations in *any* domain. We hope our observations encourage those working with more complex heterogeneous agent organizations to investigate and report their performance over a wider range of conditions. Recognizing when a multiagent system will be operating in its organizational sweet spot is helpful in deciding how much effort should be spent in designing and using an agent organization as well as for explaining situations where using an agent organization results in little observed benefit (because the system is operating outside the sweet spot). We have observed that coordination and cooperation amounts, lost work opportunities, and span-of-control capabilities all contribute to sweet-spot performance benefits.

Understanding a multiagent system's organizational sweet spot is important, not just for understanding organizational control opportunity and effectiveness, but when considering if organizational adaptation might be worthwhile [25–27, 12]. Sweet-spot understanding is also important in open, sociotechnical settings when designing an organization (and sizing that design appropriately) for agent recruitment. Identifying where a multiagent system is operating in relation to its organizational sweet spot is important to any discussion or analysis of organizational suitability, performance, or effectiveness.

Acknowledgment This material is based in part upon work supported by the National Science Foundation under Award No. IIS-0964590. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. M. S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):70–80, Jan. 1981.
2. D. D. Corkill and V. R. Lesser. The use of meta-level control for coordination in a distributed problem-solving network. In *IJCAI-83*, pages 748–756, Karlsruhe, Federal Republic of Germany, Aug. 1983.
3. L. Gasser and T. Ishida. A dynamic organizational architecture for adaptive problem solving. In *AAAI-91*, pages 185–190, Anaheim, California, July 1991.
4. E. H. Durfee and Y. pa So. The effects of runtime coordination strategies within static organizations. In *IJCAI-97*, pages 612–618, Nagoya, Japan, Aug. 1997.
5. K. M. Carley and L. Gasser. Computational organization theory. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, chapter 7, pages 299–330. MIT Press, 1999.
6. B. Horling and V. Lesser. A survey of multi-agent organizational paradigms. *Knowledge Engineering Review*, 19(4):281–316, Dec. 2004.
7. B. Horling and V. Lesser. Using quantitative models to search for appropriate organizational designs. *Autonomous Agents and Multi-Agent Systems*, 16(2):95–149, 2008.
8. M. Sims, D. Corkill, and V. Lesser. Automated organization design for multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 16(2):151–185, Apr. 2008.
9. J. Slight and E. H. Durfee. Organizational design principles and techniques for decision-theoretic agents. In *Proceedings of the Twelveth International Joint Conference on Autonomous Agents and Multi-Agent Systems* (AAMAS 2013), pages 463–470, May 2013.
10. A. S. Rao and M. P. Georgeff. BDI agents: From theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems* (ICMAS'95), pages 312–319, San Francisco, California, June 1995.
11. D. D. Corkill, E. Durfee, V. R. Lesser, H. Zafar, and C. Zhang. Organizationally adept agents. In *Proceedings of the 12th International Workshop on Coordination,*

Organization, Institutions and Norms in Agent Systems (COIN@AAMAS 2011), pages 15–30, May 2011.
12. D. Corkill, C. Zhang, B. D. Silva, Y. Kim, X. Zhang, and V. Lesser. Using annotated guidelines to influence the behavior of organizationally adept agents. In *Proceedings of the 14th International Workshop on Coordination, Organization, Institutions and Norms in Agent Systems* (COIN@AAMAS 2012), pages 46–60, June 2012.
13. D. Corkill, C. Zhang, B. D. Silva, Y. Kim, D. Garant, V. Lesser, and X. Zhang. Biasing the behavior of organizationally adept agents (extended abstract). In *Proceedings of the Twelveth International Joint Conference on Autonomous Agents and Multi-Agent Systems* (AAMAS 2013), pages 1309–1310, May 2013.
14. H. Kitano and S. Tadokoro. RoboCup-Rescue: A grand challenge for multi-agent and intelligent systems. *AI Magazine*, 22(1):39–52, 2001.
15. P. Cheeseman, B. Kanefsky, and W. M. Taylor. Where the really hard problems are. In *IJCAI-91*, pages 331–337, Sydney, Australia, Aug. 1991.
16. S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random boolean expressions. *Science*, 264:1297–1301, May 1994.
17. R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic 'phase transitions'. *Nature*, 400:133–137, July 1999.
18. L. Gasser. An overview of DAI. In *Distributed Artificial Intelligence: Theory and Praxis*, pages 9–30. Springer, 1993.
19. L. Sterling and K. Taveter. *The Art of Agent-Oriented Modeling*. MIT Press, 2009.
20. O. Boissier and M. B. van Riemsdijk. Organizational reasoning agents. In S. Ossowski, editor, *Agreement Technolgies*, volume 8, chapter 19, pages 309–320. Springer, 2013.
21. M. B. van Riemsdijk, V. Dignum, C. Jonker, and H. Aldewereld. Programming role enactment through reflection. In *Proceedings of the 2011 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology* (WI-IAT'11), pages 133–140, Lyon, France, Aug. 2011.
22. J. G. March and H. A. Simon. *Organizations*. John Wiley & Sons, 1958.
23. B. Horling and V. Lesser. Analyzing, modeling and predicting organizational effects in a distributed sensor network. *Journal of the Brazilian Computer Society,* special issue on Agent Organizations, 11(1):9–30, July 2005.
24. C. Zhang, S. Abdallah, and V. Lesser. Integrating organizational control into multi-agent learning. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems* (AAMAS 2009), volume 2, pages 757–764, Budapest, Hungary, May 2009.
25. J. F. Hübner, J. S. Sichman, and O. Boissier. Developing organised multi-agent systems using the MOISE+ model: Programming issues at the system and agent levels. *International Journal of Agent-Oriented Software Engineering*, 1(3/4):370–395, 2009.
26. A. Staikopoulos, S. Saudrais, S. Clarke, J. Padget, O. Cliffe, and M. D. Vos. Mutual dynamic adaptation of models and service enactment in ALIVE. In *Proceedings of the Third International Models@Runtime Workshop*, pages 26–35, Toulouse, France, Sept. 2008.
27. T. B. Quillinan, F. Brazier, H. A. Frank, L. Penserini, and N. Wijngaards. Developing agent-based organizational models for crisis management. In *Proceedings of the Industry Track of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems* (AAMAS 2009), pages 45–51, Budapest, Hungary, May 2009.