

Towards Sustained Team Effectiveness

Niek Wijngaards, Masja Kempen, Annika Smit and Kees Nieuwenhuis

Delft Co-operation in Intelligent Systems Laboratory – Thales Research and Technology
Nederland

P.O. box 90, NL-2600 AB, Delft, The Netherlands

{niek.wijngaards; masja.kempen; annika.smit; kees.nieuwenhuis}@icis.decis.nl

Telephone: +31-15-2517867; Fax +31-15-2517801

Track: Organizations

Abstract. Collaboration environments impose high demands on humans and artificial systems. Especially during critical tasks team members, including humans, artificial systems and other (sub-) teams, require support to guarantee their continued effectiveness. Effectiveness of individuals and teams is an important ingredient for organizational effectiveness, managerial decision quality, as well as for maintaining organizational awareness. In this position paper we introduce our conceptual view on realizing sustained team effectiveness, in which both the measurement of effectiveness and team management play an important role. A unified, interdisciplinary approach facilitates measuring effectiveness in more complex organizations.

1 Introduction

Highly dynamic, or even chaotic, environments are often encountered when a disruptive event occurs; consider a car-crash in a tunnel involving a fuel-truck. Suddenly the (tunnel) environment becomes unpredictable, normal courses of action may not yield expected results, and performance of teams as well as individuals is affected. Nevertheless, individuals and teams are expected to effectively address crisis situations over a period of time. In our view, an organization is a collection of teams and individuals, in which effectiveness plays a crucial role, for an overview see e.g. [1]. We loosely define the term effectiveness as the degree to which a team is successful in reaching its goals/objectives. In this paper we focus mostly on individuals and teams, as they provide us with insights which can be translated to organizations.

Teams are often considered to consist of humans, e.g. [2, pp. 126-127] defines a team as "a distinguishable set of two or more people who interact dynamically, interdependently, and adaptively toward a common and valued goal/object/mission, who have each been assigned specific roles or functions to perform, and who have a limited life span of membership". In our opinion, intelligent systems such as agents

and robots can also be team-members, equivalent in status to humans. This is in contrast to a large amount of system-level teams and agent research, which concentrates on agent-based support for individual human team members, see e.g. [3], [4] and [5]. In our view, agents (whether software entities on a network or robots) may also take the initiative and give orders to human (and other agent) team members. In essence, we approach a ‘team’ as an actor-agent community.

An *actor-agent community* is a particular type of organization that involves collaboration of multiple participants, including humans and artificial systems, for the realization of a common mission or for the support of a shared process [6]. Within a community there are social rules that members adhere to, and there is communication, sharing of responsibility and a certain distinct identity among the community members. From a human perspective, an actor-agent community is not unlike any conventional human community – the same traits apply. From a technical perspective an actor-agent community contains distributed systems and processes that have autonomous and anticipatory capabilities – software systems that can be referred to as agents or agent systems. Actor-agent communities are typically involved in complex collaborative decision making processes, such as the day-to-day air traffic management. These are characteristic settings where humans and artificial systems are foreseen to collaborate in the near future.

In our view actor-agent communities need to be able to operate in the real world. This requires the ability to adapt to changes and unforeseen events. There is a need to be able to operate in highly dynamic situations under high degrees of uncertainty. Such a realistic domain is *crisis management*, in which both humans and artificial systems are involved, such as victims, rescuers, observers, and decision makers, all working together to mitigate the situation as quickly as possible [6]. This implies that a team consists of heterogeneous team members of potentially equal status (cf. ‘mixed initiative taking’).

In crisis situations normal operational conditions change radically. A crisis cannot be predicted (otherwise it is not a ‘crisis’), yet preparations can be made (e.g., [7] and see the SEVESO II Directive of the European Community). Although the exact crisis is unknown, certain aspects of a crisis can be identified, e.g. constraints regarding the location of the incident, including availability of resources, victims, geographical setting; team structure, including team members and their skills, team resources such as tools and team norms; culture, as team members may originate from different ‘host’ organizations with different cultural identities, social norms, etc.

Crisis management involves addressing a number of interrelated issues:

- Time-criticality: time is a critical factor in decision-making, de-escalation of incidents, treatment of victims, restoration of normal operating conditions, etc.
- Performance fluctuations: during a crisis, performance of teams and machines changes (usually: degrades), e.g. because of (mental) fatigue, reduced alertness, resource depletion, etc.
- Incomplete situation awareness: dispersed, partial, information about (parts of) the situation, unreliable and faulty observations, etc.
- Changes in courses of action: chaotic environments or situations yield (unpredictable) changes to current goals, plans, schedules, problems at hand, etc.
- Team-(re)composition: during a crisis, teams may be formed, changed, and disbanded.

- Alignment: choreography and coordination of teams require understanding time-critical issues, performance, shared awareness, changes in courses of action and composition of teams.
- Organizational sustainability: organizations need support for their evolution and adaptation at multiple levels of abstraction to address chaotic environments such as crisis situations, while still adhering to necessary levels of coherence and coordination to retain levels of effectiveness.

In this position paper we briefly explore how to support sustained effectiveness of (organizations consisting of) teams consisting of heterogeneous team members, including humans, artificial systems and other (sub-) teams, in crisis situations. Section 2 addresses our view on sustained team effectiveness. Section 3 describes our initial model for team management, which is a key element in realizing team sustainability. Section 4 proposes a number of future research directions.

2 Sustained Team Effectiveness

Sustained team effectiveness is, from our perspective, a basic team property, which plays an important role in the usefulness, robustness, employability and composition of teams over a period of time. We explicitly assume that the effectiveness of individuals, teams and organizations changes over time – which may be manageable to a certain extent. In this section we describe our views on effectiveness and sustained effectiveness.

A good source for research on effectiveness stems from Psychology, in which an individual's effectiveness is often termed 'performance'. The level of performance is stated in terms of behavioral measures such as reaction times and errors (false positives and false negatives). Of major importance is the individual fitness level, which is often labeled as 'vigilance' [8]. With respect to vigilance, it is not only important to measure general performance levels (e.g. average performance), but also the changes in performance over time are critical. That is, a decrement in performance (even when the average performance level is not very low) is indicative of declining vigilance of an individual [9], [10]. Whenever a vigilance decrement sets in, the deterioration is bound to get worse in the near future. So, for an individual, vigilance is determined in the following manner: performance (reaction times, errors) and performance change over time.

Interestingly enough, independence between physical and mental aspects of vigilance has been demonstrated [11]. That is, after demanding mental tasks, the mental component of vigilance is declined. Hence, subsequent mental task performance is decreased. However, purely physical tasks can still be performed adequately. The opposite also holds true: After strenuous physical tasks, performance on subsequent physical tasks is worse, but performance on mental tasks is generally not affected.

Besides the more basic issue of vigilance, other individual human variables also determine performance to a certain extent. Skills and expertise are known to have a considerable impact on performance [12]. Even when a person is highly vigilant, he

or she may not be able to perform well if faced with a totally unfamiliar situation. For example, a skilled air-traffic operator would probably not know how to extinguish a fire, not even in a vigilant state.

Furthermore, especially in a crisis domain, performance is not only dependent on intrinsic aspects, but is also dependent on other aspects such as tools and external resources. For example, a firefighter may be both mentally fit (good and stable performance on mental tasks) and physically fit (good and stable performance on physical tasks), but if he does not have a proper suit or a full fire extinguisher, he will not perform as expected. Related to both aspects is whether the individual is in a relevant situational context: Is the right person at the right place at the right time? A vigilant, highly skilled and fully equipped fire fighter still needs to be near a fire in order to be able to extinguish it.

The abovementioned research illustrates a separation of physical and mental tasks. This is not so surprising, because physical systems, such as a human, operate within different dynamical regimes under normal conditions. Physical tasks, such as scrolling, and mental tasks, such as writing a paper, are assigned to different subsystems of the brain. In this case one could argue that the subsystems are most likely completely independent, meaning that the performance of one subsystem carrying out the mental task is not affected by the subsystem carrying out the physical task. In sum, these tasks do not interfere with one another. However, in real, more complex settings such as air traffic control operation, this independence of subsystems in the brain that carry out physical and mental tasks might not be plausible, due to the complexity of the cognitive tasks that need to be performed. Here, the physical and mental components might actually be coupled.

An example of coupling is the writing of a research paper, where the physical effort of clicking, scrolling and moving your head to read comments in the margin of a research paper interferes with the high-level mental activity of creating a nice piece of text that aims at getting your message across. Moving your head from the text field to the margin field – a physical task – involves a mental task, namely that of conscious motor control of neck and eye muscles. This places a burden on the other ongoing mental task: the creative process of writing the paper. Both physical and mental subsystems obstruct each other in such cases, due to their (un)fortunate coupling. Concluding, mental and physical task alignment, ensuring a highest possible independence of both subsystems, is first of all a pre-requisite for avoiding stress and strain, as well as mental and physical overload. Secondly, optimal alignment should sustain good task performance, as both sub-systems will minimally interfere with each other. From a cognitive ergonomic perspective it is a challenge to provide an empirical account of when, why and how such alignment is achieved for which subsystems. Furthermore, it is interesting to embody psycho-physical findings in future applications, systems and teams [13].

In short, effectiveness is dependent on physical and mental aspects of an individual, its tools and resources, cognitive ergonomic design factors, and situational context. Part of our future research is to investigate whether it is possible to measure and maintain effectiveness of teams of humans, software agents and robots. Our initial stance is that these constituents of (human) effectiveness are applicable to artificial entities. For example, consider vigilance; although vigilance may not seem directly applicable to a software agent, its processing capacity and supportive

resources (e.g., memory, electricity, fuel, and bandwidth) may deplete over time. This corresponds with one of the definitions of human vigilance, which refers to the availability of capacity or resources [8]. The use of vigilance as one of the effectiveness parameters thus appears justified for both actors and agents. Despite the fact that similar principles apply, there of course remain structural differences between actors and agents. They will differ in actual task performance; quality and speed of performance will not be identical on all tasks. However, this will not interfere with our statement that their effectiveness can be *measured* in a similar manner.

We suggest that the fitness level of human *teams* can be measured in a similar manner as on the individual human level. That is, we intend to apply the same criteria for the team level as defined for the individual level. A unified approach facilitates measuring effectiveness in more complex organizations, including teams consisting of individuals and (sub-)teams. Determining the fitness of individuals and teams also requires a mechanism that monitors the actions/progress of individuals and teams in relation to their joint goals. E.g., in [14, p. 1] “Team situation awareness involves the team's assessment (i.e., perception, comprehension, and projection) of the current situation, which can include the surrounding environment (including any equipment or systems), the task, and the team itself.” We would like to extend this approach from human teams to ‘hybrid’ teams (consisting of actors and agents (software and/or robots), which can be monitored in a similar fashion to measure effectiveness.

An important issue concerning measurements is that sustained team effectiveness is a desired high-level emergent property of a (possibly ad-hoc or changing) team or organization. This implies that the effectiveness of individuals and teams needs to be *measured* in time-critical circumstances. The measurements of effectiveness need to be such that they can be used in time by the right actors, agents or teams to manage themselves and others. Being able to measure effectiveness implies both attaching values to ‘effectiveness’, as well as having methods to determine these values. In addition, to guarantee a specific level of effectiveness, a number of ‘levels’ of effectiveness need to be distinguished. The relationship between required (minimal) effectiveness to cope with a certain class of problems in specific situations needs to become explicit as well.

Sustained effectiveness basically entails managing team effectiveness over time according to some criteria. Unfortunately, changes to team effectiveness are neither easily planned nor predicted, as the environment is essentially unpredictable and the criteria are not easily determined. In addition, time-criticality requires pragmatic approaches to measuring effectiveness of individual humans and agents as well as teams, and establishing pragmatic criteria (including a desired minimal level of effectiveness). Nevertheless, methods and techniques are required to positively influence effectiveness of individuals and teams. In the next section we elaborate on effectively managing ad-hoc and dynamic teams.

3 Sustained Team Management

Sustained effectiveness, by nature, is an emergent property (which may or may not manifest itself), of both individuals and teams. Sustained effectiveness even manifests at the organizational level. The Aisin crisis in the Toyota Group [15] illustrates the possibility that without direct high management control, sustained effectiveness is feasible, even in a large crisis situation. Sustained effectiveness may be measured, to some extent, but cannot be directly controlled: it is not a simple ‘parameter’ of any human or artificial system. In organizational literature, ‘management’ is responsible for organization effectiveness, e.g. see [1]. We currently assume that at each level in the organization, team management is concerned with measuring and influencing effectiveness. In this section our initial model for sustained team management is described. First, a conceptual model is introduced, after which team management and its relation to sustained effectiveness is elaborated in some detail. This section is concluded by a discussion regarding measuring sustained effectiveness.

3.1 Towards a Model

To study sustained team effectiveness, a model is needed which represents teams, their composition, and the role of sustaining effectiveness. Our simple initial model is shown in Figure 1, in which team composition and team context are outlined. A team is a compositional construct, which can be composed of any combination of specific individuals (actors and/or agents) and sub-teams. In this compositional approach, no strict hierarchy is enforced: a team may be sub-team of multiple other teams, and similarly an individual may be a team-member of multiple teams. With this approach we intend to model dynamic organizational structures – although this needs further research to understand the limitations. We currently assume that it is not necessary to have a ‘top-most’ team which has a compositional relation with every individual in the entire organization. Similarly, we do not enforce that, at the ‘bottom’ of composition relations, teams can only consist of individuals; we leave open the possibility that a team ‘believes’ that it has a sub-team as a team-member, but this sub-team is not (yet) populated with individuals. How a sub-team is represented as a team-member is also part of future research.

In the model depicted in Figure 1, the composition of one team is shown. The team consists of two individuals (an agent and a human actor) and two sub-teams (which may again consist of actors, agents, and/or sub-sub-teams). A team has an associated ‘team management’ process, which is shared by all team-members. At this point, it is not important how this management process is realized by specific team-members; multiple options exist and will be explored in future research. On the one hand, team management is responsible for real-time dynamic (re)scheduling of allocation of tasks to humans and agents in order to solve problems given certain goals. On the other hand, there is the responsibility for acquiring and analyzing relevant status information of all members including their effectiveness, and changes of their effectiveness over time. These two responsibilities overlap; effectiveness is also a basis for (re)scheduling task allocations. For example, the fitness or vigilance that is determined of individuals and sub-teams serves as a basis for task allocation. An air-

traffic operator who is very tired (i.e. not vigilant) will not be allocated a task that requires high vigilance. Likewise, a highly effective sub-team will be required to take the responsibility of a high-priority, demanding set of tasks.

A team operates in a specific environment. Goals and issues that are relevant for the situation at hand are common to all team members: the well-known joint or shared goals, commitments, intentions, etc. The shared environment, in turn, is part of the entire (global) environment. Within the global environment resides an entity (or multiple entities) that fulfill(s) the role of ‘external manager(s)’; this is explicitly separated from the (internally shared) team management process. It is assumed that each team and each individual is capable of understanding and reacting to internal and external management directives and needs for reports. Each individual or team may also fulfill the role of ‘manager’ for another individual or team, possibly external to its own team. So the management-relationships do not have to adhere to a specific team-composition structure. In addition, note that team management processes need not be heavy-weight; it may be virtually non-existent – which places a burden on external managers to manage the internals of those teams.

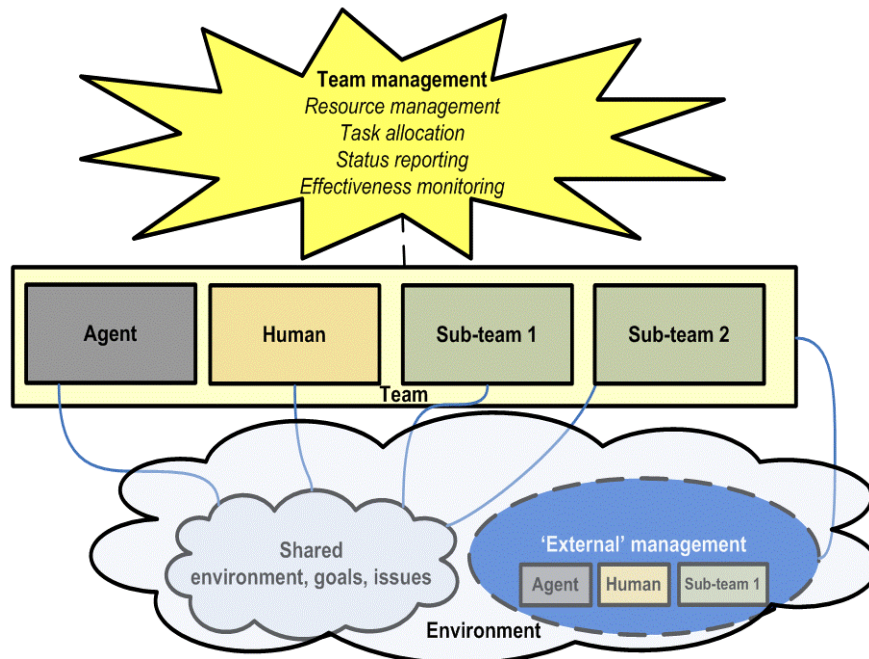


Figure 1. A conceptual model of sustained team effectiveness.

3.2 Management Scenarios

In this sub-section three example scenarios illustrate the application of our model to sustained effectiveness of individuals and teams. In these examples, the management process related to a team is shown to involve sub-processes for strategic deliberation,

task allocation, team-composition, choreography and coordination, dealing with external managers and maintaining organizational synergy. In essence, a (team) management process entails a feedback loop, in the simplest form involving team and situation monitoring, deliberation and action effectuation (e.g., compare with the well-known OODA loop of decision making cycles: observe-orient-decide-act). The team as a whole follows specific strategies to deal with issues arising from the current crisis within their shared environment. The initial strategy can be an arbitrary choice, but is nevertheless constrained by the resources available and possible rules and regulations imposed by the host organization(s) of the team-members. Team management keeps track of the available resources, task allocations and continually monitors the status of individual entities' performance, team goals, and information from and reports to external management.

For all scenarios, assume that a team has already been formed and currently addresses (part of) a crisis situation. The focus is on team management deliberations, irrespective of how the team management process is realized (by one or more team members). All team-members are currently assumed to be cooperative, reliable, and have non-conflicting norms, etc.

In the first example scenario, team management aims at sustaining effectiveness of its own team members. The crisis situation places a heavy burden on all team members, and the overall effectiveness is slowly but surely decreasing. A simple, opportunistic management strategy to maintain an overall acceptable level of effectiveness within the team is to allocate new tasks directly to those team members that have the most suitable vigilance level, without any explicit negotiations or inquiries about status. This opportunistic task allocation strategy may be particularly appropriate in time-critical circumstances. In addition, the effectiveness of team members may be (in)directly influenced. For example, current measurement information shows that one human team member in particular is showing signs of decreased vigilance. It is a matter of strategic deliberation when the right moment arises to give this human team member a 'bogus' high priority task which involves taking a rest break, and the human's other tasks can be postponed or delegated to other team members. Another predictable decline in team effectiveness caused by resource depletion is addressed by charging a team member with an extra task to acquire additional resources. As an illustration of differences between humans and artificial team members, consider an (artificial) agent team member, whose effectiveness in the current situation is too low to be acceptable. As its skill set is deemed insufficient for the current task and problem, this software agent team member is required to update its knowledge-bases and acquire additional functionality to increase its effectiveness immediately, as other team members depend on the performance of this specific team member.

In the second example scenario, team management addresses the issue that team effectiveness cannot be increased by influencing its current team members; other team members need to be acquired, and possibly a number of current team members need to be removed to prevent unnecessary complexity in the team's organizational structure. Prospective team members may for example be found by means of brokers (specialized actors or agents who trade in team members, e.g. external managers may fulfill this role as they may have more complete organizational awareness), by

proximity in the environment (an opportunistic approach), or by searching in databases (yellow pages) and by ‘word of mouth’.

In the third example scenario, an external manager coordinates collective actions of multiple teams. For example, the team management of the current team (team 1) considers its team to be sufficiently effective, and reports this to its external manager. However, the external manager (which may be a team by itself) is also informed by another team (team 2) that team 2 is currently significantly reduced in effectiveness because of a skill-set deficiency. Based on information concerning the locations of teams 1 and 2, the external manager decides that team 1 may be in a favorable position to fulfill the request of team 2. Team 1 is contacted by the external manager, and is provided with information on team 2, its predicament, and is ordered to select a suitable team member to move to team 2.

Each of these example scenarios involves a monitoring process. Maintaining current and accurate information on the local, intra-team situation as well as on (relevant parts of) the environment, is a non-trivial issue. This issue is compounded by the need of external management for reports on the effectiveness and other aspects of teams and individuals. Criteria need to be established for information distribution, e.g. on a ‘need to know basis’, to prevent both information overload and a higher likelihood of timely arrival of task-relevant information. In [16] an information distribution system is proposed which is capable of providing an actor or an agent with task-relevant information. With such a system in place, necessary information and/or knowledge are reported back and the team’s strategy is adapted accordingly.

In addition, sustaining team effectiveness ultimately implies measuring the effectiveness of the entire organization involved in the crisis resolution. Measuring effectiveness on an organizational scale means that an external manager is aware of the entire organization. The design of an information system that supports the exchange of organizational information is presented in [17]. With our model we intend to provide a scalable approach to time-critical monitoring and control in ad-hoc, dynamic, organizations.

3.3 Measurement of Sustained Effectiveness

The above examples illustrate the role of measuring effectiveness of individuals and teams over time. Although time can be measured to a certain degree of accuracy in a distributed system, measuring effectiveness remains a challenge, especially when the measurement techniques need to be applicable in time-critical, resource-sparse organizations. A major challenge is how to come up with the right indicators that determine whether the performance of the individuals, teams, and the overall organization is optimal. A first step is to devise performance measurement techniques at the individual and team level. In addition, the contextual sensitivity of effectiveness needs to be addressed; measuring effectiveness is dependent on tasks, goals, situation at hand, available individuals, teams, tools and resources, etc. Note that this is in contrast with vigilance, which can be measured out of context as it reflects, for humans, a basic energetic level of information processing capacity [8]. However, the impact a certain level of vigilance will have on performance, depends on the actual

task that needs to be performed. It is to be expected that the more specific the context becomes, the accuracy of the measurement of effectiveness increases.

The techniques for measurement should have a minimal impact on the effectiveness of individuals and teams, otherwise they defy their purpose. An example technique to determine the fitness or vigilance level of an actor involves measuring behavior. This may be accomplished on an objective task performance level, but also on a subjective scale by means of (short) alertness questionnaires. It has been shown that these subjective measures can be remarkably accurate as they correlate highly with sensitive vigilance measures such as brain responses measured in the EEG and they distinguish between mental and physical components of vigilance [8]. Although vigilance can thus be determined with a short alertness questionnaire and with performance measures on tasks, these tasks should not be intrusive and should be naturally integrated in the work environment (such as checking incoming mail for an operator or answering the mobile phone for a police officer). An open question remains how to measure “team or organizational vigilance” in chaotic and dynamic environments.

We explore the practical uses of measuring effectiveness in time-critical, chaotic situations. Merely measuring the time it takes before a police officer answers his mobile phone does not yield sufficient information. An officer who responds rather slowly might still be vigilant, but busy aiding civilians. Of equal importance is the fact that a police officer may not be able to answer to a questionnaire in a full-blown crisis situation. It is probably necessary to develop two (or more) styles or modes of measurement. The first mode entails performance and subjective measures that may include tasks that are not strictly part of the work environment itself (such as questionnaires). The second mode involves an aggregation of ‘naturalistic’ work-related tasks.

With respect to effectiveness of individuals and teams, it is important to distinguish levels of effectiveness on different dimensions. There are various facets of effectiveness (fitness of the individual, of a team, (higher-order) goal satisfaction, etc.) and these facets together determine the overall effectiveness. That is, effectiveness should be determined in a sensitive manner (not just high or low). These different aspects or dimensions should be explored and determined in order to gain proper insight into why and how overall team effectiveness may be lacking.

Moreover, effectiveness should not be expressed as a single value; it is a multi-faceted concept, which loses expressiveness (and comparative usage) when used as a single value. This is similar to valuations of the ‘trust’ concept. Trust as a single-value does not help to differentiate between e.g. the belief in an agent’s competence versus willingness: important aspects for decision making [18]. To continue the comparison with ‘trust’: The trust in another agent’s capacity to fulfill a certain task is based on aspects of the task and the other agent. Similarly, we expect effectiveness to consist of a number of aspects, which can also be determined (perhaps to some extent) for individuals, teams, problems, situations, resources and their relations. It then becomes possible to, for example, use aspects related to a specific problem to specify necessary levels of effectiveness that have to be met by a specific team. Differentiation in aspects is also expected to facilitate prediction of change in effectiveness, e.g. by task completion or management actions.

A potential problem in using the same aspects (or concepts) to measure effectiveness of both actors and agent concerns the abstraction level of descriptions. One of the problems in the interaction of humans with computers (e.g., agents) is the fact that computers are described at a much lower level of abstraction than humans, who are supposed to work with them [19]. In order to realize actual collaboration between actors and agents within a community, it is necessary to describe their behavior in a situational context dependent manner and on related levels of abstraction – even if this involves additional antropomorphisation of agents. Another problem is that there is no a priori reason to suppose that the interaction between actors and agents can be abstracted to a similar enough level, because of their inherently different technological make-up.

4 Future Research

This position paper outlines our progress towards sustained team effectiveness – in our opinion a basic element in managing team performance over time. The approach and issues presented manifest in at least three interrelated and interdisciplinary research projects regarding possible influences of the emergent aspect of sustained team effectiveness. These three research projects are intended to further refine our model for team management within the domain of crisis management in the context of the interdisciplinary ICIS research program on interactive collaborative information systems. Multiple heterogeneous individuals are incorporated, including actors, agents, and robots.

The first research project concerns instruments for enhancing effectiveness within one team. An example is task allocation on the basis of vigilance levels of human actors and artificial agents. The second research project involves instruments for enhancing effectiveness of teams on the level of team formation. This project involves team formation and re-composition, with many types of entities (actors, artificial agents) and roles. The third research project focuses on increasing effectiveness on an ‘inter-team’ level, that is sustaining effectiveness over multiple teams distributed in the environment. A major challenge is to address the combination of increasing scale of organization size and crisis escalation together with the need for time-critical information flows.

Our approach to modeling and analyzing team management in dynamic organizations needs to be related to other, existing, approaches in literature and practice: a major component of our future research. For example, the well-known model for agents and teams [20], STEAM and its successor Machinette, provide support for communication and coordination within a team, where agents support humans and intend to maximize overall team utility. Our approach is an extension, both in the autonomy of the agents (from supportive to equivalent team member) and in the management of effectiveness of individuals and teams. In addition, our approach needs to be compared to other research disciplines regarding sustaining effectiveness (or utility) of teaming of humans and artificial systems; example research areas include cybernetics [e.g. 21], psychology [e.g. 22], cognitive ergonomics [e.g. 23] and robots [e.g. 24].

In addition, the impact of a number of important aspects in the design of distributed systems needs to be investigated. For example, the role of trust in relation to measuring effectiveness in specific, and management processes in general; interoperability in terms of communication, understanding, norms and culture; and security, privacy and malicious intent.

Acknowledgements

This work is supported by the Dutch Ministry of Economic Affairs, grant nr: BSIK03024. The ICIS project is hosted by the DECIS Lab (<http://www.decis.nl>), the open research partnership of Thales Nederland, the Delft University of Technology, the University of Amsterdam and the Netherlands Foundation of Applied Scientific Research (TNO).

References

- 1 Baker, K. A. and Branch, K. M. (2002). Concepts Underlying Organizational Effectiveness: Trends in the Organization and Management Science Literature, In: *Management Benchmark Study*, Chapter 1, Office of Planning & Analysis, Department of Energy, USA, <http://www.sc.doe.gov/sc-5/whatwedo/com-benchmark.html>.
- 2 Salas, E. Dickinson, T. L., Converse, S. A., and Tannenbaum, S. I. (1992). Toward an understanding of team performance and training. In: Swezey, R.W. and Salas, E. (Eds.), *Teams: Their training and performance* (pp. 3-29). Norwood, NJ: Ablex.
- 3 Sycara, K. and Lewis, M. (2004). Integrating intelligent agents into human teams. In: Salas, E. and Fiore, S. (Eds.) *Team Cognition: Understanding the Factors that Drive Process and Performance*, Washington, DC: American Psychological Association.
- 4 Lenox, T., Lewis, M., Roth, E., Roberts, L., Shern, R. and Rafalski, T. (1998). Support of Teamwork in Human-Agent Teams. In: *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*, Oct 12-14, San Diego, CA, pp. 1341-1346.
- 5 Lewis, M. and Sycara, K. and Payne, T. (2003). Agent Roles in Human Teams. In: *Proceedings of AAMAS-03 Workshop on Humans and Multi-Agent Systems*.
- 6 Wijngaards, N., Nieuwenhuis, K. and Burghardt, P. (2004). Actor-Agent Communities in Dynamic Environments. In: *Proceedings of the workshop ICT Agents*, November 2004, TNO Defence, Security and Safety, The Hague, The Netherlands.
- 7 Lin, Z. and Carley, K. (2001), Organizational Design and Adaptation in Response to Crises: Theory and Practice, In: *2001 Academy of Management Best Papers Proceedings*, 2001: B1-B7, Washington, D.C.
- 8 Smit, A. S. (2004). *Vigilance or availability of processing resources: A study on cognitive energetics*. PhD thesis, Radboud University Nijmegen.
- 9 Smit, A. S., Eling, P. A. T. M., and Coenen, A. M. L. (2004). Mental effort causes vigilance decrease due to resource depletion. In: *Acta Psychologica*, 115, 35-42.
- 10 Parasuraman, R. (1985). Sustained attention: a multifactorial approach. In: Posner, M. J. and Marin, O. S. M. (Eds.), *Attention and performance* (492-511). New York: Erlbaum.
- 11 Smit, A. S., Eling, P. A. T. M., and Coenen, A. M. L. (2003). Mental and physical effort have diverse effects on vigilance. In: *Journal of Psychophysiology*, 17 (Suppl. 1), S64.
- 12 Parasuraman, R., and Davies, D. R. (Eds.) (1984). *Varieties of attention*. Academic Press, Inc, Florida.

- 13 Salden, A. H. and Kempen, M. H. (2004). Sustainable Cybernetics Systems - Backbones of Ambient Intelligent Environments, In: Remagnino, P, Foresti, G. L. and Ellis, T. (Eds.), *Ambient Intelligence- A Novel Paradigm*, Springer, November 2004.
- 14 Cooke, N. J., Stout, R. and Salas, E. (2001). A Knowledge Elicitation Approach To The Measurement Of Team Situation Awareness. In: McNeese, M., Endsley, M. and Salas, E. (Eds.). *New Trends in Cooperative Activities: System Dynamics in Complex Settings*, pp. 114-139. Santa Monica, CA: Human Factors.
- 15 Nishiguchi, T. and Beaudet, A. (1997). Self-Organization and Clustered Control in the Toyota Group: Lessons from the Aisin Fire. *Report #w-0167a, Massachusetts Institute of Technology International Motor vehicle Program*.
- 16 van Someren, M., Netten, N., Evers, V., Cramer, H., de Hoog, R. and Bruinsma, G. (2005). A trainable information distribution system to support crisis management. In: Carle, B. & Van de Walle, B. (Eds.) *Proceedings of the Second International ISCRAM Conference*, Brussels, Belgium, April 2005.
- 17 Oomes, A. H. J. and Neef, R. M. (2005). Scaling-up support for emergency response organizations. In: Carle, B. and Van de Walle, B. (Eds.) *Proceedings of the Second International ISCRAM Conference*, Brussels, Belgium, April 2005.
- 18 Wijngaards, N.J.E., Boonstra, H.M. and Brazier, F.M.T. (2004). The Role of Trust in Distributed Design, In: Brazier, F.M.T. and Gero, J.S. (Eds.), *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, volume 18, Special Issue on Intelligent Agents in Design, pp. 199-209.
- 19 Wilpert, B., and Qvale, T. (1993). *Reliability and Safety in Hazardous Work Systems: Approaches to Analysis and Design*, Lawrence Erlbaum Associates: East Sussex, UK.
- 20 Schurr, N., Okamoto, S., Maheswaran, R. T., Scerri, P. and Tambe, M. (2004). Evolution of a Teamwork Model. In: *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*, Cambridge University Press.
- 21 Wolpert, D. and Tumer, K. (1999). An Introduction to Collective Intelligence. In: *Handbook of Agent Technology*, AAAI Press/MIT Press.
- 22 Rasker, P. (2002), *Communication and Performance in Teams*, PhD Thesis, University of Amsterdam
- 23 Klein, G. Woods, D. D., Bradshaw, J. M., Hoffman, R. R., and Feltovich, P. J. (2004). Ten Challenges for Making Automation a "team Player" in Joint Human-Agent Activity, In: *IEEE Intelligent Systems*, Vol. 19, No. 6, November/December 2004.
- 24 Nourbakhsh, I., Sycara, K., Koes, M., Yong, M., Lewis, M. and Burion, S. (2005). Human-Robot Teaming for Search and Rescue, *IEEE Pervasive Computing: Mobile and Ubiquitous Systems*, January, 2005, pp. 72-78.