13th ICCRTS: C2 for Complex Endeavors

"Reducing Network Load through Intelligent Content Filtering**"**

Topic 7: Network-Centric Experimentation and Analysis

Topic 9: Collaborative Technologies for Network-Centric Warfare

Topic 2: Networks and Networking

Trude Hafsøe and Frank T. Johnsen

Point of Contact

Trude Hafsøe

Norwegian Defence Research Establishment (FFI)

P.O. Box 25

NO-2027 Kjeller

Norway

+47 63 80 74 31

trude.hafsoe@ffi.no

Reducing Network Load through Intelligent Content Filtering

Abstract

Future international military operations will be more complex than traditional operations undertaken by just one nation; military units from different nations will have to cooperate with not only with each other but also with local governments and civil organizations in order to reach common goals and to ensure a shared understanding of each other's task and domain responsibility. One characteristic of such endeavors is that each organization brings with it its own information and communication systems. Interconnecting these communication systems will lead to an increase in the total amount of information available to users of these systems. One of the main challenges when building an information infrastructure to support such operations is to ensure information superiority; all users must get access to the information they need to perform efficiently, while at the same time avoiding that the user is flooded with irrelevant information. Making sure that only relevant information is transmitted is even more important in tactical systems, where communication resources are very limited. This paper describes the use of several different types of content filtering as a measure for reducing the network load, and presents the results of our experiments with content filtering in disadvantaged grids performed at NATO CWID 2007.

Keywords: content filtering, tactical networks

Introduction

International military operations, such as those performed by the NATO Response Forces, require that a number of participants from different nations and organizations work together towards a shared purpose. Mission effectiveness depends on the participants' ability to communicate both effectively and efficiently with all cooperating partners, thus a common information infrastructure is essential.

A common information infrastructure for NEC operations needs to be able to ensure that all users are supplied with information that is both sufficient and relevant enough for them to be able to make appropriate decisions at all times. Such an increase in the amount of information that is available to users can cause problems both at the cognitive level and at the network level. This paper discusses how content filtering can be used to reduce the impact increased volumes of information can have on the network. In particular, tactical links have low bandwidth available, and only relevant information should be sent over the network to limit the possibility of congestion due to irrelevant data. Maintaining information superiority, which is the capability to collect, process, and disseminate an uninterrupted flow of information while exploiting or denying an adversary's ability to do the same [7], means that proper information management is critical. Ensuring that only relevant information is transmitted over the network helps maintain information superiority, in that irrelevant information is not allowed to disrupt the information flow by overflowing the network

Content filtering can be used to alleviate network congestion by removing information that is not relevant to the user. Several different types of content filtering exist, depending on the type of the data and how the data is used. This paper presents results of experiments performed at NATO CWID 2007. In these experiments we tested how content filtering can be used to ensure that only relevant information was supplied to tactical users, and thereby avoiding overloading the network and saving bandwidth resources at the same time.

The remainder of this paper is organized as follows: First we present the tactical system and the data format used in the experiments, followed by a description of different types of content filtering. Then we present the experiment setup and results.

Background and Motivation

Using content filtering is one of several measures that can be employed in order to increase the information infrastructures ability to adapt to changing network and battlefield conditions. [1] states that this kind of adaptability is one of the requirements next-generation military information systems need to fulfill.

In order to test the usability of content filtering, we performed several technical trials at NATO CWID 2007, as part of a larger Service Oriented Architecture (SOA) experiment. We SOA-enabled an experimental tactical system called NORMANS Advanced, as this allowed us to test not only content filtering alone, but also confirm that the content filtering techniques we tested works in conjunction with Web services.

NORMANS Advanced

The NORMANS [2] concept includes a C4I system that is designed to take the roles of the different type of users into account. Dismounted soldiers will in the future act as sensors, effectors and decision makers, and their C4I equipment, both hardware and software, must reflect their main tasks. The NORMANS C4I concept has a modular approach based on a main navigation and communication module, named NORMANS-light, for all private soldiers in a section. A more advanced commander system (NORMANS advanced) uses digital maps, friendly force tracking and the ability to mark red force observations to help improve situational awareness for more advanced users. The NORMANS C4I concept is based on voice and data communication within the sections using a simplified data transmission protocol. In addition, voice as well as data can be sent between sections using IP and a tactical messaging system.

At NATO CWID we used a slightly modified version of the NORMANS advanced software, as the use of a proprietary protocol complicates interoperability. Using standard based solutions makes the task of interconnecting systems easier, so we modified the software to communicate by inputting and outputting XML formatted data.

Filtering

To maintain information superiority in a coalition force it is paramount that all necessary and relevant information is disseminated throughout the network. Thus, it becomes important to identify the information that is indeed relevant, and transmit only that over the network. Which information that is relevant will vary from user to user depending on their role and what the information will be used for. In order to perform correct filtering, the system performing the filtering must be aware of the needs of its users, and filter accordingly. In our trials the filtering was done based on a profile that specified which information was most important to the user.

How can we identify information as relevant? There are many factors to take into consideration. For example, some information may only be relevant within a certain area of operations, and thus it should be disseminated only in that geographical location or to users outside the area that specifically request that information. Some information may change frequently, for example position information, whereas other information can change less frequently or even be entirely static. In such cases messages containing status updates have different requirements as to how often they need to be transmitted. If network resources are scarce, then knowledge of the importance of the information can help the system prioritize by sending the most important information first, and delaying or perhaps even entirely discarding the less important information. Which information each unit needs, is first and foremost a question of which role the unit has. In some cases filtering an entire message or part of a message will save network resources while still ensuring that the recipient gets all the relevant information that it requires. Security issues, trust and clearance are also important aspects, and filtering should also be used to stop classified information from exiting a system. This latter aspect is currently subject to research and has a lot of open issues still which are beyond the scope of this paper. See [4] for an overview of some of the security related filtering experiments at NATO CWID 2007.

In summary, we have several aspects to consider when disseminating information. The most important aspect is that only necessary and relevant information should be received by the units. There are many factors that can be used for filtering; a non-exhaustive list is presented below:

- Geographical filtering
- Frequency based filtering
- Priority based filtering
- Role based filtering
- Security label filtering

Furthermore, the filtering can be of two types, in that one can filter

- Entire messages, or
- Parts of messages.

In our experiments at NATO CWID 2007 we used a combination of geographical and frequency based filtering. The information we considered in the experiments was only tracking information, and thus we used filtering of parts of messages to ensure that only relevant tracks were delivered to the unit in the (simulated) field. In the following section we discuss the experiments and filter functionality in detail.

Implementation challenges

Having discussed some of the different issues of filtering above, we now turn our attention to the challenges that arise when one considers implementing a filtering scheme. I.e., we need to decide *how* and *where* the filtering should be done. The "how" of filtering is basically a matter of choosing which technique(s) to implement, for example a combination of geographical and frequency based filtering as we used for tracking information in our experiments. Which kind of filtering is best to use will depend on the kind of information the message contains. The challenge here is to identify the recipient's needs when designing the system, and performing filtering accordingly. How to describe these needs should be a matter of discussion and standardization within NATO, and a further discussion of these challenges are beyond the scope of this paper. After having decided which policy to employ, there is the issue of where the functionality should be implemented.

The "where" of filtering is a matter of placing the filtering functionality in the NEC information infrastructure. The easiest way to filter information is in the receiving unit. That unit may know which data is relevant to present to its user, and can discard other information. This requires no state information in the network or in the information producers, thus leading to low system complexity. However, for each piece of unimportant information that is discarded in the end-system, a corresponding amount of bandwidth has been wasted in transmitting this information all the way from the producer to the receiver. Ideally, only information that is relevant according to the chosen policy should be injected into the network. If one can perform the filtering where the information is produced, then this is optimal in two ways; firstly, no bandwidth is wasted, and secondly, only relevant information is received by the end-system terminals. However, implementing the filtering policy in every potential information producer may be infeasible. Filtering requires some processing for the system to find out whether the information should be transmitted or not by inspecting the information and comparing it to the policy. As such, this will put higher requirements on the computational capabilities of these systems. In any case, proxies should be employed between networks to ensure better use of resources [5]. The proxies can for example function as security guards [4], something that will be needed on the way towards full-fledged NEC to secure the information flow. If we need proxies anyway, then perhaps one should just implement the filtering functionality there and reduce the complexity of the end-systems? System implementation complexity is reduced by centralizing the filtering functionality in proxies, but such a solution leads to an increase in bandwidth

consumption between the producer and the proxy since the data transmitted between these are unfiltered. It should be noted that the proxies will need policy information for each kind of unit that is to receive information. Furthermore, the proxy must be able to recognize and process different kinds of message formats that can pass through it. This means that the proxies will become potential bottlenecks in the network due to the computational complexity of message processing.

In short, we have discussed the three places to perform filtering:

- Filtering in the end-system terminals
    - Low complexity
    - Stateless
    - High bandwidth use
- Filtering in proxies
    - High complexity – must know all combinations of end-system terminal and message formats and the corresponding filter policy
    - Proxy may need to keep state (for example position information for each receiving unit in the case of geographical filtering)
    - Reduces bandwidth use across networks
- Filtering in the message producing system
    - Medium complexity – must know all end system-terminals and corresponding filter policy
    - May need to keep state
    - Best bandwidth utilization since no unnecessary information is injected into the network

Basically, filtering in the end system should be avoided since it wastes network resources, and especially on the tactical level bandwidth is scarce. Seemingly, filtering in the message producing system is the best option. However, proxies also have an important benefit over that of filtering in the producer: If the producer sends information to recipients with different capabilities, then it must filter once for each type of recipient. A proxy, on the other hand, will be closer to the recipient, and as such potentially have fewer types of recipient to filter for.

In our experiments we used the proxy filtering approach. The local HQ track store was on a high capacity LAN together with the proxy server which was connected to the (emulated) disadvantaged grid. We focused exclusively on disseminating tracking

information, and as such we had a relatively simple proxy solution: Our proxy kept state about each recipient (last reported position). It received all the information from the track store at regular intervals, and would then perform geographical filtering of the tracks based on the state it kept before sending the regional tracks over the tactical network to the NORMANS unit. The proxy also performed frequency based filtering as one of the filter types used did not transmit all data at the same interval, but rather updated the most relevant data more frequent than other data.

Experiments and Evaluation

Defence R&D Canada have performed a series of technical trials related to the dissemination of operationally important information in congested tactical radio subnets using their Low Bandwidth Test Bed [3]. Among the experiments performed is a test of dynamic reduction of network load by using content filtering techniques, as described in [1]. The experiment involved using an information management rule to determine whether or not to suppress replication messages. These messages contained a unit's report of its own position, and the rule used was based on how far the unit had moved since it previously reported its own position. Each node would use this rule to make an autonomous decision to either broadcast or suppress its own position at given time intervals. This means that the type of content filtering done in this experiment was a type of frequency based filtering, but the information management rules used were geographically based. The decision whether to perform filtering or not was made locally, which means that the required state could be maintained by each unit independently.

In our experiments units reported their own position, and the position data was gathered by a central unit, and distributed to all units. Because the unit reporting its own position was in constant movement, we did not perform filtering of the unit's own position reports. We concentrated on filtering data that was being sent out to the units, as these messages could contain position data for all other units in the battlefield, and were thus significantly larger in size than the position reports transmitted by each unit individually. We performed this filtering as close to the source as possible in order to save bandwidth on the simulated tactical links in addition to avoiding flooding units with information they did not want according to their profile. Allowing an intermediate node to perform filtering made our experiments more complex, as the intermediate node had to have updated information about the location of each unit in order to correctly perform geographical filtering. The intermediary did this by intercepting the reports sent by each unit containing their location, and maintaining an overview of the last known location of each unit.

In our experiments at NATO CWID 2007 we looked into the use of content filtering for a blue force tracking application, using NFFI-formatted data. We performed two different kinds of filtering, namely geographically based filtering of complete tracks, and filtering of optional information within tracks. A simple form of geographical track filtering is using a fixed zone filter to remove all tracks that are outside the unit's area of operation. Geographical filtering can also be used in combination with a second content filter that reduces the frequency of track reports.

Geographical filtering - Fixed zone filter

The fixed zone filter consists of a simple filter mechanism on the server side which is performed on each track. ”Relevance” in this filter is defined as tracks within a certain distance of the soldier. For each track the distance to the last known position of the unit is calculated. If the distance is greater than a certain number (fixed, but configurable in the filter) then the track information is not sent. All tracks closer to the unit than this are reported. Such filtering is important to ensure that no unnecessary information is sent. Figure 1 shows unit placement on the battlefield for a tactical user (left) and a local HQ user (right) respectively. The local HQ has a complete overview of the situation, while a fixed zone filter is used to limit the information sent to the tactical user. This means that the tactical user only gets notified of other units that are within its area of operation.

Tactical unit display

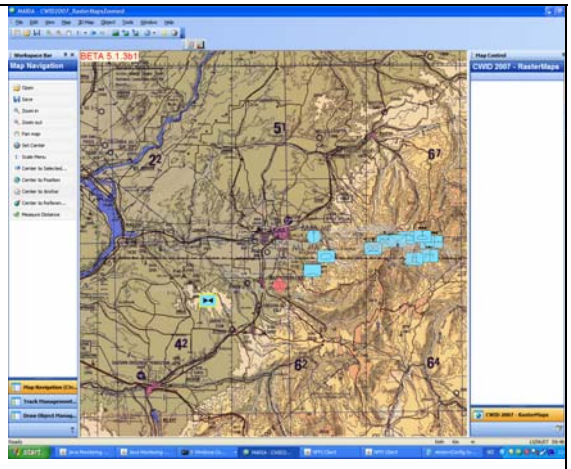

Local HQ screenshot

Figure 1 Fixed zone filter

Geographical filtering - Zone ring filter

The zone ring filter is similar to the fixed zone filter in that it uses distance as its filter metric. The idea is to save bandwidth while at the same time providing the unit with an overview of a larger section of the battlefield. This technique can be used when bandwidth limitations makes it impossible to transmit information about all relevant units as frequently as needed. The zone ring filter is optimized to allow for more frequent updates of tracks that are closer to the client than those that are further away. While the fixed zone filter uses one ring as its zone (a track is either inside (report it) or outside (don’t report it) the ring), the zone ring filter uses three rings. These rings are arranged in

such a way that the inner ring is updated most frequently, followed by the second ring, and finally the third ring. Information about tracks outside the third ring will not be sent. This is comparable to the fixed zone filter, which actually only has the functionality of ring 3.

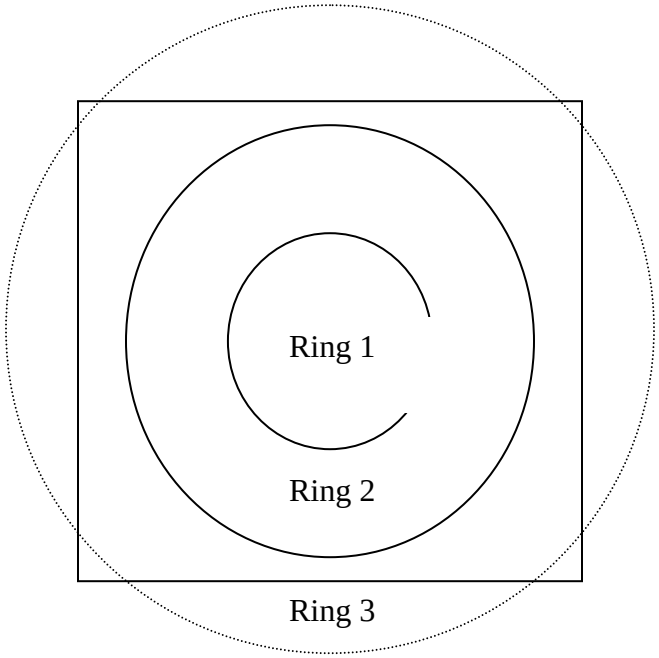

Figure 2 Zone ring filter layout

In Figure 2 we give an example of the zone ring filter: The rectangle is the display of the NORMANS system. The display is always centered at the soldier's position. This position is reported back to the central service every 40 seconds. Based on this position, the service will send its filtered track information back at certain intervals (configurable), with different multiples of the interval for the different rings (also configurable). The way we used the filter at NATO CWID was to configure rings 1 and 2 so they fit inside the map view of the unit, and configure ring 3 so that it was just outside what was possible to visualize. That way, the unit would receive frequent updates for units positioned inside ring 1 and less frequently for units that are inside ring 2, but outside ring 1. Updates about units that fall between rings 2 or 3 were rarely sent, and no information outside ring 3 was ever sent. Note that the rings are adjacent and do not overlap: When track information inside ring 2 is updated the information inside ring 1 is not sent (unless ring 1 and 2 have the same update frequency configured). The same applies to information from within ring 3, which is also independent of ring 1 and 2.

Figure 3 shows an example of the map display for a tactical user on the left and the local HQ on the right. At first glance the two images seem to report the exact same

information, but the local HQ has more frequently updated information for the units that are far away from the tactical user.

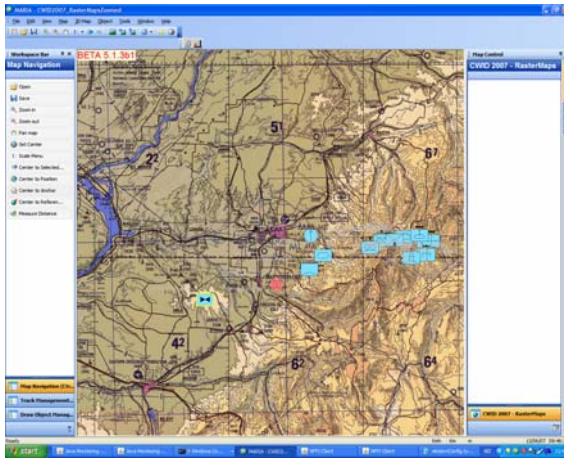

Tactical unit display *Local HQ screenshot*

Figure 3 Zone ring filter example

Filtering optional fields

We have now discussed two filters for tracks. We can also filter information within the tracks themselves, optional fields which may be of lesser importance to the soldier. This form of content filter is discussed below.

At CWID we transmitted all track information as NFFI, which has some mandatory and a lot of optional fields. The NORMANS Advanced tactical system can only utilize a very small subset of the information that can be represented in an NFFI message. This means that if we transmit the full message, a large part of the information will be discarded by the recipient as not relevant. In order to save bandwidth we employed optional field filtering by removing all the irrelevant information at the server side before transmitting the data.

Results using filtering of optional fields combined with the zone ring filter

In our experiments we used the NIST Net[1] network emulator package for emulating the tactical link. Using this software, the link was limited to a bandwidth of 2.4 Kbit/s, which is representative of the bandwidth one can expect when using a radio network designed for speech traffic. The time it takes to transmit a package over this link will vary depending on other traffic that is using the same link, and the numbers given below are typical of the ones we observed during the experiment.

When transmitting track updates over a tactical network like the one used in the experiments, even a small number of tracks per message will quickly fill the link. The link usage can be reduced by either sending messages less frequent, or by reducing the size of the messages.

Our experiment with the fixed zone filter was aimed at determining what the maximum update frequency is when using a speech channel. Applying the fixed zone filter reduced the number of tracks in the messages significantly enough to allow for an update frequency of 30 seconds. This update frequency used up most of the available bandwidth, but a somewhat lower update frequency that leaves more link capacity free will in many cases be sufficient.

Further reduction of the bandwidth usage can only be achieved by applying a stricter filtering method or by reducing the update frequency. However, reducing the update frequency means that the risk of tracking information becoming outdated increases. We investigated how to better utilize the limited bandwidth available by applying the zone ring filter described above without having to compromise too heavily on accuracy.

Table 1 shows some examples of the measured transmission time of NFFI track updates of varying size. When using the zone ring filter, the transmission of the NFFI tracks was split up so that the most frequently updated tracks, which in our case were 5 tracks, took about 7 seconds to transmit. This means that these tracks could be updated every 15 seconds, while at the same time leaving enough free bandwidth to allow tracks in the other two zones to be updated at least once per minute.

Tracks in NFFI message	NFFI message size	Wire message size, i.e. with compression	Time traverse a 2.4 Kbit/s link
13	10789 bytes	2246 bytes	10.0 s
7	5904 bytes	1707 bytes	8.0 s
5	4322 bytes	1509 bytes	7.2 s

Table 1 Sending filtered NFFI messages over an emulated tactical network

[1] The NIST Net software is freely available at "http://www-x.antd.nist.gov/nistnet/".

The exact update frequencies that can be used in an operational scenario will of course depend on the number of tracks that fall within the various zones of the filter, and the total bandwidth available. At NATO CWID we were operating in a controlled, experimental environment. The numbers given here should thus be considered as an example intended to illustrate the effects of the different types of filters.

The information overflow problem

As mentioned in the introduction, international military operations cause an increase of available information. This can cause problems not only on the network level, but also when it comes to the user's ability to process the information she receives, known as information overflow. As noted in [6], the predominant problems associated with overload of information is that there is more information available than can be absorbed and understood within a time span of any single individual. This can cause the recipient to overlook critical information.

Having some information overload is not necessarily bad: A skilled user, with the proper training, can learn to overcome information overload and in a team, such information can be shared. This corresponds to the thoughts in this older study [8], where it is pointed out that team members can perform better in high workload situations when there is a partial overlap in roles between them. So, with trained personnel, information overload may not be a major problem for a team to perform its tasks efficiently. However, for less trained personnel, a pre-processing of data prior to dissemination can help guide them towards making the right decisions. It is crucial to understand the difference between recognizing and ignoring significant information that can result in either making an informed, strong decision or an ill-informed one [6]. The content filtering techniques described in this paper may also be used to alleviate the information overflow problem, but a full evaluation of the effects content filtering has on information overflow is beyond the scope of this paper.

Conclusion

Using content filtering is one of several measures that can be employed in order to increase the information infrastructure's ability to adapt to changing network and battlefield conditions.

To maintain information superiority in a coalition force it is paramount that all necessary and relevant information is disseminated throughout the network. Thus, it becomes important to identify the information that is indeed relevant, and transmit only that over the network. As a means to achieve this, we have discussed several different types of filtering: Geographical filtering, Frequency based filtering, Priority based filtering, and more. Furthermore, the filtering can be of two types, in that one can filter entire messages, or parts of messages.

Filtering in the end system should be avoided since it wastes network resources, especially on the tactical level where bandwidth is scarce. Filtering in the message producing system may be the best option. On the other hand, proxies have an important benefit over that of filtering in the producer: If the producer sends information to a many recipients, all with difference capabilities, then it must filter once for each type of recipient. Since a proxy will typically be placed at the edge of a network, it is much more likely that the recipients in this network will be of similar types. For instance, all users in a tactical network are likely to have similar limitations, such as low available bandwidth and limited power supply. The means that a proxy often has fewer different recipient types to filter for, which in turn mans that the proxy filtering implementation can be simpler.

As a proof-of-concept we have presented our experiments from NATO CWID 2007, where we successfully tested our implementation of a combined geographical and frequency based filter for track information in a proxy. The filter was based on a set of rings with different frequency assigned to each ring. We called this a *zone ring filter*.

References

[1] Gibb, H. Fassbender, M. Schmeing, J. Michalak, J. E. Wieselthier, *Information Management over Disadvantaged Grids*, Final report of the RTO Information Systems Technology Panel, Task Group IST-030 / RTG-012, RTO-TR-IST-030, 2007

[2] R. Lausund, S. Martini, *Norwegian Modular Network Soldier (NORMANS)*, FFI Facts, November 2006 http://www.mil.no/multimedia/archive/00086/FFI-FACTS-NORMANS-EN_86447a.pdf

[3] Defence R&D Canada Valcartier, *High Capability Tactical Communications Network-HCTCN TD*, DRDC Valcartier Fact Sheet IS-226-A http://www.valcartier.rddc-drdc.gc.ca/poolpdf/e/164_e.pdf

[4] R. Haakseth, T. Gagnes, D. Hadzic, T. Hafsøe, F.T. Johnsen, K. Lund, B.K Reitan, *Experiment Report: "SOA – Cross Domain and Disadvantaged Grids" – NATO CWID 2007*, FFI-Raport 2007/02301, ISBN 978-82-464-1272-6

[5] T. Hafsøe, F.T. Johnsen, K. Lund, A. Eggen, *Adapting Web Services to for Limited Bandwidth Tactical Networks*, 12th ICCRTS, June 2007, Newport, RI, USA

[6] J. Carreno, G. Galdorisi, R. Goshorn, A. Siordia, *Maintaining Situational Awareness in Large, Complex Organizations*, 11th ICCRTS, 2006, Cambridge UK

[7] J.E. Miller, J. Dussault, *Accessing and Sharing: Facets of Addressing Information Overload*, 11th ICCRTS, 2006, Cambridge UK

[8] D.L. Kleinman, D. Serfaty, *Team performance assessment in distributed decision making*. In: Gibson, R., Kincaid, J.P. and Goldiez, B. Editors, 1989.Proceedings of the Interservice Networked Simulation for Training Conference University of Central Florida, Orlando, FL.