14th ICCRTS

”C2 and Agility”

COMMAND WITHOUT COMMANDERS

Topic 5: Experimentation and analysis

Berndt Brehmer

Point of contact:

Berndt Brehmer

Department of Military Studies
Swedish National Defence College
P.O. Box 27805
SE-115 93 Stockholm
SWEDEN

Telephone: +46 8 55342837

Fax: +46 8 55342600

e-mail: berndt.brehmer@fhs.se

COMMAND WITHOUT COMMANDERS

Berndt Brehmer
Department of Military Studies
Swedish National Defence College

Analyses of command as a function, defined as *focus and convergence* by Alberts (2007) and *direction and coordination* by Brehmer (2007), suggest that this function can be fulfilled without a commander. The most radical alternati-ve is self synchronization. This paper presents a general paradigm for the study of self synchronization which can be implemented in the laboratory as well as in the field. It lists the variables that constitute the context in which command should be studied, variables that then need to be operationalized in actual studies. Here it the paradigm is embodied in a computer simulation, which enables us to study and measure theextent to which a number of persons selfsynchronize to achieve a goal (the application is fire fighting). The results of a series of experiments with the paradigm show that self synchronization can be demonstrated, that it is more effective that traditional forms of hierarchical command under time pressure, and that it is enhanced by free communication in a network as well as by "blue force tracking". The results suggest that the network works "as advertised" and that the paradigm is useful as a means to study the effects of different factors on self synchronization and to achieve a cumulative body of research.

In our recent analysis of command and control (C2) it is defined as the function that provides *direction and coordination* for military operations (Brehmer, 2007). Alberts (2007) makes a similar point, but uses a different terminology, when proposing that command achieves *focus and convergence*.

The important point here is, however, not what command is supposed to achieve but the view that it should be analyzed as a *function* rather than as an *activity*. This opens up new possibilities for thinking about how direction and coordination in a military operation can be achieved. The traditional answer: by a commander, is only one possibility. Alberts (2007) reaches the same conclusion when he emphasises the need to distinguish between command as a verb and command as a substantive. When we consider command as a verb, as we usually do, the question of who commands raises immediately, and we are stuck with the traditional commander-centric view of C2. Thinking of command as a substantive, on the other hand, which amounts to thinking about it as a function, opens up new possibilities.

One of these new possibilities is suggested by the concept of self synchronization. It points to a new direction in which we could look for an answer to the question about how direction and coordination is achieved, for if self synchronization is possible, it is truly a matter of a force achieving direction and coordination without a commander.

Self synchronization is one of the supposed benefits from the new Network Based Defence (NBD) now being implemented by many defence forces. It is a possibility often mentioned but seldom explained in any detail. Nor have there been many empirical demonstrations of the phenomenon, despite its central role in NBD. One reason for this may be that we do not have a conceptual framework for analyzing self synchronization, and that we have lacked a way of operationalizing the problem for research. In this paper, I will try to remedy these problems by proposing a general paradigm for the study of self synchronization and present an operationalization of that paradigm that allows us to study the phenomenon empirically. Finally, I will describe four experiments which show that self synchronization is possible as well as some of the variables that affect it, thus illustrating that command can indeed be achieved without a commander.

What is self synchronization?

Alberts, et al. (1999) define self synchronization in the following way:

> *One example of this highly decentralized C2 calls for lower-level decisionmakers to be guided only by their training, understanding of the commander's intent, and their awareness of the situation in relevant portions of the battlespace. In some variants of the concept there is provision for management by exception (i.e., the commander can negate lower-level decisions on an exception basis).* (Alberts, et al., 1999, p. 219)

This definition seems to me to say both too much and too little. It says too much in that it includes factors that are presumably preconditions for self synchronization, such as training and situational awareness. It says too little in that it does not define the phenomenon in a way that allows us to distinguish it from other phenomena. For research purposes, it is best to keep the preconditions, i.e., the theory, separate from the definition of the phenomenon. If we do not, we may end up with a vicious circle when we try to test our theories and demonstrate the phenomenon. Here, we attempt that by starting with an operational definition of self synchronization, i.e., a definition in terms of what needs to be observed to have observed the phenomenon:

> *Selfsynchronizing is observed when a number of units achieve the direction and coordination necessary to handle a mission without a commander doing the directing and coordinating.*

That there is no commander to provide direction and coordination means that the units must be guided by their understanding of what their mission requires of them in the situation at hand. This understanding must involve not only what their own unit should do, but also what it should do in concert with the other units in the situation.

We may think of self synchronization as the result of the meeting between a task and the competence that the various units involved have to handle that task. An example will clarify what this means. The example involves fighting forest fires, and it is, admittedly, an abstract view of fire fighting, more suited to our problem of developing a paradigm for self synchronization than as a guide to actual fire fighting, but I ask the reader to bear with me here[1].

---

[1] Some readers may have preferred a military example, but that would have been unnecessarily complex, and not really necessary since self synchronization as a phenomenon is not tied to military circumstances. It is a phenomenon that may, or may not take place in a force with different units and no commander. The military context is just one on many contexts where it can occur. There is, however, considerable advantages in limiting the initial discussion and studies to a situation where the "enemy", i.e., the fire in this case, behaves in a more predictable fashion than the enemy in a military context would do, and then move on to the more complex military case. The self synchronization phenomenon itself is most likely the same in both fire fighting and military circumstasnces, but the problem that must be handled by a selfsynchronizing force is, of corse, different.

Consider a forest fire. Four fire fighting units (FFU) arrive at the fire. Neither of the units can handle the fire on its own, so their mission requires coordination of their efforts. However, no FFU commander has the authority to take command of all four units and coordinate their activities. Moreover, each unit commander can see only part of the fire. If the units manage to put out the fire, it means that they have been able to coordinate their efforts, for without coordination, they could not have handled the fire. This is our paradigm example, or defining case, of self synchronization. The research question is what is required for a force of separate units to achieve this

To successfully attack this question, we need a paradigm, a conceptual framework that makes it possible to discuss the problem in general terms and to include specific cases in a common framework which defines the essential characteristics of the self synchronization phenomenon.

What is a paradigm?

Dictionary.com Unabridged defines paradigm as an example serving as a model pattern. While originally a concept used in grammar, it has obtained a wider application especially through the work of Thomas Kuhn (1962) who used it to signify the fundamental thought patterns that dominate a field of study at a given point in time. These thought patterns define the

- Fundamental questions that are asked about the phenomena of interest
- What answers and results are relevant, and
- How experiments are to be conducted

In short, a paradigm defines a field of study in the widest sense. Paradigms are not true or false, and they are never tested as such. They only provide the frame within which scientific work is conducted. They are nevertheless important because they set conditions for progress and for judging progress.

A paradigm is not a theory. Many theories may compete within a paradigm, but the paradigm nevertheless sets the limits within which the theories are developed. It is therefore important that it can be articulated, so that it can be communicated, discussed and criticized. It is the purpose of this paper to provide an articulation of a paradigm for the study of self synchronization so that we can make progress in studying it.

A paradigm for the study of self synchronization

Our first step towards a paradigm for the study of self synchronization is to find a definition in terms of the situation which requires self synchronization. Here, we take that definition from our earlier work on distributed decision making, a problem for which self synchronization is one possible answer.

As described by Brehmer (1991) a problem that requires distributed decision making for its solution has the following characteristics:

- It is too large for a single unit to handle so it requires coordinated efforts from a number of units
- The situation is dynamic requiring both planning and execution
- Each unit owns part of the resources that are needed to handle the problem, but no unit has complete control over all resources
- Each unit has a limited view on the problem, and no unit can achieve an overall view of the problem without input from the some of the other units
- No unit commander has the authority to coordinate the other units

This description covers a large class of problems, ranging from the fire fighting and the military C2 problems that provide the motivation for this paper, to the kinds of problems faced in today's civil-military operations where military and civilian units must cooperate to solve the complex problems that they face. It also applies to many problems in modern industrial production (Rasmussen, Brehmer & Leplat, 1991). It thus fulfils the requirements for a paradigm in that allows us to discuss a wide range of the problems making it possible to accumulate data from observations in many different situations and tasks.

The next step is to operationalize the conceptual paradigm in an experimental paradigm where the factors in the conceptual paradigm can be given empirical meaning. Our first attempt to do this was D3FIRE (Svenmarck & Brehmer, 1991)[2].

D3FIRE: An experimental paradigm for the study of distributed decision making and self synchronization

D3FIRE is an operationalization of the paradigm described above. It incorporates all five features listed there in a "manned computer simulation", as so called microworld (Brehmer & Dörner, 1993). The problem facing the participants is that of fighting forest fires. D3FIRE allows up to four participants in an experiment. The general concept is illustrated in Figure 1.

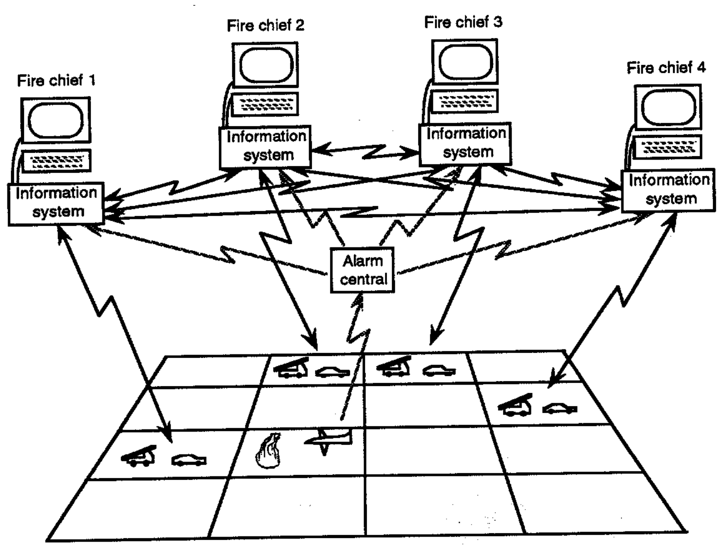

Figure 1. The D3FIRE concept.

[2] This paradigm has subsequently been further developed into a more general experimental tool called C3FIRE by Granlund (2002). For our present purposes, D3FIRE is sufficient so we stay with that paradigm here. Anyone interested in obtaining C3FIRE for experiments should visit Granlund's home page www.c3fire.org.

Each participant sits in front of a personal computer, and the four computers are connected via a Local Area Network (LAN). Figure 2 shows what the participant sees on his or her screen. There we find a map covering the area. Each participant has his or her own window on the map and can see only what goes on in that window. The size of the window can be changed as the experiment requires, allowing the participant anything from a very small view of one square to a view of the whole map. In all of the experiments reported here, the size of this window was 9 squares, as illustrated in the figure. To the left, there is a weather report providing information about the strength and direction of the wind. Immediately below that there is a window that shows the latest reported position of each fire fighting unit (= the position of each participant, as in these experiments each participant has only one asset). The bottom window is for writing and receiving e-mail messages if the design of the experiment requires communication to be done by e-mail. Other forms of communication are possible as well, but e-mail is convenient as a log of all e-mail traffic is constructed automatically by the program. Experiments with D3FIRE have been run with both voice communication and e.mail communication. The general results do not differ between these two conditions. The number of messages sent is

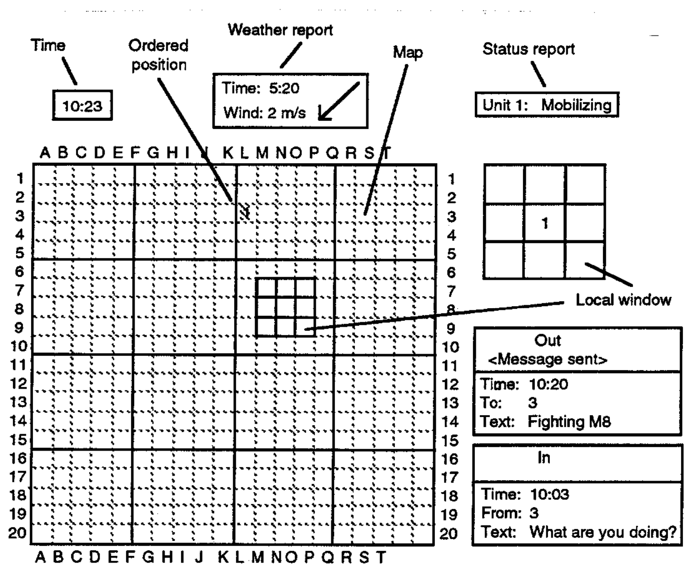

Figure 2. The D3FIRE interface.

about three times higher with voice communication than with e-email, but the relative frequency of different kinds of messages sent remains the same (Svenmarck & Brehmer, 1992).

One of the difficult problems when doing experiments with microworlds is to find useful and reliable dependent variables. D3FIRE provides two kinds of measures: *measures of performance* in the form of the number of cells lost to fire and time taken to extinguish it, and *measures of communication*: number of messages sent and received by each participant and the kind of information sent and received. In addition, the position of all FFUs and the cells that are burning and extinguished is logged continuously. As for the messages, a classification system comprising four classes of messages is used: questions, information, orders and miscellaneous (mainly chit-chat). The questions and information messages are further broken down into four subcategories according to the subject: reports about fire, about activity, own location and intention (where the unit will go). This classification system allows for reliable classification of all the messages, without an unduly large proportion ending up in the miscellaneous category. It is also possible to derive process measures based on the moves and the measures. For example, Brehmer (1997) in an experiment with a hierarchical architecture, demonstrated that handling the dynamics of the fire is a major problem for the participants. This was shown by comparing the orders given by the centrally located participant (“the commander”) and the position of the fire when the unit to which the order had been given reached its destination. The results showed that the orders did not take account of the actual development of the fire, so that the unit obeying an order often found itself in the middle of the fire, rather than at front of it, making effective fire fighting difficult. The D3FIRE program allows the experimenter to define different communication architectures by introducing constraints on who can communicate with whom. It is also possible to vary the decision rights of the various participants.

In the present experiments, the time between updates of the information on the screen was set to 30 sec., and interval that creates moderate time pressure (cf., Brehner, 1997).

D3FIRE was designed to study coordination and synchronization processes in as pure a form as possible. Therefore, the actual fire fighting process is made as simple as possible: a unit will go to its destination and if there is fire there, or when fire reaches this destination, it will start to fight it. A unit can be ordered to leave a cell before the fire in that cell has been extinguished if the participant so wishes. The details of the extinguishing process are handled automatically by the unit, and it makes sure that it has the resources required to handle the fire in a cell[3]. The only aspect that the participant has to learn and remember is that it makes a certain time to extinguish fire in a cell.

A first step towards a conceptualization of self synchronization

For experimentation, an experimental paradigm is not sufficient. We also need some preliminary conceptualization of the nature of the phenomenon to be studied, a preliminary theory to give some guidance in setting up experiments. The preliminary conceptualization guiding the present work is the following.

A force starts with an understanding of its mission and it has the ability to translate this into tasks. In D3FIRE, this means that each participant comes to the experiment with a basic understanding of fires, how the wind affects how they spread and that it is necessary to coordinate ones activity with that of other participants (FFU commanders) to stop the fire from spreading and to eventually extinguish it. This means that the participants are assumed to understand the need to cover the whole fire front so that the fire will not escape and envelop them. Their coordination can be based, either on some shared understanding of the situation (of the location of the fire, how it spreads and what the other FFUs are doing), or, if such an overall view cannot be achieved, upon the realization that any hole in the front needs to be closed.

[3] In C3FIRE complications relating to resources can be introduced.

This gives the basic requirements for self synchronization: an overall view of the situation that can serve as a basis for deciding where one can contribute, or, in the absence of this, a limited view where gaps that require action can be detected.

A further consideration has to do with time. Synchronization cannot be done instantaneously, it requires time. This means that the moment when units have to plan to be in place for the concerted activity is sometime in the future, rather than now. Synchronization will therefore have to be based on some conception of what will happen, and what will be needed, sometime in the (near) future. Information about the future with respect to the fire requires a (mental) model of the fire that can be used to predict its future development. Information about the future positions of the units with which synchronization must be made is provided by the *intentions* that the commanders of these units have. We therefore expect that communicating these intentions would be an important aspect of the synchronization process.

These considerations suggest that the variables of interest for experimentation have to do with the extent to which the unit commanders can obtain the information they need so that they can detect where their efforts are needed. This must be information that makes it possible to obtain a picture of future synchronization needs. The experiments to be presented below are concerned with various aspects of this problem, as explained in the context of each experiment.

Empircal demonstration

In this section, we report four experiments using D3FIRE that provide an empirical demonstrations that direction and coordination can be achieved without a commander and some of the variables that affect the self synchronization process. Earlier studies with D3FIRE (Brehmer, 1997; Brehmer & Svenmarck, 1995) have been concerned with organizational factors, and they have shown that networked teams without a commander are more effective than hierarchically organized teams under time pressure but not when the teams have ample time for their decisions. Then, a traditional hierarchical organization is more effective. The present studies, in contrast, are concerned with the effects of different forms of information and communication, factors that are directly relevant to self synchronization.

Experiments: Some preliminary consideration

One of the problems in doing experiments here is to find an adequate control condition, or base line. Since each experimental group starts with a message from an alarm central that there is a fire in a given location, they will not behave randomly, but will go towards the fire and they will thus extinguish some of the fire, even if they do not synchronize their behaviour at all. In these experiments we use a control condition with four FFUs. One participant is in charge of each. Each participant has a limited window (9 cells) on the forest and there are no opportunities for communication. This control condition probably leads to some underestimation of the effects of self synchronization, but that is, of course, better than having a control condition that leads to an overestimation of the effects. The same control condition is used in all experiments that employ a control. This means that in our experiments, we examine the effects that our independent variables have in addition to what synchronization goes on in the control condition.

In all experiments, the participants were university undergraduates aged 20-25, who were paid to take part in the experiments. They worked together in teams of four participants each, and in the experiments there were either six or four such teams in each condition as described for each experiment. In each condition, they first went through a 20 minute familiarization period. They then extinguished three experimental fires, with a time limit of 20 minutes per fire.

In each condition the measure used to evaluate performance is the mean number of cells lost to fire (over the three experimental fires).

As there are few units in each condition (four or six[4]) traditional statistical analysis in the form of significance tests is not very meaningful. Instead the results are evaluated by means of Cohen's d, a measure of effect size (Cohen, 1992). Cohen's d expresses the size of the effect as a proportion of the standard deviation. Thus a Cohen's d of 0.50 means that effect of the experimental variable was half a standard deviation, and a Cohen's d of 1.0 means that effect was one standard deviation (quite unusual in behavioural science experiments). In view of the fact that these experiments are intended as demonstrations, rather than definitive studies, this form of analysis should be sufficient.

Experiment 1: The effect of being networked

One of the most fundamental questions in self synchronization in network-centric operations is the effect of actually having a network that makes communication possible. This is the topic of the first experiment. It compares a condition where the four participants in each quartet have access to a network and are able to communicate freely among themselves via e-mail, with the control condition where they do not have a network and cannot communicate. No participant was given any authority over the others, so there was no commander. There were six quartets of participants in each condition.

The results were as expected: The networked group outperformed the no network group. In the former 68.3 cells were lost to fire while in the latter 82.7 cells were lost. This represents a strong effect (Cohen's d = 0.84).

Figure 3, which shows the frequency of messages of different kinds sent per minute, gives

[4] Each unit is made up of four participants, but the unit of observation is not the individual participant but the team of four. Thus, it takes 16 participants to obtain four units of observation.

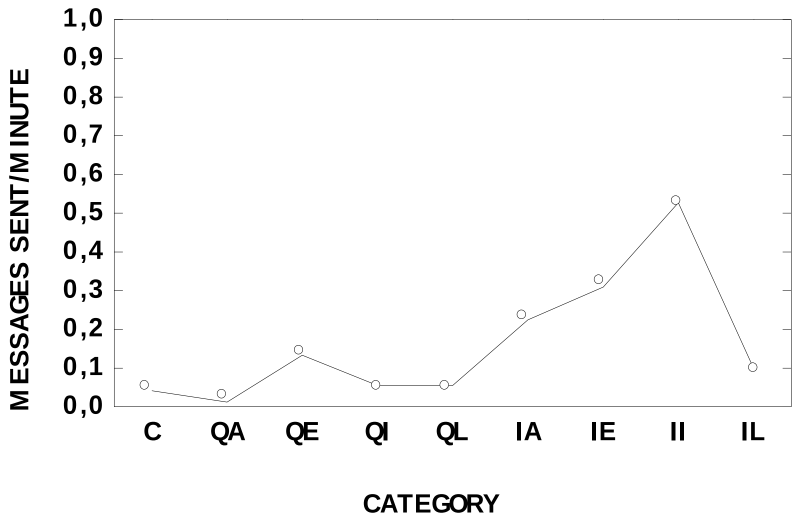

Figure 3. Frequency of different kinds of messages sent per minute by each participant in the networked condition. C denotes commands (when one participant told another participant to do something), Q denotes questions, and I denotes messages giving information. A denotes Activity, E fire, I intention and L position.

some information about how the participants achieved this. First, note that there are very few commands. Thus, no single participant seems to assume the role of commander, directing the other participants in the quartet. Instead, messages giving information about intentions (messages of the type: I will go to X) and information location of fire have the highest frequencies, suggesting that the participants understand the need to have information about the future position of the other FFUs for their synchronization. There are relatively few messages giving information about own position (and few questions about the position of someone else as well). We will return to the problem of what is needed for synchronization in Experiments 3 and 4.

In our next experiment, we compare the no commander condition with a condition where one of the participants has been assigned the role of a commander.

Experiment 2. Effect of making one participant the commander

This experiment compared the experimental, networked condition in Experiment 1 with a networked condition where one of the participants (who was chosen randomly) was appointed commander of his/her quartet of FFUs. The participants were told that they should obey the commands from the commander. In both conditions each participant had access to the network and could communicate with every other participant in the quartet via e.mail. There were six quartets of participants in each condition.

The results showed that the performance was *worse* when there was a commander than when there was no commander (74.6 vs. 68.3 cells lost to fire). The effect is moderately strong (Cohen’s d = 0.52). This result may seem surprising, but it was not unexpected. As noted above, Brehmer (1997) found in an earlier study with D3FIRE that commanders under these circumstances had considerable problems understanding the dynamics of the fire. This was shown by the fact that their commands were not very effective. Specifically, the commanders would underestimate the rate at which the fire would spread, and their commands would place the FFUs under their command behind the fire front, rather than at the front, making their efforts to extinguish fire ineffective. Apparently, the problem of finding the best position for fighting fire is much better handled locally by the individual FFU commanders on the spot than by a centrally placed commander.

Experiment 3. Effect of providing blue force tracking

In Experiment 3, each participant was given continuous information about the position of each of the other FFUs, but they were not allowed to communicate with each other. This corresponds to information provided by “blue force tracking”. They were given information about when and where a the fire started, but they were given no further information about the fire, except from what they could see in their “window” and what they could infer from the positions of the other FFUs. There were four quartets of participants in the blue force tracking condition. The performance in the blue force tracking condition was compared to the standard control condition with no network. In the former condition, 74.8 cells were lost on the average, while in the control condition 82.7 cells were lost. Cohen’s d = 0.45 indicating a weak to moderate effect. Compared to a network with opportunities for communication, blue force tracking only apparently leads to lowered effectiveness. In the blue force tracking condition 74.8 cells were lost, while in the networked condition 68.3 cells were lost, on the average. The effect is effect is moderately strong (Cohen’s d = 0.60).

This result is not surprising. As we noted above, synchronization generally cannot be done on the basis of information about current position. It takes time, and must be done with an eye to the future positions of those with whom one wants to synchronize. The best information about the future position of a unit is given by the intention of the unit commander, something that our participants seem to understand (see Figure 3 above). In Experiment 4, we test this hypothesis by examining the effect of augmented blue force tracking which provides information also about this factor.

Experiment 4. Augmented blue force tracking

In Experiment 4 we compared the performance of a condition with what may be called augmented blue force tracking with that in the standard, no network control condition. In the augmented blue force tracking condition, the participants were given information not only about the position of each of the other FFUs but also about the latest command that that had been given to each of them. Specifically, the position was given by displaying the number of the FFU in the cell where it currently was, and the intention was displayed by means of an arrow pointing from the cell where the unit was to the cell where it had been ordered to go. There were four quartets of participants in the experimental condition and six quartets in the control condition. Performance was better in the augmented blue force tracking condition than in the control condition, 63.6 in the experimental condition vs. 82.7 in the control condition. This is a strong effect yielding a Cohen’s d of 1.06. This represents an improvement from just providing the positions from 74.8 cells lost on the average with ordinary blue force tracking to 63.6 cells lost with augmented blue force tracking. This is again a strong effect (Cohen’s d = 1.16). However, compared to the ordinary network condition with full communication, it represents only a small improvement from 68.3 cells lost in the networked conditions to 63.6 cells lost with augmented blueforce tracking. This is a relatively small effect (Cohen’s d = 0.27). This is not so surprising as we have seen in Figure 3 that the participants in the networked condition communicated both their positions and their intentions, so the augmented blue force tracking probably did not add very much information compared to what the participants could obtain from their ordinary communication.

Summary of the results of Experiments 1 – 4

The results of these experiments offer few surprises to those who have followed the writings on network-centric warfare. Thus, they show that providing a network improves effectiveness. Moreover, the results show that blue force tracking has positive effects if there are no other means of communication, but it seems less effective than a network with opportunities for full communication. As expected, augmented blue forces tracking allowing for communication of intentions is more effective than simple blue force tracking, but it is only a small improvement over the effectiveness of a network with opportunities for unrestricted communication, as would be expected from the fact that the participants communicate their positions and intentions in the network condition also. Interestingly, but not surprisingly, making one of the participants in each quartet a commander with the authority to coordinate efforts of the team does not improve performance. The explanation seems to be that the commanders are not able to handle the task of finding the best position for the units to fight the fire as well as they themselves are able to do on the basis of their local information.

Conclusions

The results of the experiments agree with our expectations and with what would be expected from earlier writings on the possible effects of the network (e.g., Alberts, et al., 1999) in that they show that access to a network has strong and positive effects on effectiveness. The fact that adding a commander reverses the effect suggests that the positive effect in the networked condition without a commander is indeed a result of self synchronization, an interpretation that is further strengthened by results with respect to what the participants communicate. Synchronization seems to take place on the basis of information about positions and, especially, intentions, just as it does in the augmented blue force tracking condition in Experiment 4.

The most important conclusion is, perhaps, not these results in themselves, but the fact that D3FIRE produces effects of the variables that should be effective in self synchronization experiments: It thus provides us with a tool for investigating the self synchronization phenomenon and the variables that affect it in detail, thus removing some of the mystery surrounding the phenomenon.

D3FIRE is, of course, just one possible operationalization of the general paradigm for studying self synchronization described above. Other operationalizations could involve real military forces in actual exercises, so long as one makes sure that the exercises adhere to the conditions specified in the paradigm. Moreover, fighting fire is just one possible operationalization. Other kinds of enemies could easily be substituted for the fire, as for example in the DKE simulation (Kuylenstierna, Rydmark & Sandström, 2004). DKE is a two sided, computerized war game that has been used for experimentation in our laboratory for some time, although not yet for studies of self synchronization.

Comparison with ELICIT

The ELICIT experimental Laboratory for Investigating Collaboration, Information-Sharing, and Trust) paradigm (Ruddy, 2007) is currently the most popular paradigm for studying self synchronization phenomena. The present paradigm is not in competition with that paradigm. The experimental arrangements used in ELICIT experiments fit comfortably in the general paradigm suggested here, and may be seen as an alternative operationalization of that paradigm. The ELICIT experiments have a different focus than D3FIRE experiments, however. Being based on the data-information-knowledge framework, ELICIT experiments investigate the important problem of how a group of people, each with a different piece of the total picture, are able to arrive a shared situational understanding. But what they then do with that shared understanding and whether it actually results in self synchronization is not assessed. D3FIRE experiments, in contrast, bypass this stage (although it is possible to introduce measurements of the extent to which the participants achieve a common understanding of the situation also in these experiments) and focuses directly on the effects of the synchronization process. Thus, the two operationalizations complement each other, and together they provide greater possibilities for studying self synchronization phenomena than any of them on its own.

References

Alberts, D.S. (2007). Agility, focus and convergence: The future of command and control. *International C2 Journal, 1*, 1-30..

Alberts, D.S., Gartska, J.J. & Stein, F.P. (1999). *Network centric warfare*. Washington, D.C.: CCRP.

Brehmer, B. Time scales, distributed decision making, and modern information technology. I J. Rasmussen, B. Brehmer & J. Leplat (Eds.), *Distributed decision making: Cognitive models for cooperative work*. Chichester: Wiley, 1991.

Brehmer, B. (2007). Understanding the functions of C2 is the key to progress. *International C2 Journal, 1,* 211-232.

Brehmer, B. & Dörner, D. Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior*, 1993, *9*, 171-184.

Brehmer, B. (1997). Effects of time pressure in fully connected and hierarchical architectures of distributed decision making. I Y. Waern (Ed.), *Co-operative process management*. London: Taylor & Francis.

Brehmer, B. & Svenmarck, P. (1995). Distributed decision making in dynamic environments: Time scales and architectures of decision making. I J.-P. Caverni, M. Bar-Hillel, F. H. Barron & H. Jungermann (Red.), *Contributions to decision making. I.* Amsterdam: Elsevier. Pp 155-174.

Cohen, J. (1992). A power primer. *Psychological Bulletin*, *112*, 155-159.
Granlund, R. (2002). *Monitoring distributed teamwork training*. Dissertation. Linköping: Linköping University

Kuhn, T. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Kuylenstierna, J., Rydmark, J. & Sandström, H. (2004). Some research results obtained with DKE: A dynamic war game for experiments. *Proceedings of the 9th ICCRTS, Copenhgen, Denmark, 14-16, 2004.*

Rasmussen, J., Brehmer, B. & Leplat, J. (Eds.), *Distributed decision making: Cognitive models for cooperative work*. Chichester: Wiley, 1991.

Ruddy, M. (2007). ELICIT: The experimental laboratory for investigating collaboration, information sharing and trust. *Proceedings of the 12th ICCRTS, Newport June 19-21, 2007*.

Svenmarck, P. & Brehmer, B. (1991). D3FIRE: An experimental paradigm for the study of distributed decision making. In B. Brehmer (Ed.), *Distributed decision making. Proceedings of the third MOHAWC Workshop, Beligrate Italy May 15-17, 1991*. Roskilde: Risö National Laboratory.

Svenmarck, P. & Brehmer, B. Face-to-face communication vs. communication by means of electronic mail in distributed decision making. In B. Brehmer (Ed.), *Separate papers of the MOHAWC Project*. Roskilde: Risö National Laboratory