FFI-rapport 2008/01430

Publications related to NATO CWID 2007 SOA experiments

Frank T. Johnsen, Anders Eggen, Trude Hafsøe and Ketil Lund

Forsvarets forskningsinstitutt/Norwegian Defence Research Establishment (FFI)

11 July 2008

FFI-rapport 2008/01430

1086

P: ISBN 978-82-464-1411-9
E: ISBN 978-82-464-1412-6

Keywords

NATO CWID

Nettverksbasert Forsvar

Eksperimentering

Web services

Approved by

Anders Eggen	Project manager
Vidar S. Andersen	Director

Sammendrag

Denne rapporten inneholder de tre artiklene relatert til eksperimentene på NATO CWID 2007 som Sikker gjennomgående SOA (prosjekt 1086) har publisert på internasjonale konferanser.

English summary

This report contains three articles describing various parts of Secure Pervasive SOA's experiments at NATO CWID in 2007. These articles have been peer reviewed and published at international conferences.

Contents

1 Introduction

This report contains the three articles we have written concerning parts of our experiments at NATO CWID in 2007. The articles have been peer reviewed and published at international conferences; one at the NATO RTO/IST-083 symposium in Prague in April 2008, and the two others at the 13th ICCRTS in Seattle, WA in June 2008.

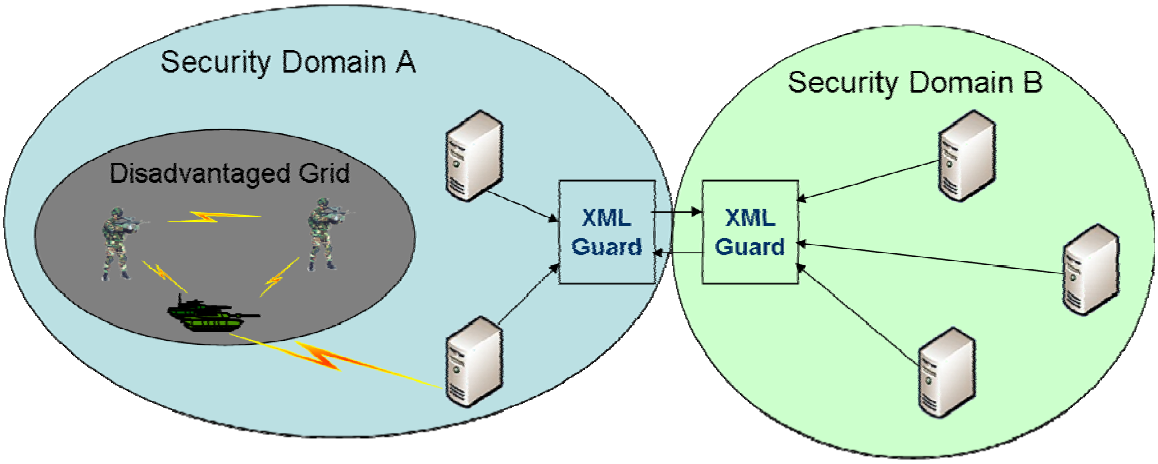

Figure 1.1 NATO CWID demonstrator scenario

Our experiments were divided into two main focus areas, as shown in Figure 1.1. One experiment part was concerned with XML security and guard solutions, while another part was concerned with the use of XML and Web services in disadvantaged grids[1]. It is the latter part that is the focus of this report, as the three publications discuss different parts of this experiment in detail:

1. Traditional Web services use SOAP over HTTP over TCP/IP. Tactical networks have different characteristics from those of the Internet, and thus can benefit from using solutions tailored to such networks. In military networks one should have a solution that can be used across different networks, and that supports store-and-forward. The latter is lacking in traditional Web services. We suggest that one can use tactical protocols as a transport for Web services in such cases, and have experimented with MMHS as a carrier. Our results are presented in Appendix A.
2. Replacing the communication protocol as described above is only one way to reduce communication overhead. We have also experimented with other methods, such as content filtering. The idea behind using content filtering is that you can remove information that is not useful, and thus reduce the amount of information that needs to be transmitted over the network. See Appendix B for our publication regarding content filtering.
3. Optimizing the transport protocol and filtering content will, as we have shown, help reduce communication overhead. However, these techniques will only solve parts of

[1] Disadvantaged grids are networks with high delay, low bandwidth, mobile units and frequent disruptions, i.e. the military communication networks at the lowest tactical level.

the problem. Since Web services are based on XML, a text-based format, the information that needs to be exchanged will be quite verbose compared to a binary format. However, the drawbacks of XML can be circumvented by using compression. See Appendix C for our evaluation of different compression mechanisms.

For an overview of the entire NATO CWID experiment, see **FFI/Rapport 2007/02301**. Further details about the disadvantaged grid experiment can be found in **FFI/Notat 2007/02063**. Please note that the articles presented in this report discuss their respective topics in more detail than those previous reports.

Appendix A Publication at IST-083 regarding MMHS as a transport for Web services

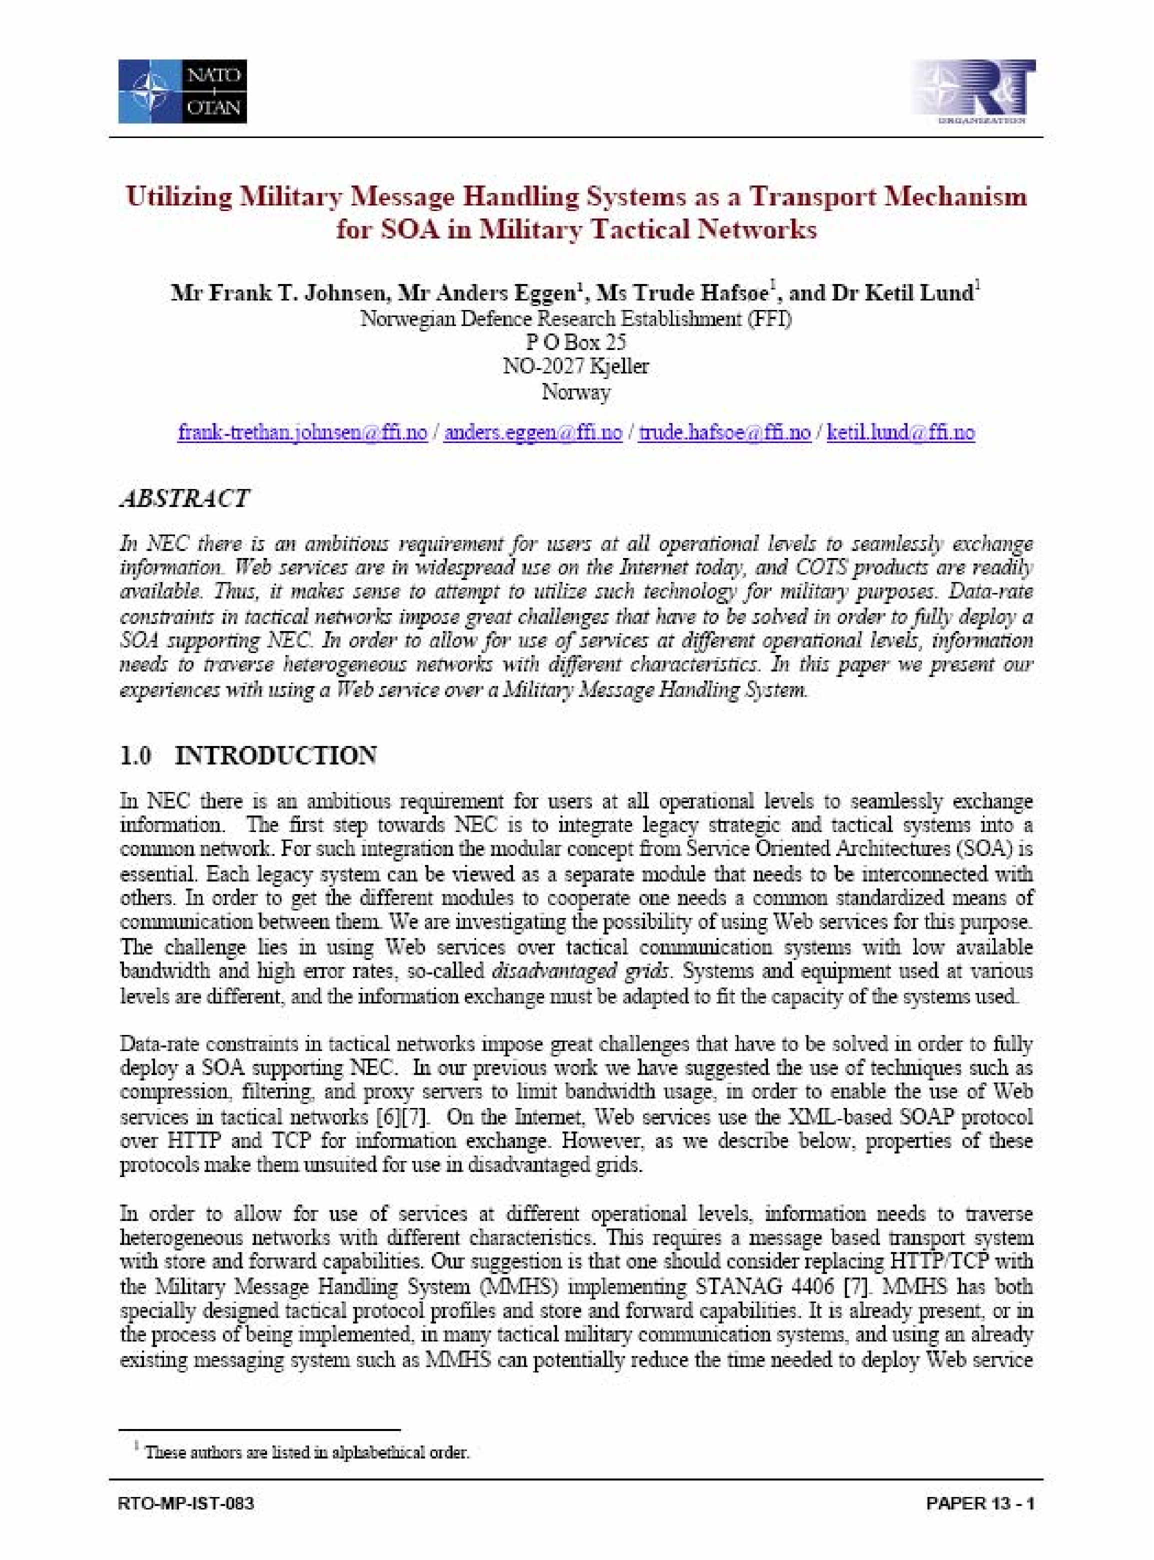

Utilizing Military Message Handling Systems as a Transport Mechanism for SOA in Military Tactical Networks

Mr Frank T. Johnsen, Mr Anders Eggen[1], Ms Trude Hafsoe[1], and Dr Ketil Lund[1]
Norwegian Defence Research Establishment (FFI)
P O Box 25
NO-2027 Kjeller
Norway

frank-trethan.johnsen@ffi.no / anders.eggen@ffi.no / trude.hafsoe@ffi.no / ketil.lund@ffi.no

ABSTRACT

In NEC there is an ambitious requirement for users at all operational levels to seamlessly exchange information. Web services are in widespread use on the Internet today, and COTS products are readily available. Thus, it makes sense to attempt to utilize such technology for military purposes. Data-rate constraints in tactical networks impose great challenges that have to be solved in order to fully deploy a SOA supporting NEC. In order to allow for use of services at different operational levels, information needs to traverse heterogeneous networks with different characteristics. In this paper we present our experiences with using a Web service over a Military Message Handling System.

1.0 INTRODUCTION

In NEC there is an ambitious requirement for users at all operational levels to seamlessly exchange information. The first step towards NEC is to integrate legacy strategic and tactical systems into a common network. For such integration the modular concept from Service Oriented Architectures (SOA) is essential. Each legacy system can be viewed as a separate module that needs to be interconnected with others. In order to get the different modules to cooperate one needs a common standardized means of communication between them. We are investigating the possibility of using Web services for this purpose. The challenge lies in using Web services over tactical communication systems with low available bandwidth and high error rates, so-called *disadvantaged grids*. Systems and equipment used at various levels are different, and the information exchange must be adapted to fit the capacity of the systems used.

Data-rate constraints in tactical networks impose great challenges that have to be solved in order to fully deploy a SOA supporting NEC. In our previous work we have suggested the use of techniques such as compression, filtering, and proxy servers to limit bandwidth usage, in order to enable the use of Web services in tactical networks [6][7]. On the Internet, Web services use the XML-based SOAP protocol over HTTP and TCP for information exchange. However, as we describe below, properties of these protocols make them unsuited for use in disadvantaged grids.

In order to allow for use of services at different operational levels, information needs to traverse heterogeneous networks with different characteristics. This requires a message based transport system with store and forward capabilities. Our suggestion is that one should consider replacing HTTP/TCP with the Military Message Handling System (MMHS) implementing STANAG 4406 [7]. MMHS has both specially designed tactical protocol profiles and store and forward capabilities. It is already present, or in the process of being implemented, in many tactical military communication systems, and using an already existing messaging system such as MMHS can potentially reduce the time needed to deploy Web service

[1] These authors are listed in alphabethical order.

based solutions in tactical networks.

At NATO CWID, an annual venue for interoperability demonstration and experimentation, we implemented a Web service using MMHS as a transport layer as a part of our experiments. In this paper we present our experiences with using a Web service over MMHS, based on our experiments at NATO CWID 2007.

The remainder of this paper is structured as follows: First, we discuss HTTP and TCP, the most common means of transport for SOAP messages in Web services, with emphasis on the drawbacks of these protocols in tactical networks. Then, we proceed to introduce STANAG 4406, and discuss the benefits of employing tactical transport protocols. These two parts form the theoretical foundation of our paper, after which we present the experiments we performed at NATO CWID where we used MMHS as a carrier for Web services traffic. The experiments were a success, thus functioning as a proof of concept. A summary section highlighting our most important findings concludes the paper.

2.0 HTTP AND TCP IN DISADVANTAGED GRIDS

The use of HTTP over TCP originates from the World Wide Web, and has later been adopted as the primary protocol combination for Web services. The SOAP messages exchanged between Web services clients and servers are sent using HTTP, which in turn utilizes TCP for reliable transfer of the messages.

HTTP is synchronous, which means that when a SOAP request is sent, the HTTP connection is kept open until the SOAP response is returned in the HTTP "acknowledgement". If the connection times out because of delays or for any other reason, there will be a problem routing the SOAP response back to the service consumer. Consequently, HTTP will not work well when used in disadvantaged grids or in a combination of heterogeneous networks. Furthermore, in disruptive networks TCP connections break, thus making the protocol less suited to the tactical environment. Such networks require asynchronous communications and protocols that are able to cope with the characteristics (e.g., data rates, delays, frequency of disconnections) of military communication networks:

- Protocols that can withstand long and variable round trip times, while at the same time having very little communication overhead.

- Store and forward capabilities, where intermediate nodes can store a message until it can be delivered to the recipient rather than discarding the message if immediate delivery is not possible.

The store and forward capability is needed for two reasons: Users connected through a disadvantaged grid can experience frequent but short communication disruptions, which can prevent a message from being delivered immediately. Having store and forward support can ensure that the message is not dropped. In addition, store and forward can be used in gateways between network types to compensate for differences in link capacity between the networks. An ordinary router is at risk of having to drop packets due to its buffers filling up faster than the packets can be transmitted out onto the lower capacity network.

3.0 STANAG 4406

In NATO, Formal Military Messaging is standardized in STANAG 4406 ed. 2 (S4406). A MMHS is responsible for the delivery, formal audit, archiving, numbering, release, emission, security, and distribution of received formal messages. In NATO, the formal messaging service is seen as the vehicle for secure mission-critical operational, military applications (email systems are not). S4406 Edition 1 is the only agreed standard to achieve interoperability between the formal messaging systems of NATO nations. Systems compatible with the S4406 standard have been and are being implemented widely by the

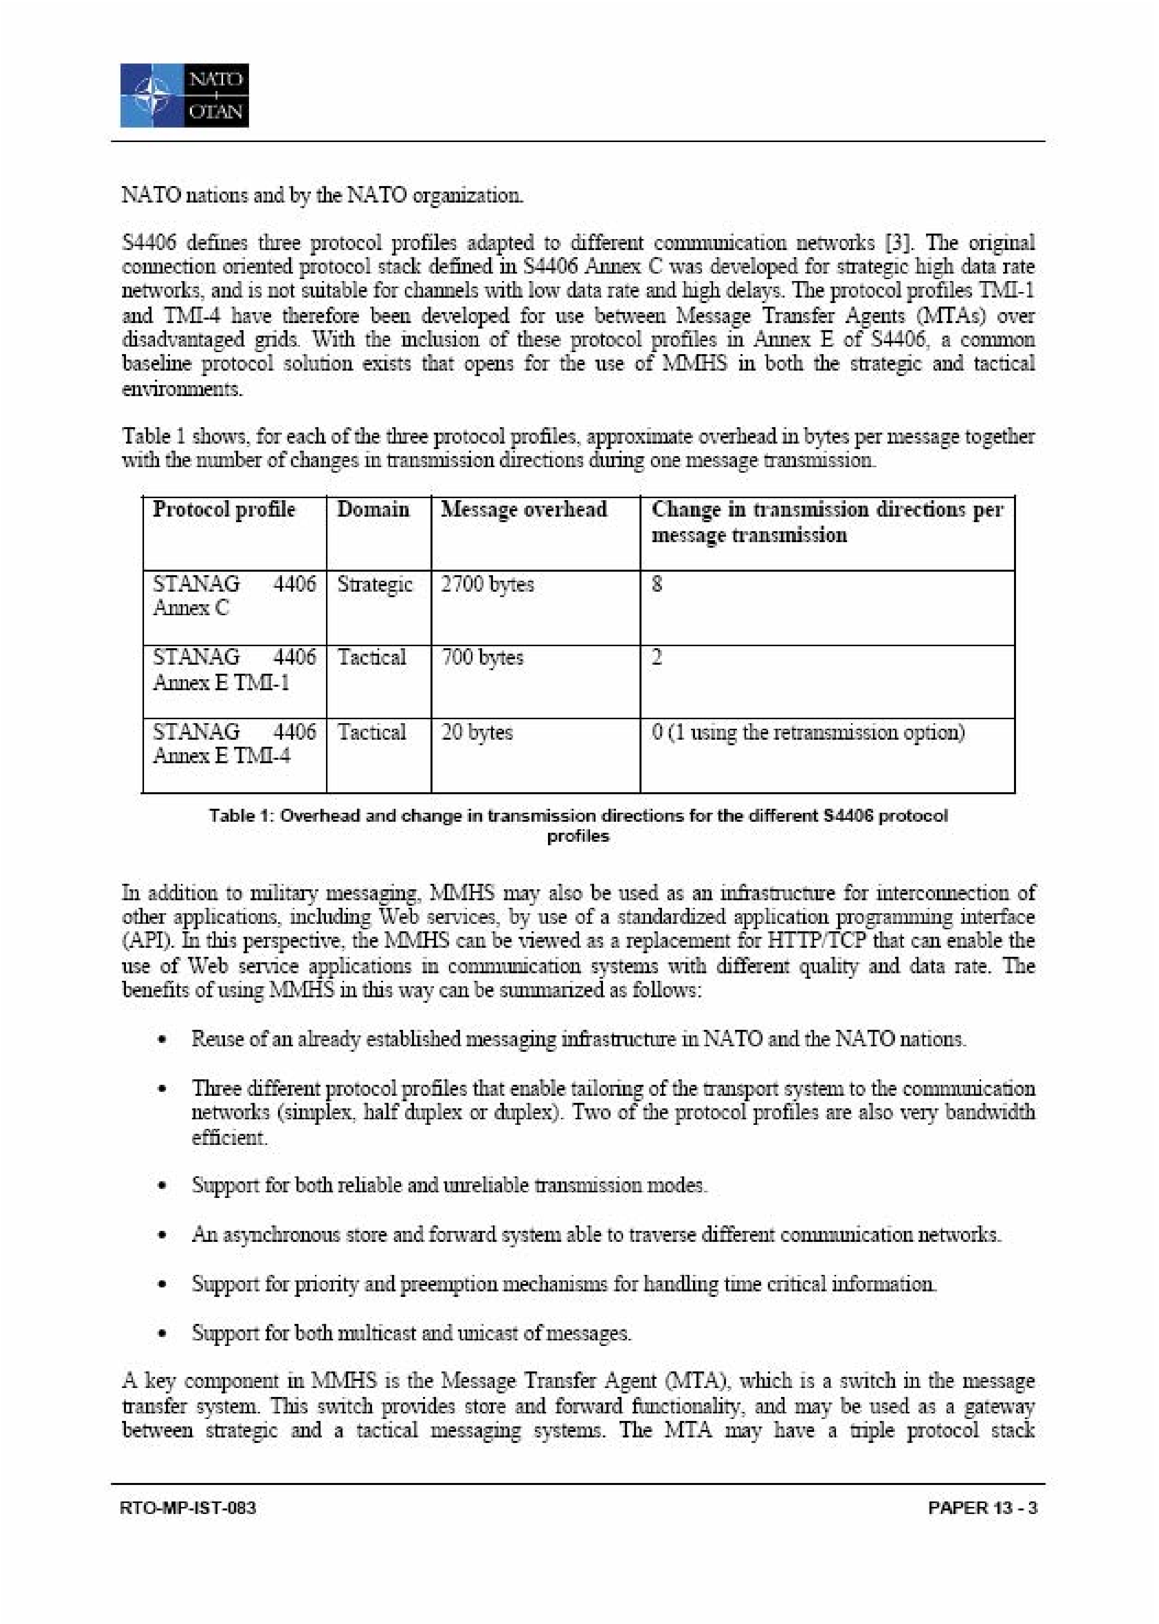

NATO nations and by the NATO organization.

S4406 defines three protocol profiles adapted to different communication networks [3]. The original connection oriented protocol stack defined in S4406 Annex C was developed for strategic high data rate networks, and is not suitable for channels with low data rate and high delays. The protocol profiles TMI-1 and TMI-4 have therefore been developed for use between Message Transfer Agents (MTAs) over disadvantaged grids. With the inclusion of these protocol profiles in Annex E of S4406, a common baseline protocol solution exists that opens for the use of MMHS in both the strategic and tactical environments.

Table 1 shows, for each of the three protocol profiles, approximate overhead in bytes per message together with the number of changes in transmission directions during one message transmission.

Protocol profile	**Domain**	**Message overhead**	**Change in transmission directions per message transmission**
STANAG 4406 Annex C	Strategic	2700 bytes	8
STANAG 4406 Annex E TMI-1	Tactical	700 bytes	2
STANAG 4406 Annex E TMI-4	Tactical	20 bytes	0 (1 using the retransmission option)

Table 1: Overhead and change in transmission directions for the different S4406 protocol profiles

In addition to military messaging, MMHS may also be used as an infrastructure for interconnection of other applications, including Web services, by use of a standardized application programming interface (API). In this perspective, the MMHS can be viewed as a replacement for HTTP/TCP that can enable the use of Web service applications in communication systems with different quality and data rate. The benefits of using MMHS in this way can be summarized as follows:

- Reuse of an already established messaging infrastructure in NATO and the NATO nations.
- Three different protocol profiles that enable tailoring of the transport system to the communication networks (simplex, half duplex or duplex). Two of the protocol profiles are also very bandwidth efficient.
- Support for both reliable and unreliable transmission modes.
- An asynchronous store and forward system able to traverse different communication networks.
- Support for priority and preemption mechanisms for handling time critical information.
- Support for both multicast and unicast of messages.

A key component in MMHS is the Message Transfer Agent (MTA), which is a switch in the message transfer system. This switch provides store and forward functionality, and may be used as a gateway between strategic and a tactical messaging systems. The MTA may have a triple protocol stack implementing both the strategic connection oriented protocol profile and the two tactical protocol profiles, and can therefore route messages between infrastructure WANs and low data rate tactical links (see Figure 1). When using the MMHS for transfer of SOAP messages, the MTAs and the Web service functionality are integrated, and there is therefore no additional delay for checking or connecting to a Message Store[2].

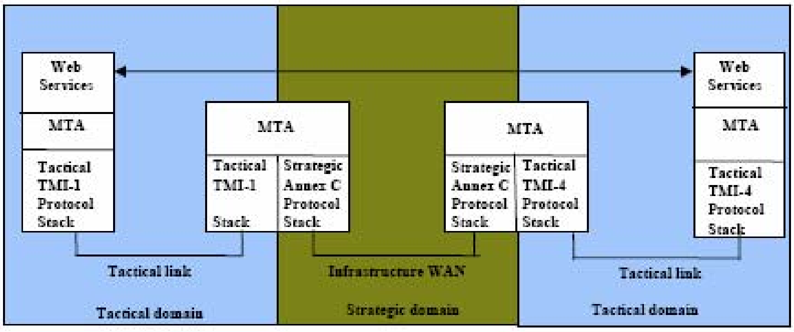
Web Services
MTA
Tactical TMI-1 Protocol Stack
MTA
Tactical TMI-1 Stack
Strategic Annex C Protocol Stack
MTA
Strategic Annex C Protocol Stack
Tactical TMI-4 Protocol Stack
Web Services
MTA
Tactical TMI-4 Protocol Stack
Tactical link
Infrastructure WAN
Tactical link
Tactical domain
Strategic domain
Tactical domain

Figure 1: Seamless interconnection of Web Services over heterogeneous communication networks by using MMHS as an overlay network

4.0 EXPERIMENT SETUP

Because of the interoperability benefits of using Web services, we want to extend its use as far out on the tactical level as possible. In our work on using Web services over disadvantaged grids, we focus on the upper layers of the protocol stack (i.e., application and transport layers). More specifically, we focus on efficient information representation and compression of XML, in addition to the use of MMHS as a transport mechanism.

In Figure 2 we show a simplified view of our experiment setup at NATO CWID 2007. The machine running the NORMANS software was connected to the local HQ through a link emulator. It was configured in such a way that network traffic between the local HQ and NORMANS was slowed down, yielding a link with 2.4 Kbps capacity. This functioned as our disadvantaged grid in the experiments we performed.

We have tested different data models in order to achieve an efficient information representation. This year at NATO CWID we have been using XML encoded NATO Friendly Force Information (NFFI) [2] data, which is a relatively compact format. On the Web service layer we have used compression, in order to reduce the size the SOAP messages that are passed between the systems. In particular, we have used Efficient XML from Agile Delta, which has proven to achieve high compression ratios.

[2] All sent and received messages are placed into one or more storages. This is determined by originator and recipient(s) of the messages. Every user has a personal storage, and every organizational unit has an official storage.

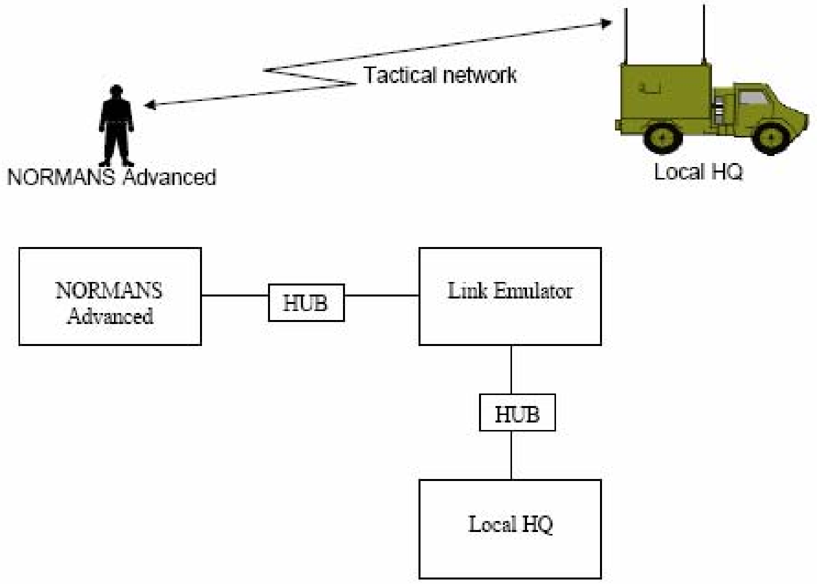

Figure 2: NATO CWID disadvantaged grid setup (scenario on top, actual experiment setup below)

Below we give a brief overview of the different experiment components. For further details about configuring and using XOmail, a link emulator, the NORMANS software, and our wrapper software, see [1].

4.1 NORMANS

The focus of our experiment was on Web services and interoperability. Thus, we chose to SOA-enable an experimental tactical system, the Norwegian Modular Network Soldier (NORMANS). We wrapped the NORMANS software in a Web service, which would communicate with the local HQ using standardized XML-encoded NFFI.

An overview of NORMANS is given in [4], and the C2I system is presented in [5]. NORMANS is a conceptual approach towards the future Norwegian soldier system. The concept includes use of legacy equipment and focuses on the need of integrating all components to a working system both for the individual soldier, and for the section and higher echelon units. The NORMANS C4I concept is based on voice and data communication within the sections using a simplified data transmission protocol. This proprietary protocol currently does not facilitate interoperability with other nations, a key concept at NATO CWID. Thus, the modified version of the software does not communicate via this protocol but by input and output of XML formatted data. The NORMANS system would report its own position to the local HQ, and receive track updates for its area of operation.

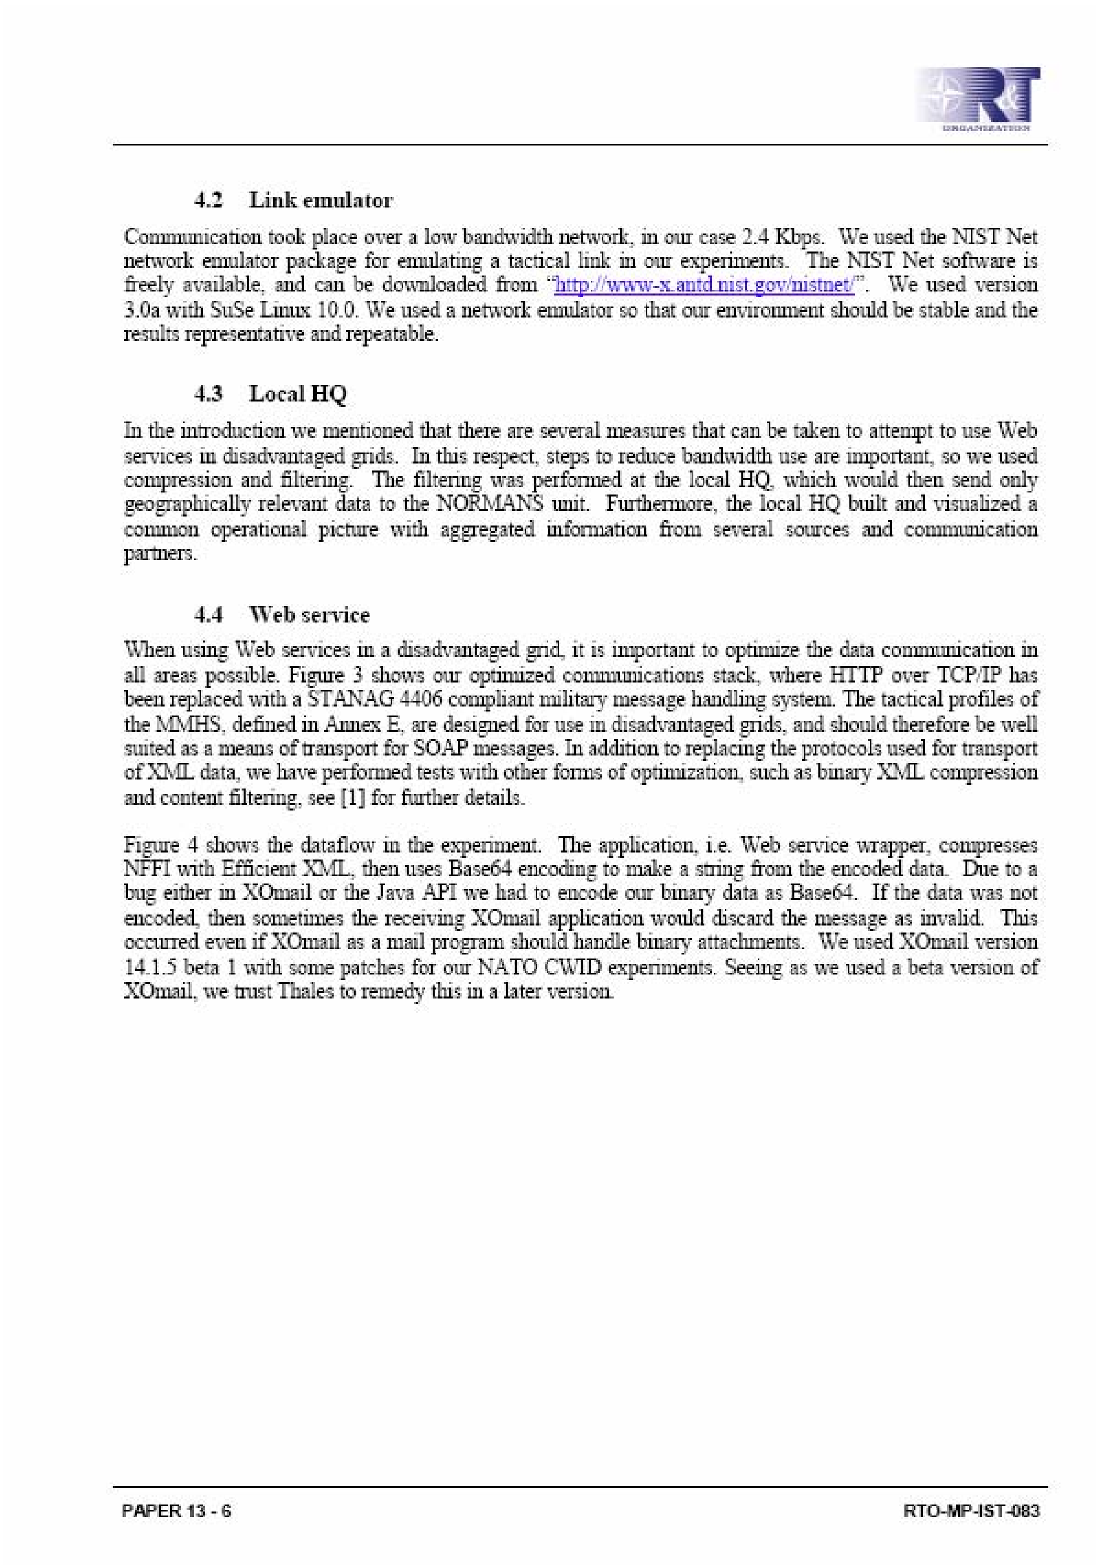

4.2 Link emulator

Communication took place over a low bandwidth network, in our case 2.4 Kbps. We used the NIST Net network emulator package for emulating a tactical link in our experiments. The NIST Net software is freely available, and can be downloaded from "http://www-x.antd.nist.gov/nistnet/". We used version 3.0a with SuSe Linux 10.0. We used a network emulator so that our environment should be stable and the results representative and repeatable.

4.3 Local HQ

In the introduction we mentioned that there are several measures that can be taken to attempt to use Web services in disadvantaged grids. In this respect, steps to reduce bandwidth use are important, so we used compression and filtering. The filtering was performed at the local HQ, which would then send only geographically relevant data to the NORMANS unit. Furthermore, the local HQ built and visualized a common operational picture with aggregated information from several sources and communication partners.

4.4 Web service

When using Web services in a disadvantaged grid, it is important to optimize the data communication in all areas possible. Figure 3 shows our optimized communications stack, where HTTP over TCP/IP has been replaced with a STANAG 4406 compliant military message handling system. The tactical profiles of the MMHS, defined in Annex E, are designed for use in disadvantaged grids, and should therefore be well suited as a means of transport for SOAP messages. In addition to replacing the protocols used for transport of XML data, we have performed tests with other forms of optimization, such as binary XML compression and content filtering, see [1] for further details.

Figure 4 shows the dataflow in the experiment. The application, i.e. Web service wrapper, compresses NFFI with Efficient XML, then uses Base64 encoding to make a string from the encoded data. Due to a bug either in XOmail or the Java API we had to encode our binary data as Base64. If the data was not encoded, then sometimes the receiving XOmail application would discard the message as invalid. This occurred even if XOmail as a mail program should handle binary attachments. We used XOmail version 14.1.5 beta 1 with some patches for our NATO CWID experiments. Seeing as we used a beta version of XOmail, we trust Thales to remedy this in a later version.

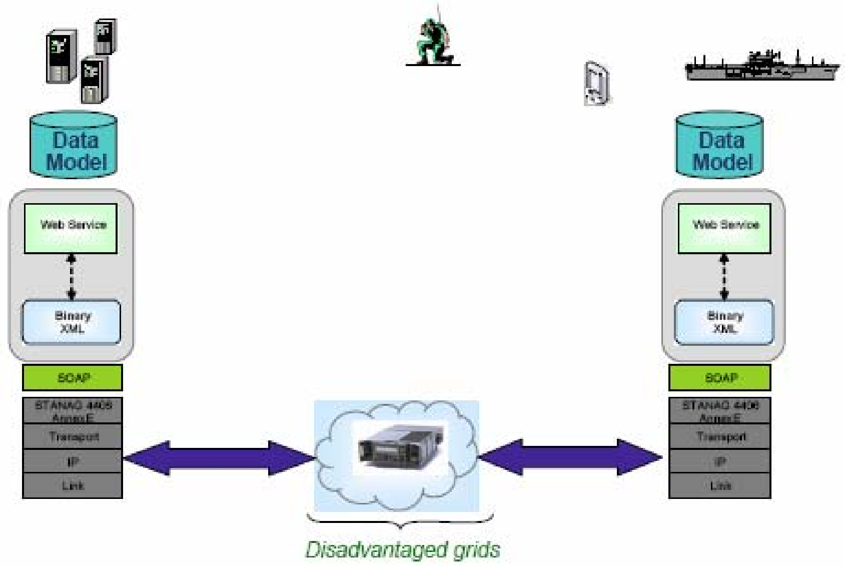

•

Figure 3: Communication between systems

After applying the Base64 encoding, the resulting string is then wrapped in a SOAP envelope. The resulting message is sent via MMHS, i.e. XOmail, using the send function in the Java API. XOmail then handles the actual transport, using one of the tactical transport protocols.

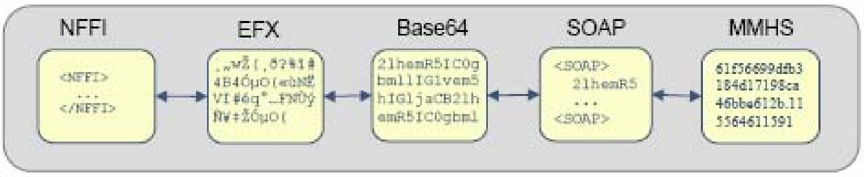

Figure 4: Dataflow in the experiment.

5.0 RESULTS

As a part of our interoperability experiments at NATO CWID we tested and evaluated the use of MMHS as a carrier for Web service traffic. MMHS has many of the qualities that are needed to ensure delivery of messages between strategic and tactical communication systems, and gives the benefit of being able to reuse an existing infrastructure for a new purpose.

Our goal in performing these tests was twofold; we wanted to evaluate the use of store and forward in general, and MMHS specifically. We also wanted to compare and evaluate the two tactical protocol profiles of S4406, TMI-1 and TMI-4, in order to investigate how the differences in overhead and directional changes affects transmission delay. Through our tests we have:

- Confirmed our hypothesis that MMHS can be used as a replacement carrier for Web services.
- Shown that using MMHS avoids the time-out problems that arise when using standard HTTP over TCP in tactical networks.
- Introduced the ability to use Web services even when communication links fail temporarily.

Furthermore, we were able to show that an existing (proprietary) service can be adapted to a SOA environment; and that by taking the necessary measures, Web services can be used in disadvantaged grids similar to the one we emulated.

We compared the two tactical profiles of S4406, TMI-1 and TMI-4, in order to establish the efficiency of the two profiles. For the measurements, we transferred documents containing NFFI-tracks, and we compressed the documents using efficient XML with built-in compression enabled. In order to simulate a disadvantaged grid, we used NIST Net configured with a bandwidth of 2.4 kbps.

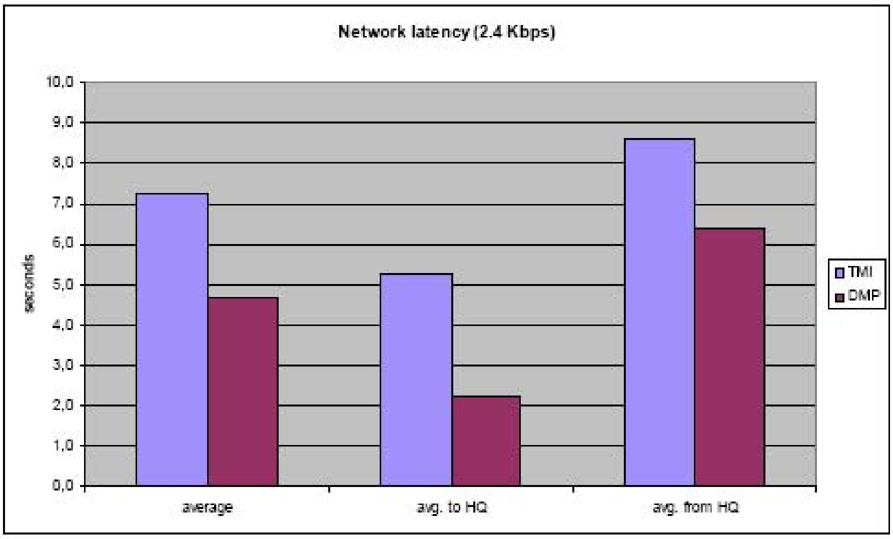

Figure 5: Message transmission delays

In Figure 5 we show the results of our experiments. The graph shows the overall average transfer time, the average transfer time from NORMANS Advanced to HQ, and the average transfer time from HQ to NORMANS Advanced. It should be noted that the documents sent to HQ contained only one NFFI-track (the soldier reporting own position), while the documents sent from HQ to NORMANS Advanced contained 20 to 25 NFFI-tracks each. Thus, the graph clearly shows the effect of compression; sending 20

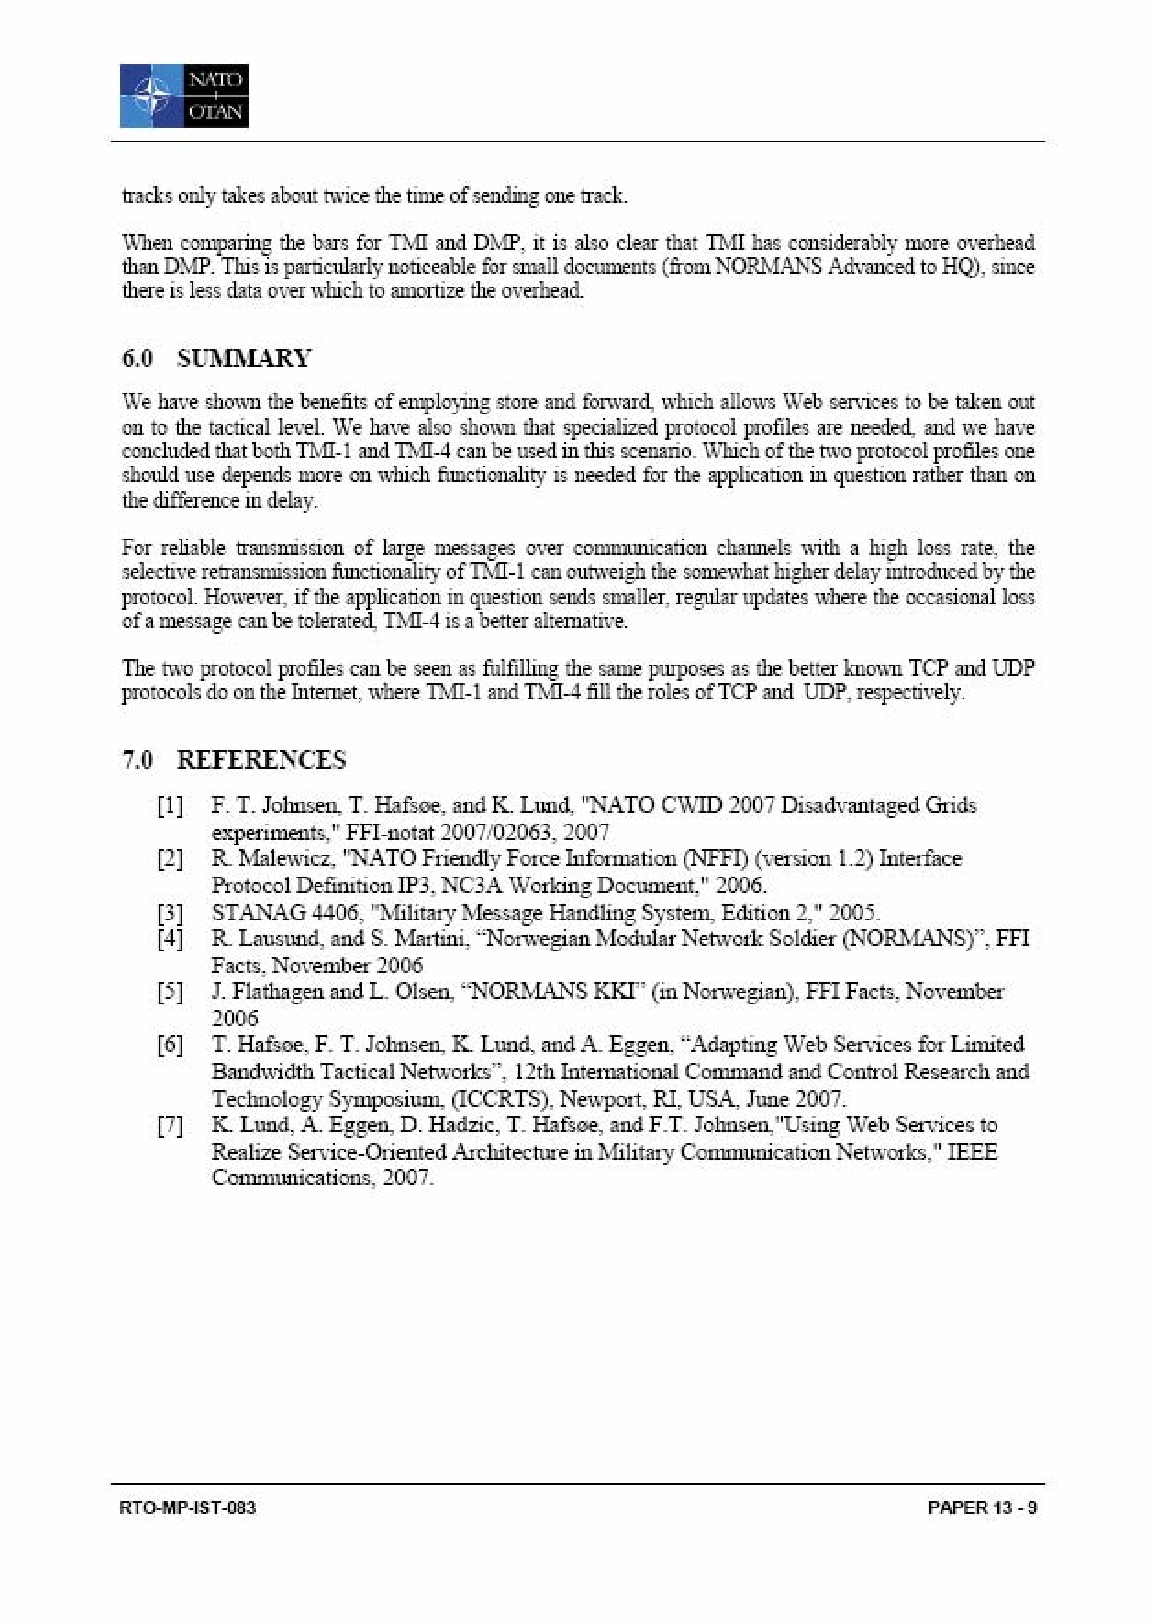

tracks only takes about twice the time of sending one track.

When comparing the bars for TMI and DMP, it is also clear that TMI has considerably more overhead than DMP. This is particularly noticeable for small documents (from NORMANS Advanced to HQ), since there is less data over which to amortize the overhead.

6.0 SUMMARY

We have shown the benefits of employing store and forward, which allows Web services to be taken out on to the tactical level. We have also shown that specialized protocol profiles are needed, and we have concluded that both TMI-1 and TMI-4 can be used in this scenario. Which of the two protocol profiles one should use depends more on which functionality is needed for the application in question rather than on the difference in delay.

For reliable transmission of large messages over communication channels with a high loss rate, the selective retransmission functionality of TMI-1 can outweigh the somewhat higher delay introduced by the protocol. However, if the application in question sends smaller, regular updates where the occasional loss of a message can be tolerated, TMI-4 is a better alternative.

The two protocol profiles can be seen as fulfilling the same purposes as the better known TCP and UDP protocols do on the Internet, where TMI-1 and TMI-4 fill the roles of TCP and UDP, respectively.

7.0 REFERENCES

[1] F. T. Johnsen, T. Hafsøe, and K. Lund, "NATO CWID 2007 Disadvantaged Grids experiments," FFI-notat 2007/02063, 2007

[2] R. Malewicz, "NATO Friendly Force Information (NFFI) (version 1.2) Interface Protocol Definition IP3, NC3A Working Document," 2006.

[3] STANAG 4406, "Military Message Handling System, Edition 2," 2005.

[4] R. Lausund, and S. Martini, "Norwegian Modular Network Soldier (NORMANS)", FFI Facts, November 2006

[5] J. Flathagen and L. Olsen, "NORMANS KKI" (in Norwegian), FFI Facts, November 2006

[6] T. Hafsøe, F. T. Johnsen, K. Lund, and A. Eggen, "Adapting Web Services for Limited Bandwidth Tactical Networks", 12th International Command and Control Research and Technology Symposium, (ICCRTS), Newport, RI, USA, June 2007.

[7] K. Lund, A. Eggen, D. Hadzic, T. Hafsøe, and F.T. Johnsen,"Using Web Services to Realize Service-Oriented Architecture in Military Communication Networks," IEEE Communications, 2007.

Appendix B Publication at ICCRTS regarding content filtering as a means to reduce communication overhead

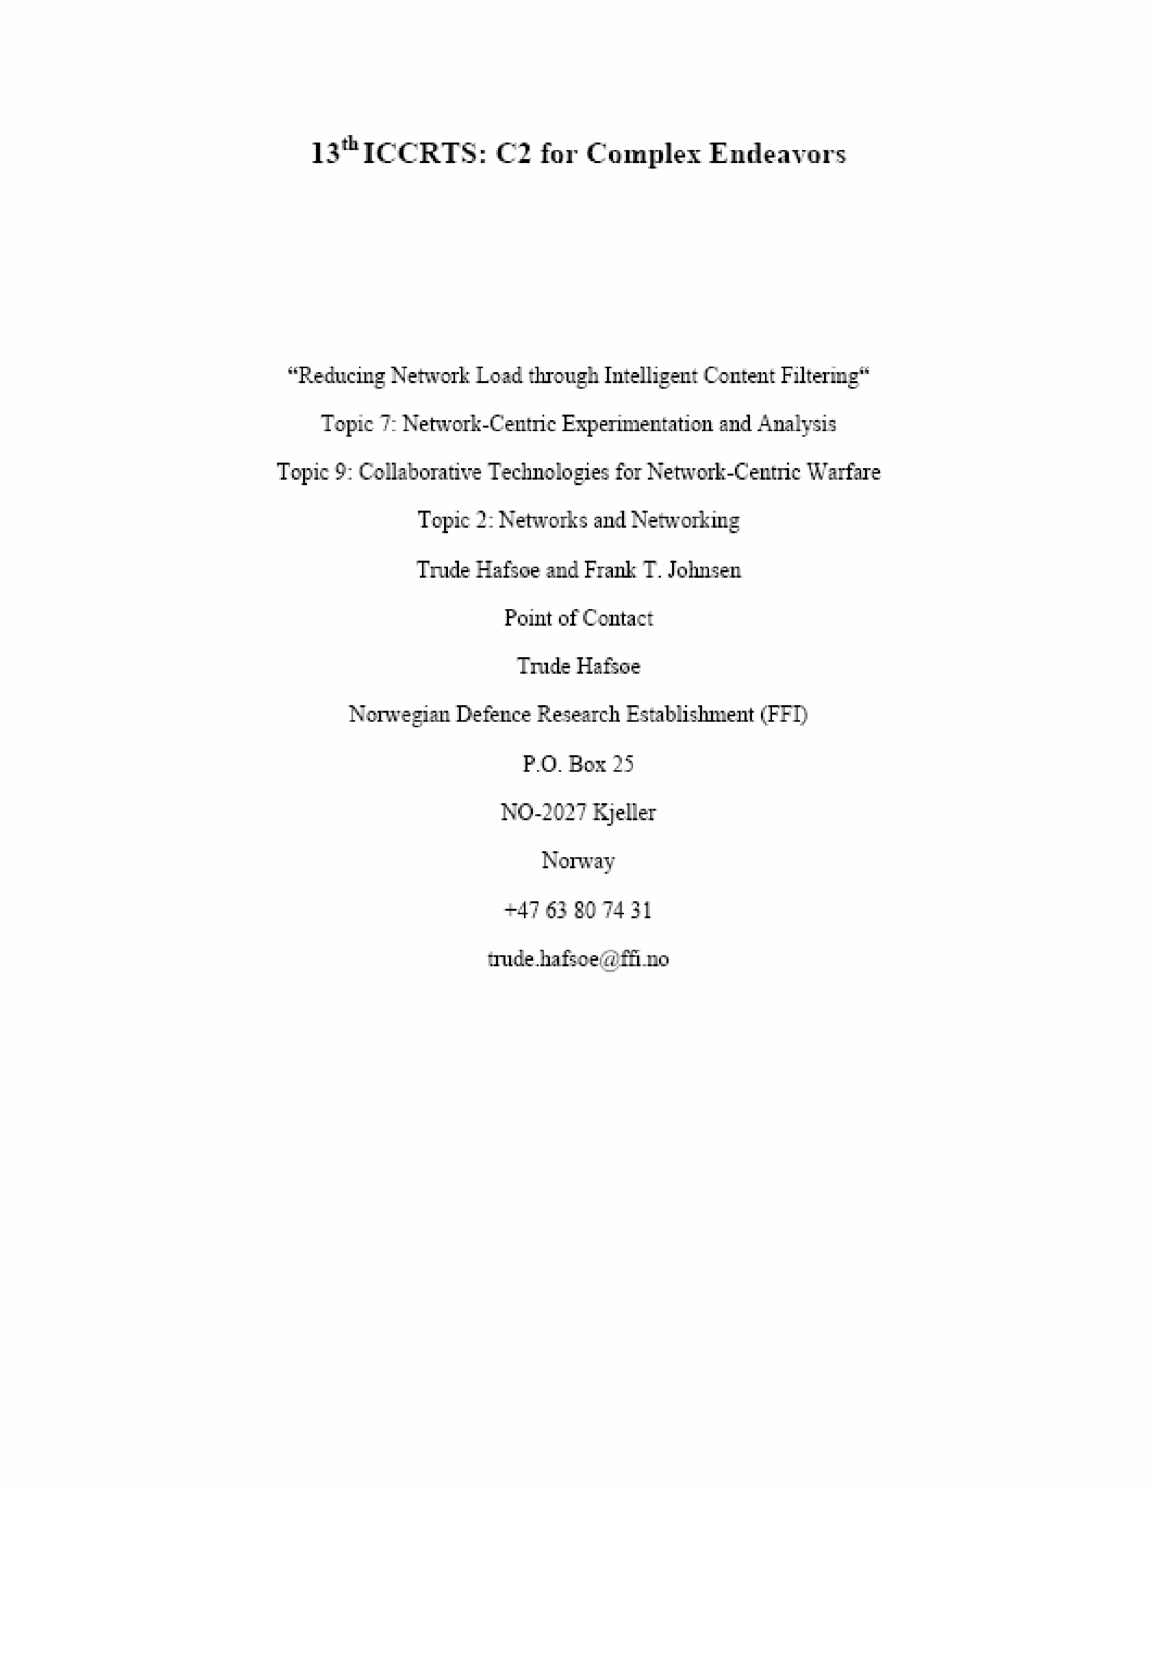

13th ICCRTS: C2 for Complex Endeavors

"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 Hafsoe 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

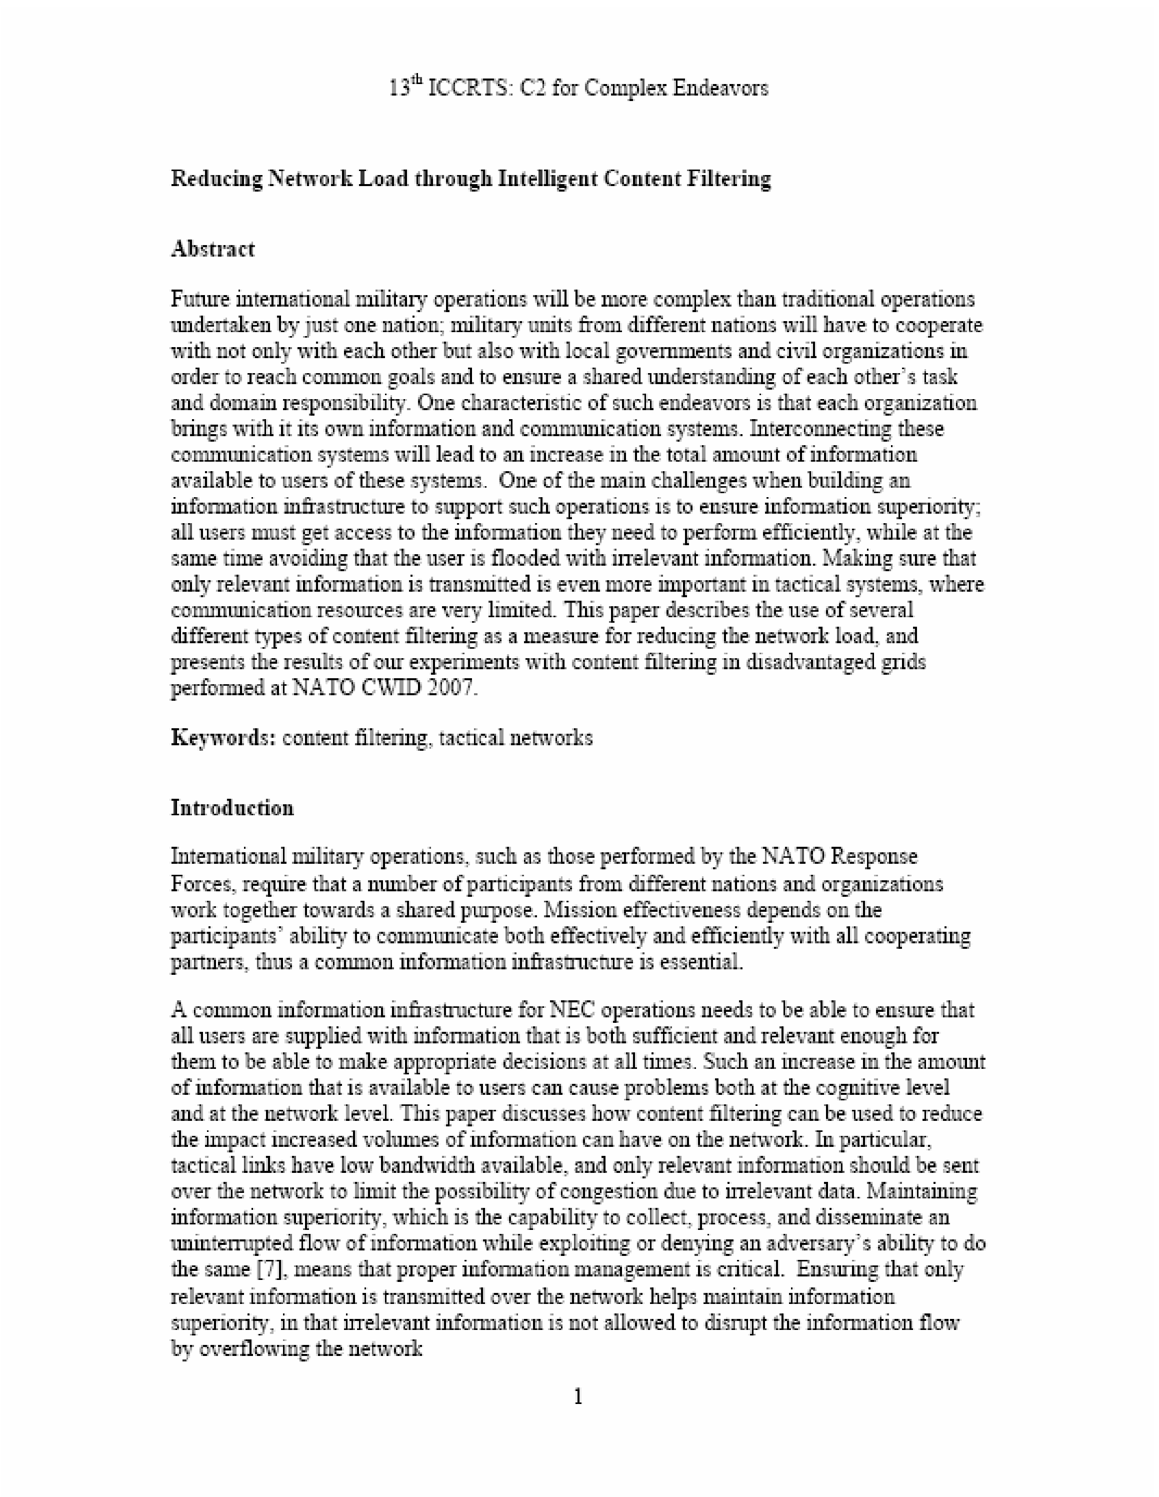

Reducing Network Load through Intelligent Content Filtering

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

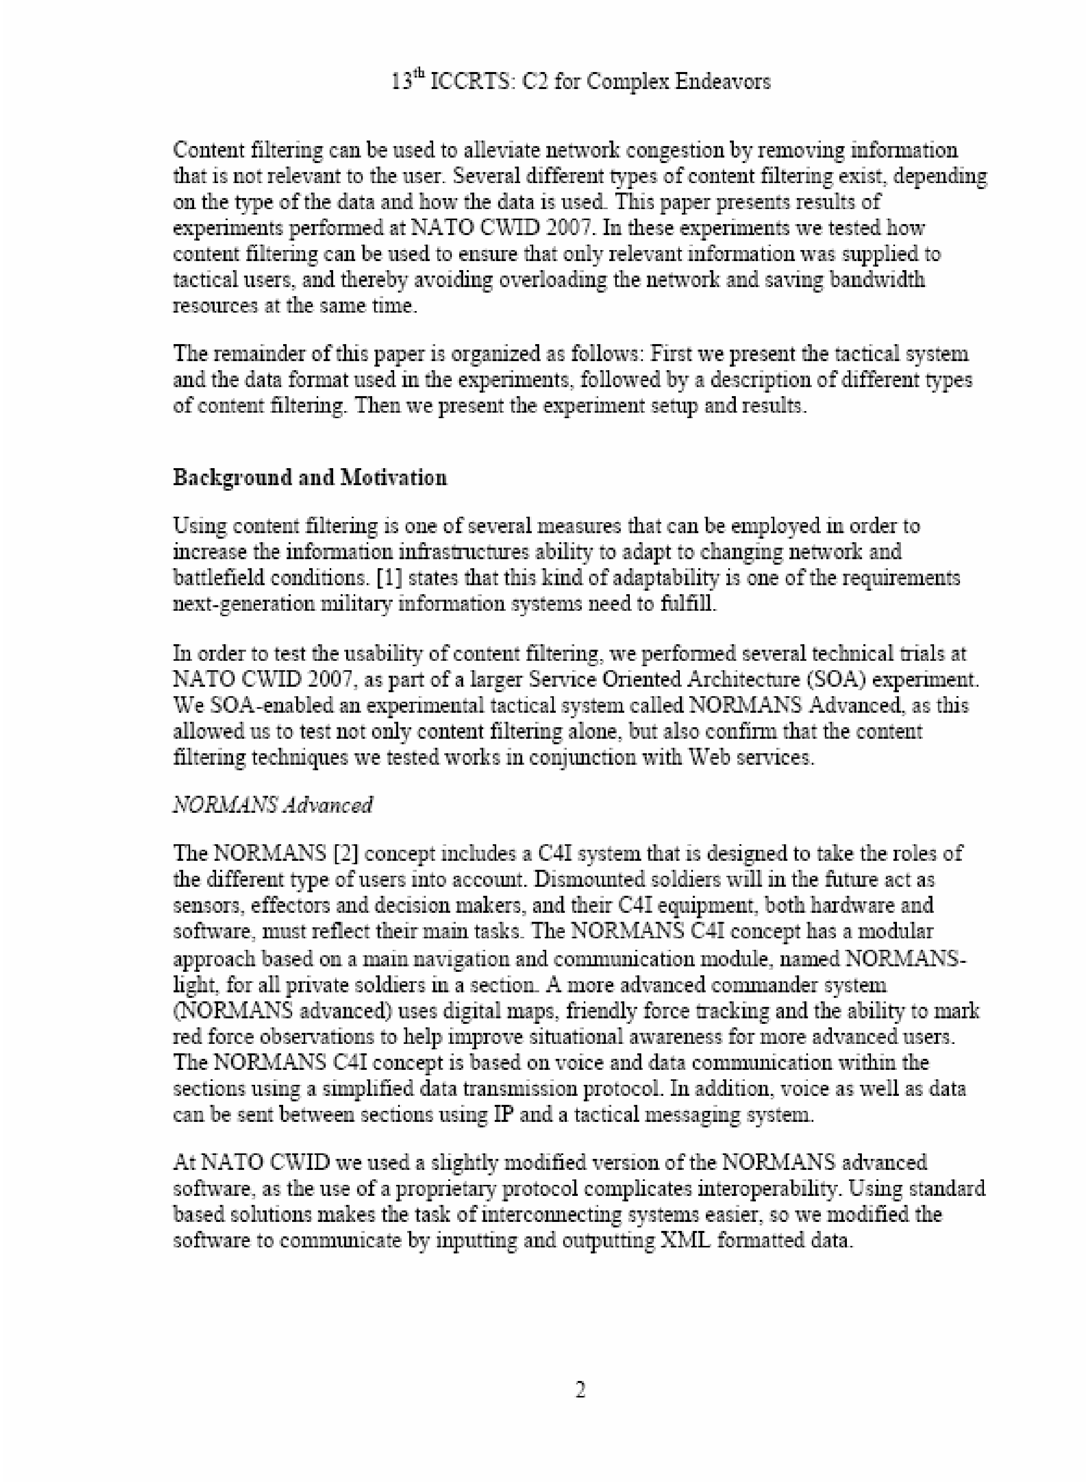

Content filtering can be used to alleviate network congestion by removing information that is not relevant to the user. Several different types of content filtering exist, depending on the type of the data and how the data is used. This paper presents results of experiments performed at NATO CWID 2007. 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.

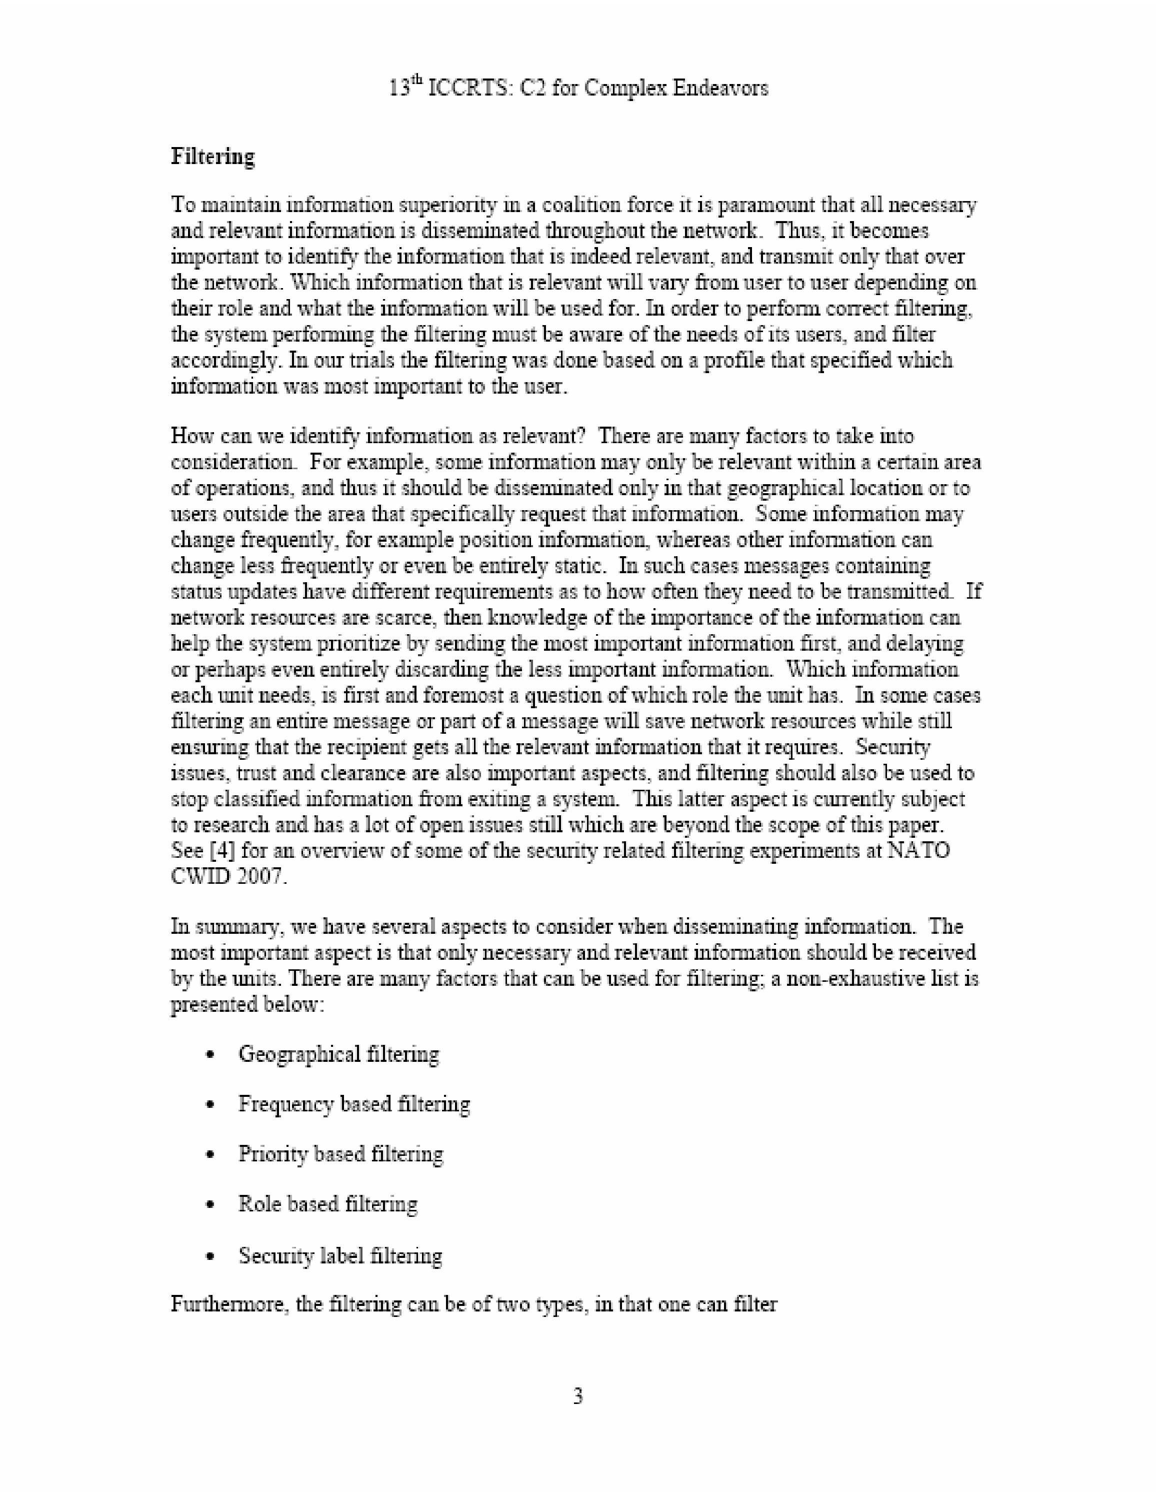

Filtering

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

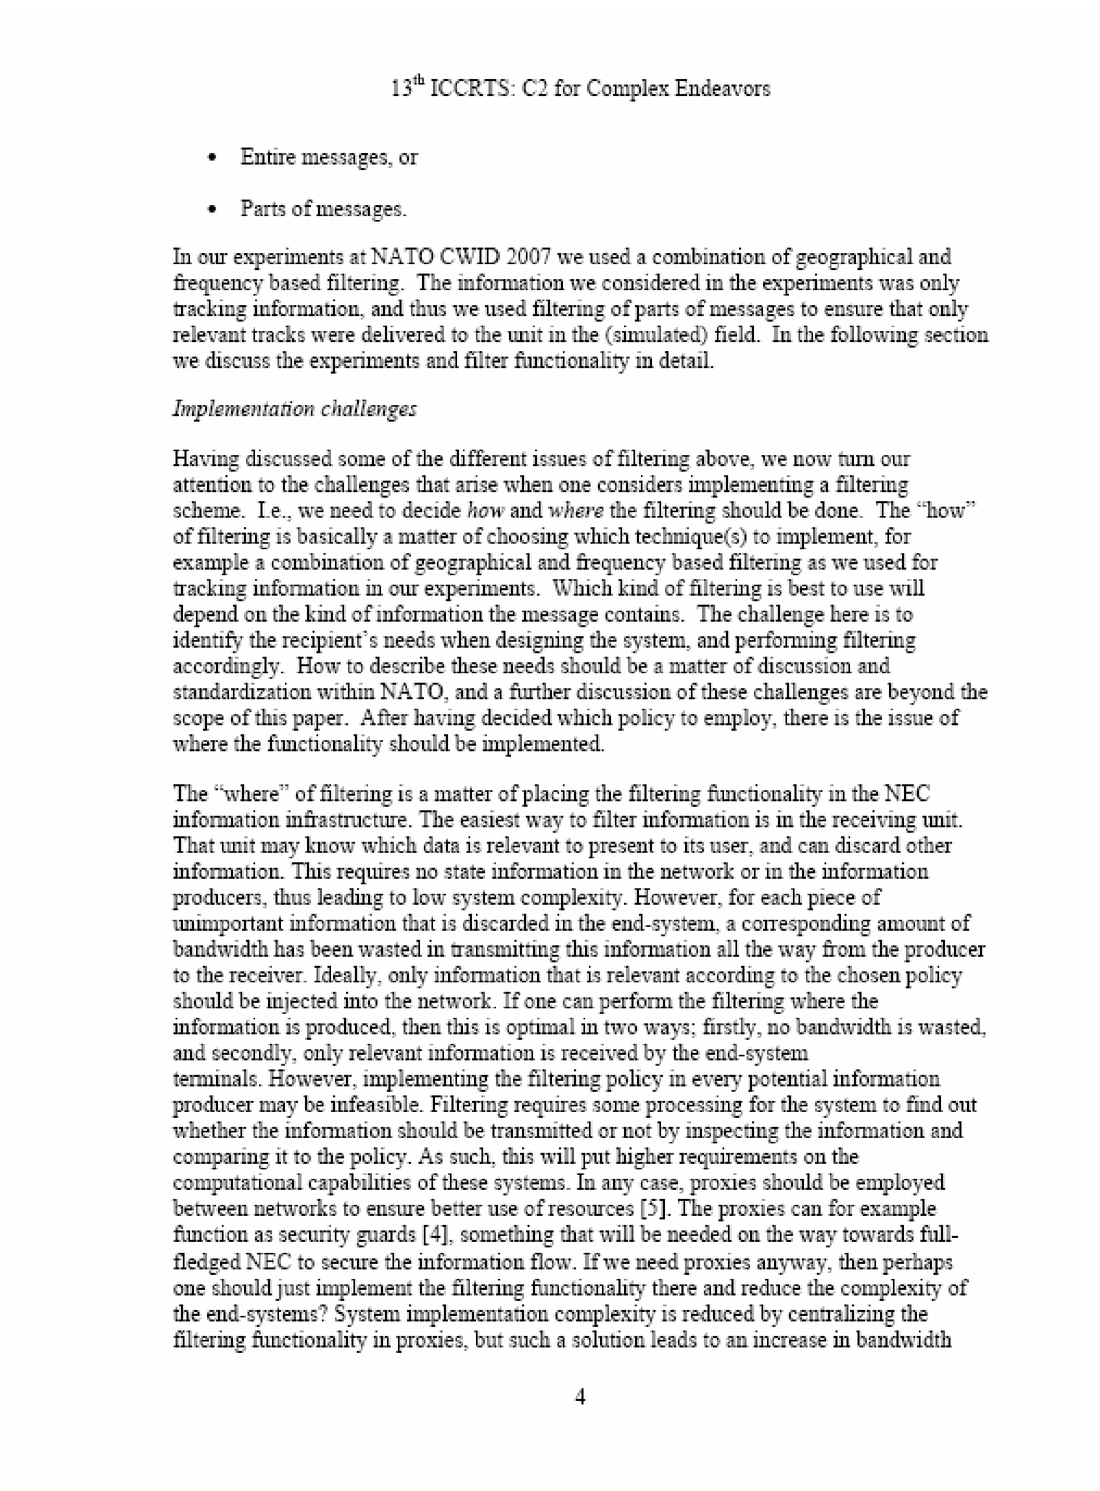

- Entire messages, or
- Parts of messages.

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

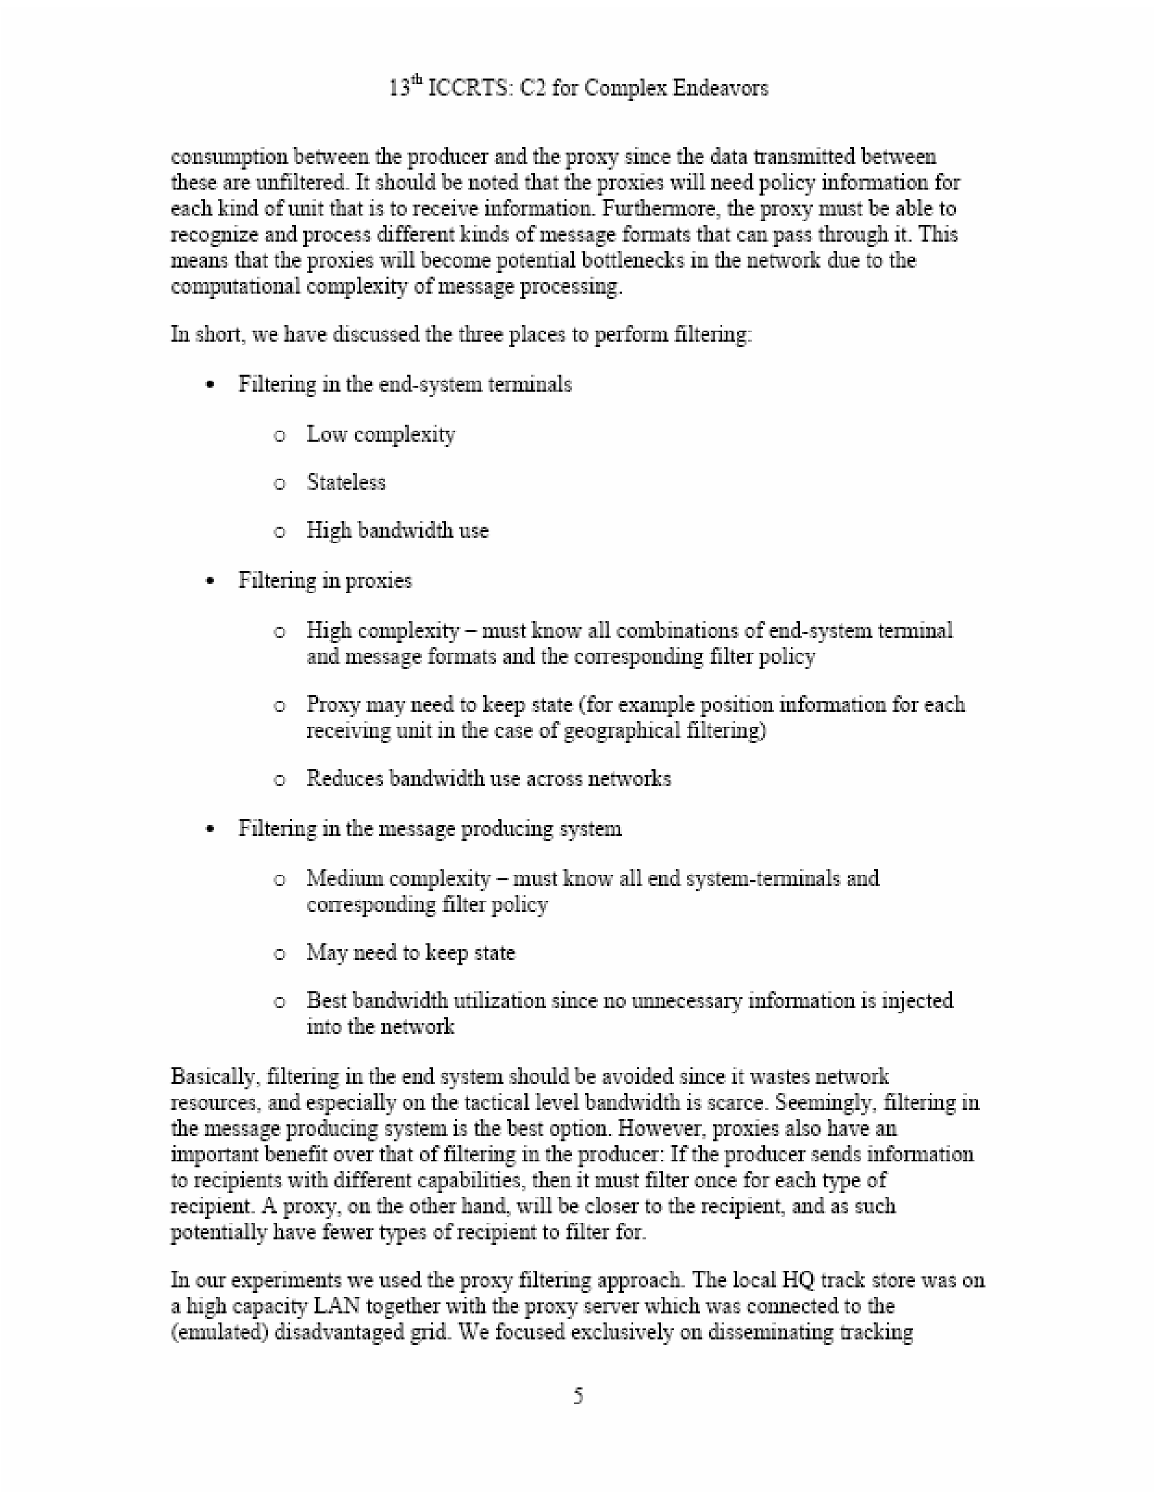

consumption between the producer and the proxy since the data transmitted between these are unfiltered. It should be noted that the proxies will need policy information for each kind of unit that is to receive information. Furthermore, the proxy must be able to recognize and process different kinds of 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

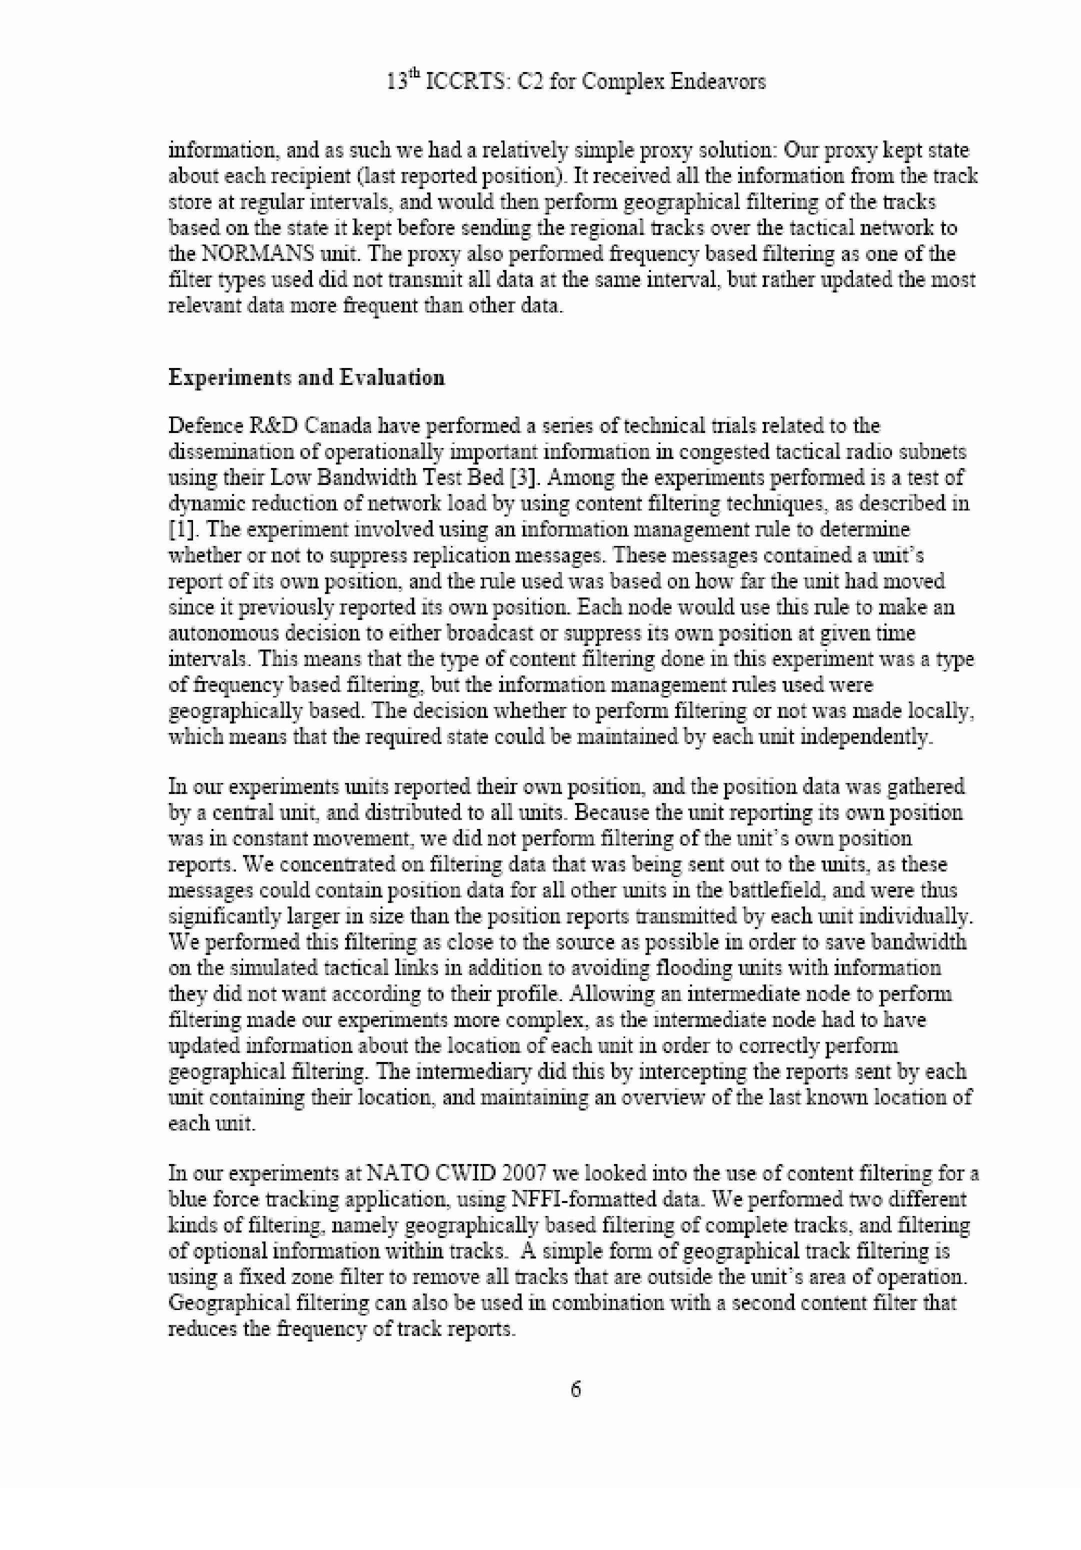

information, and as such we had a relatively simple proxy solution: Our proxy kept state about each recipient (last reported position). It received all the information from the track store at regular intervals, and would then perform geographical filtering of the tracks based on the state it kept before 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.

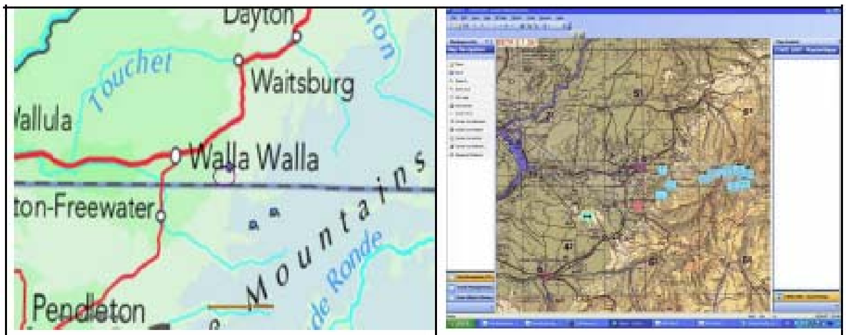

Tactical unit display

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.

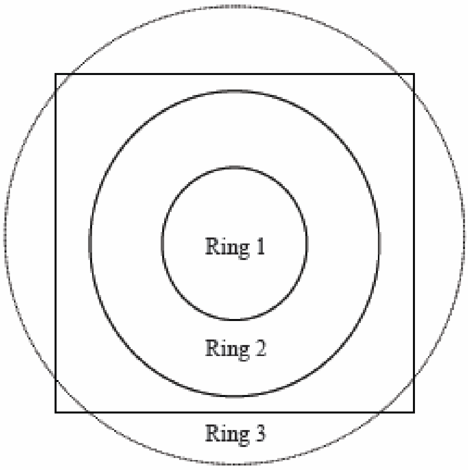

Figure 2 Zone ring filter layout

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.

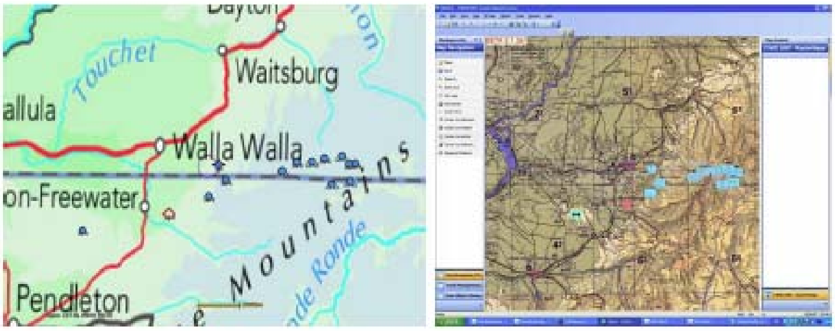

Tactical unit display

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.

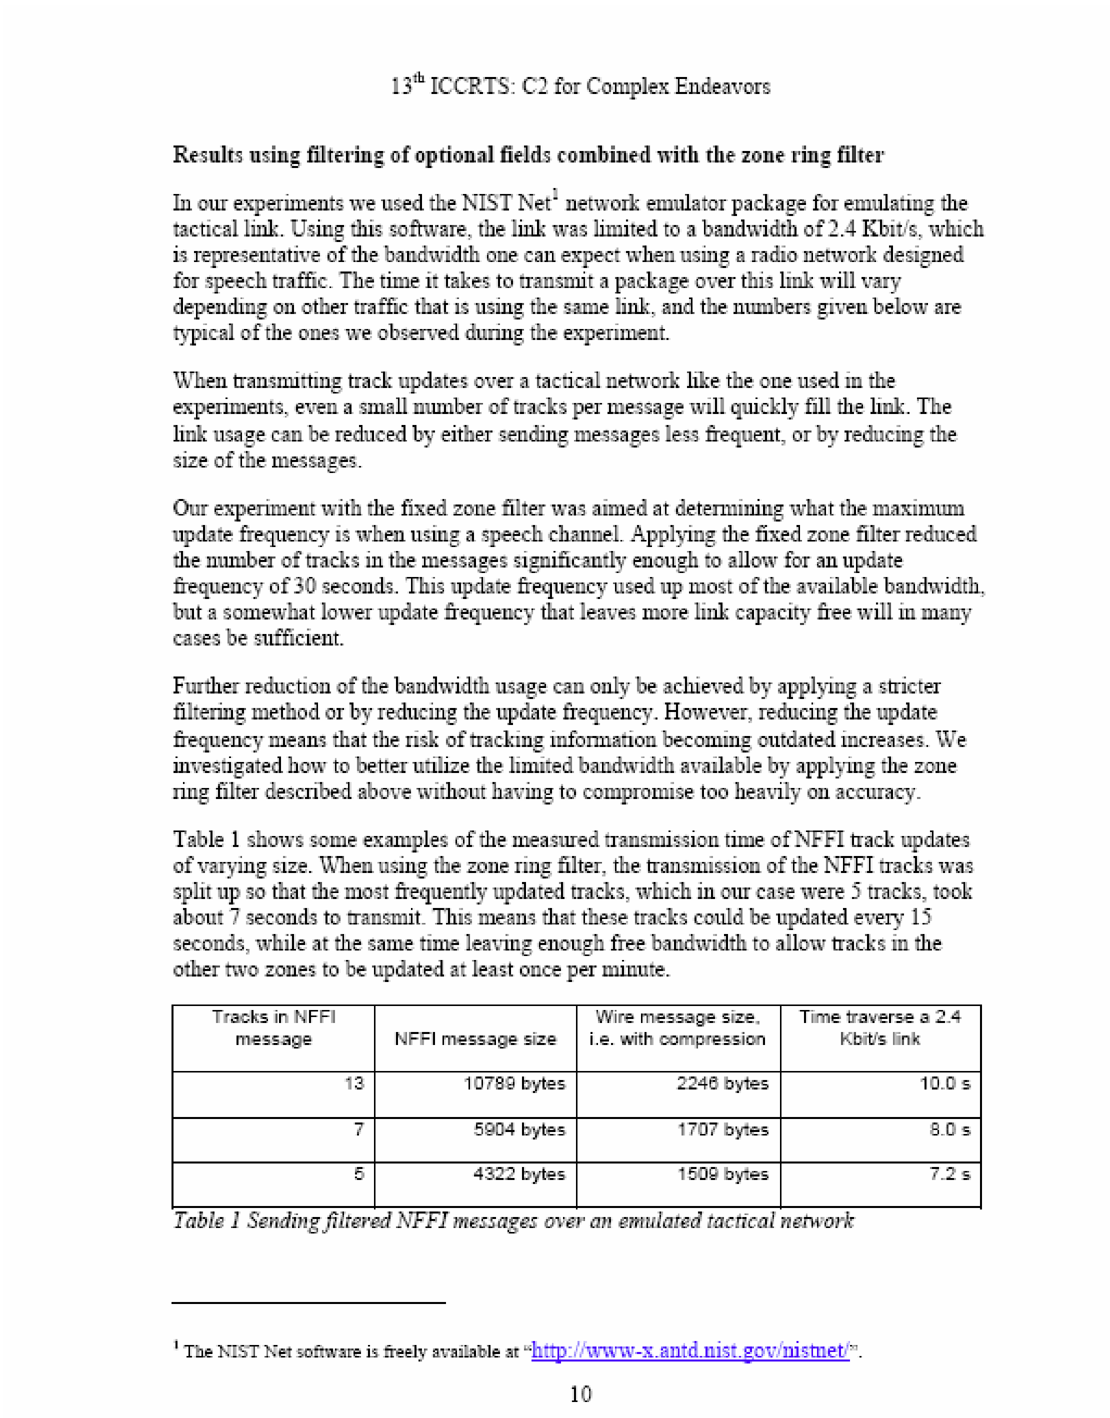

Results using filtering of optional fields combined with the zone ring filter

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/".

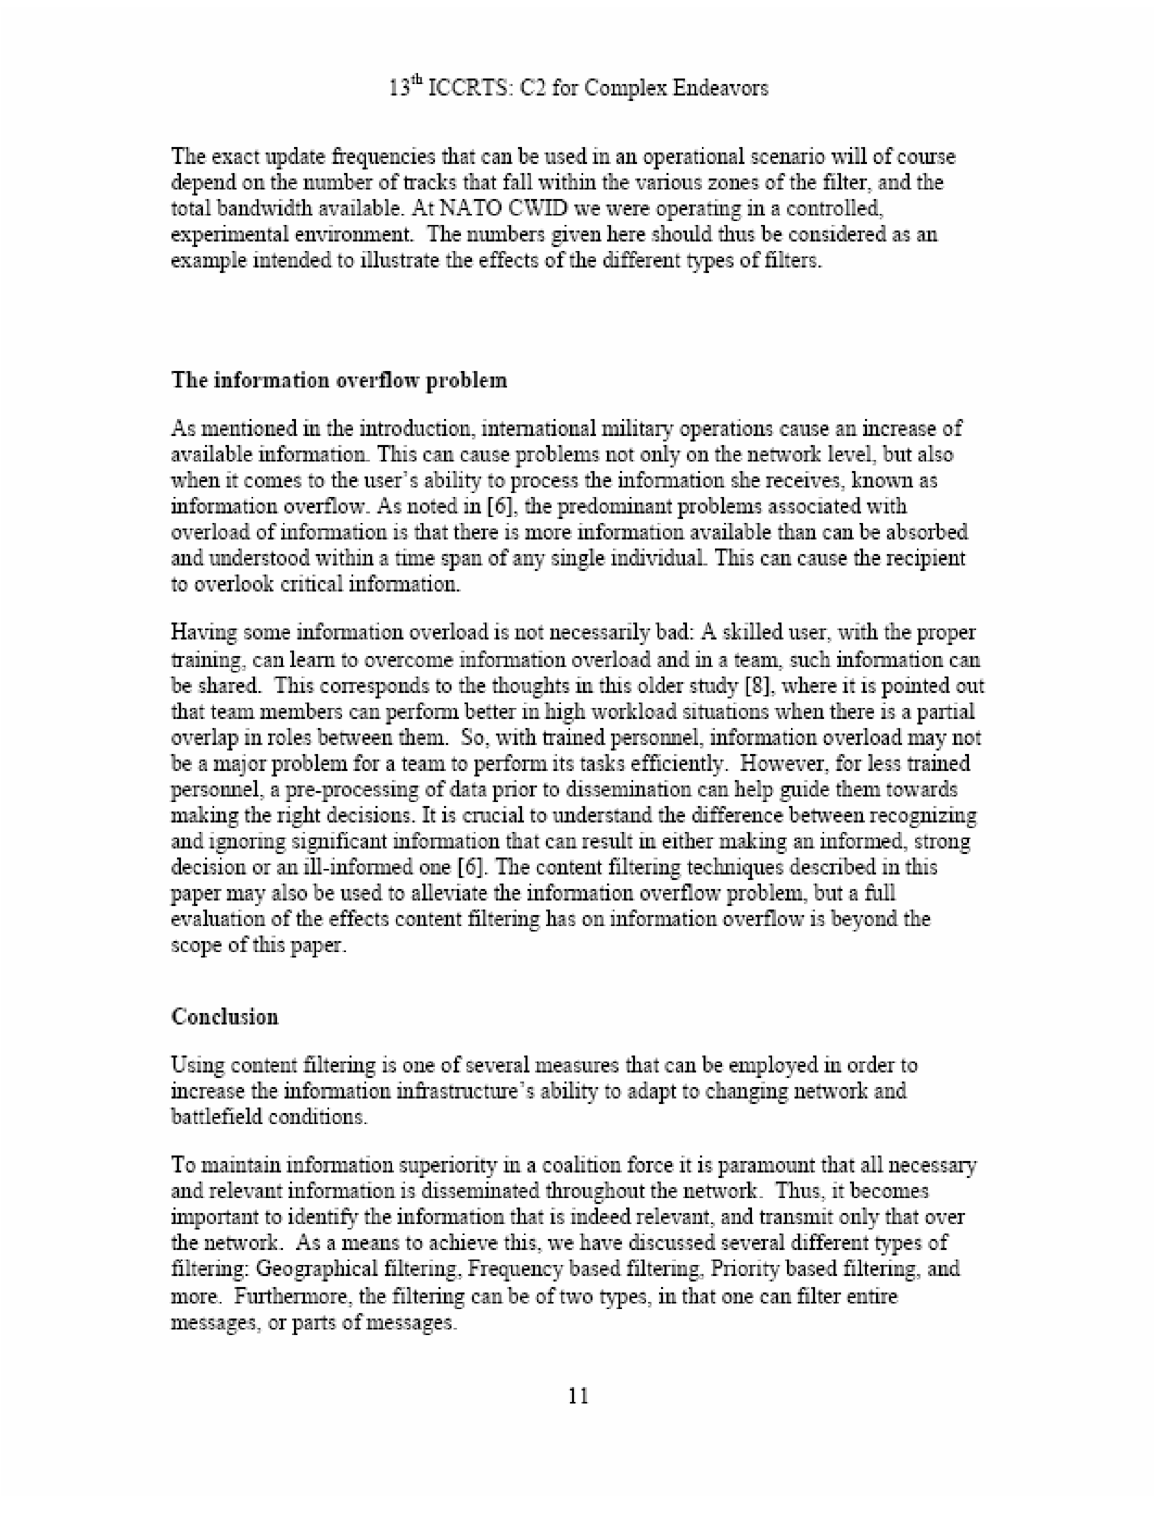

The exact update frequencies that can be used in an operational scenario will of course depend on the number of tracks that fall within the various zones of the filter, and the total bandwidth available. At NATO CWID we were operating in a controlled, experimental environment. The numbers given here 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.

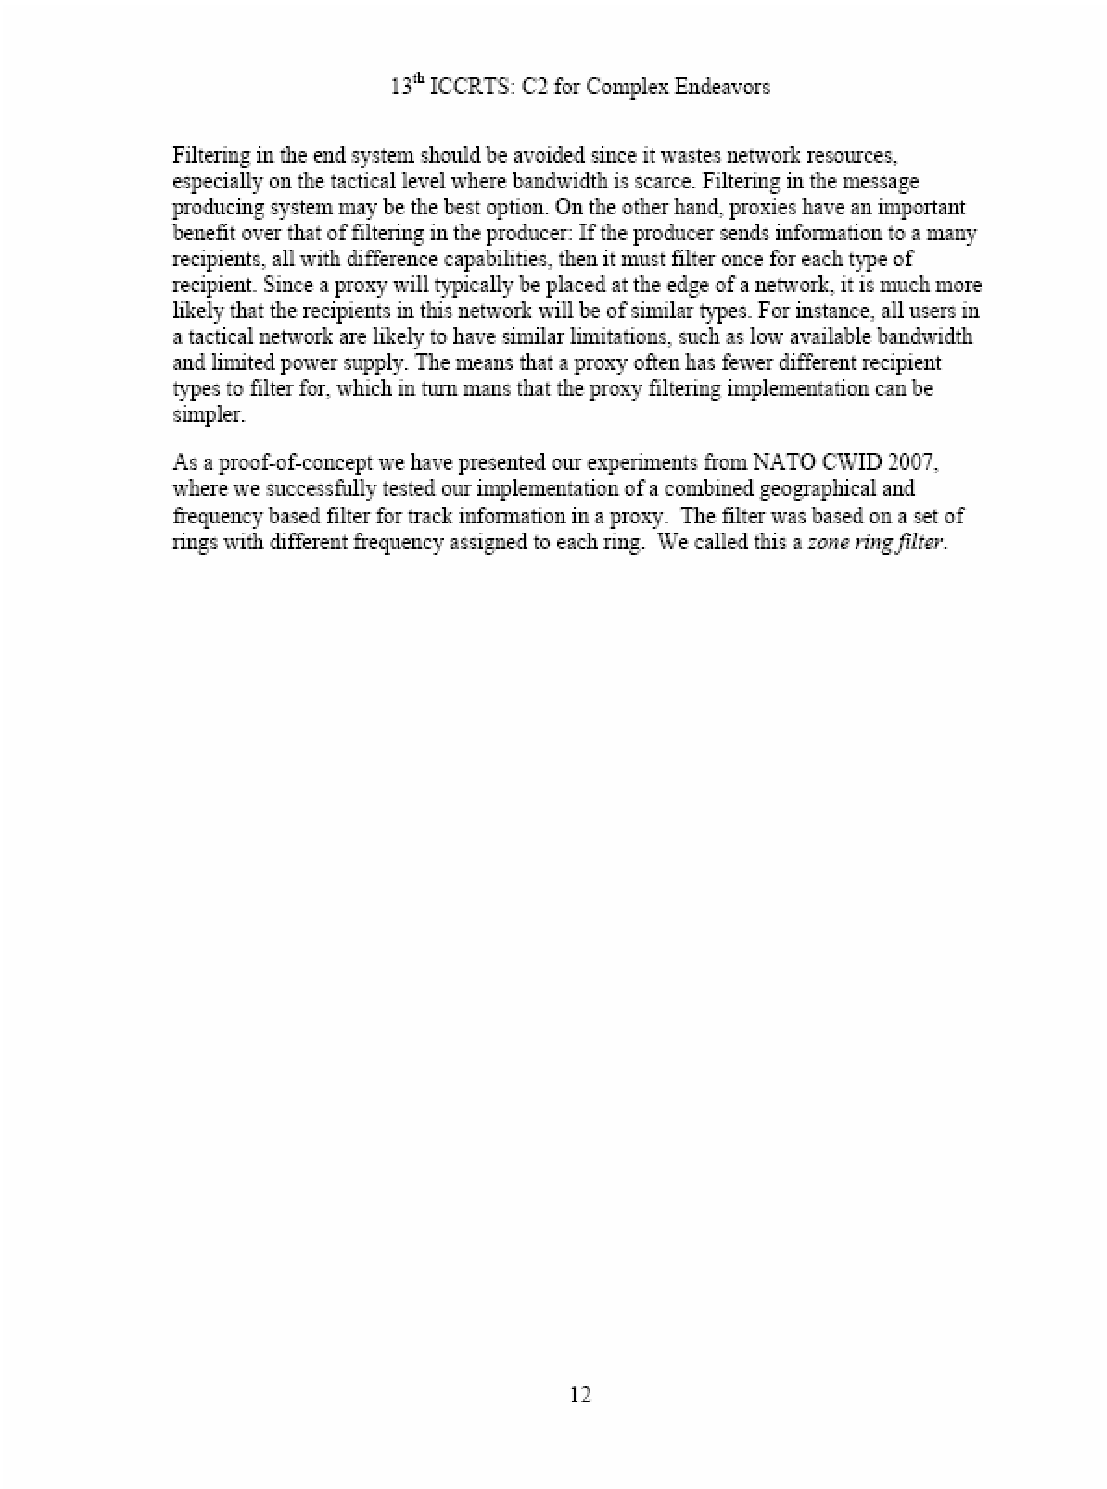

Filtering in the end system should be avoided since it wastes network resources, especially on the tactical level where bandwidth is scarce. Filtering in the message producing system may be the best option. On the other hand, proxies have an important benefit over that of filtering in the producer: If 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*.

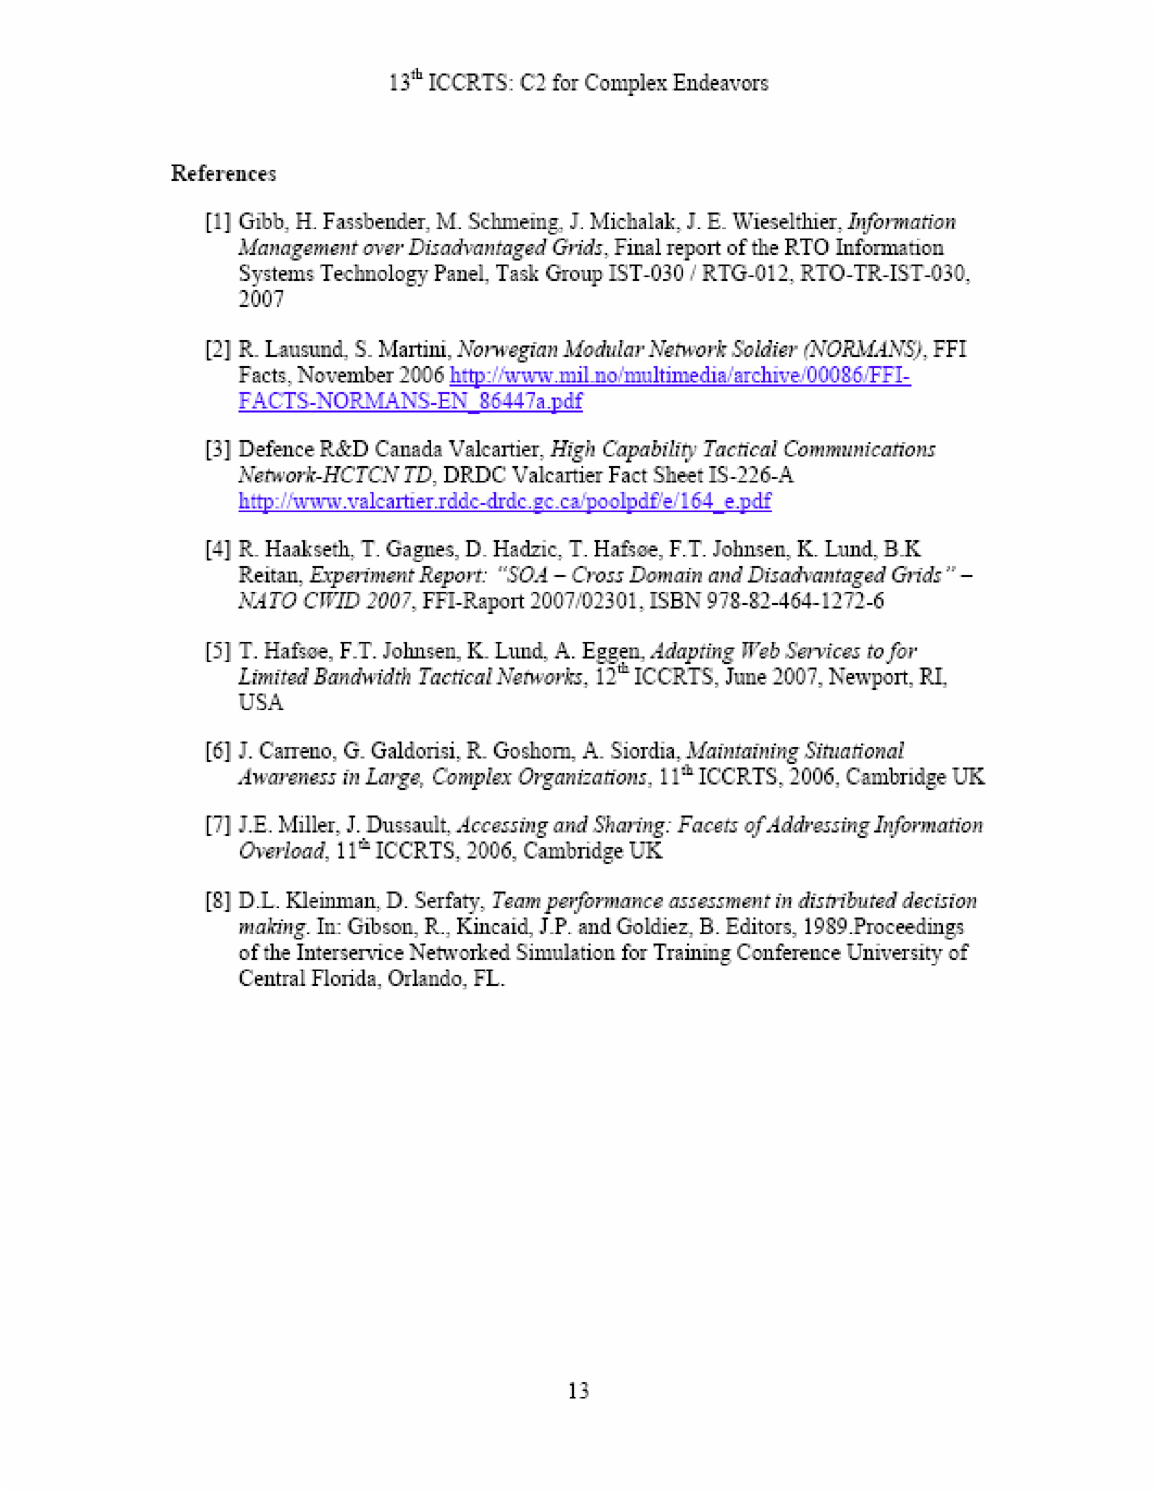

References

[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.

Appendix C Publication at ICCRTS regarding compression as a means to reduce communication overhead

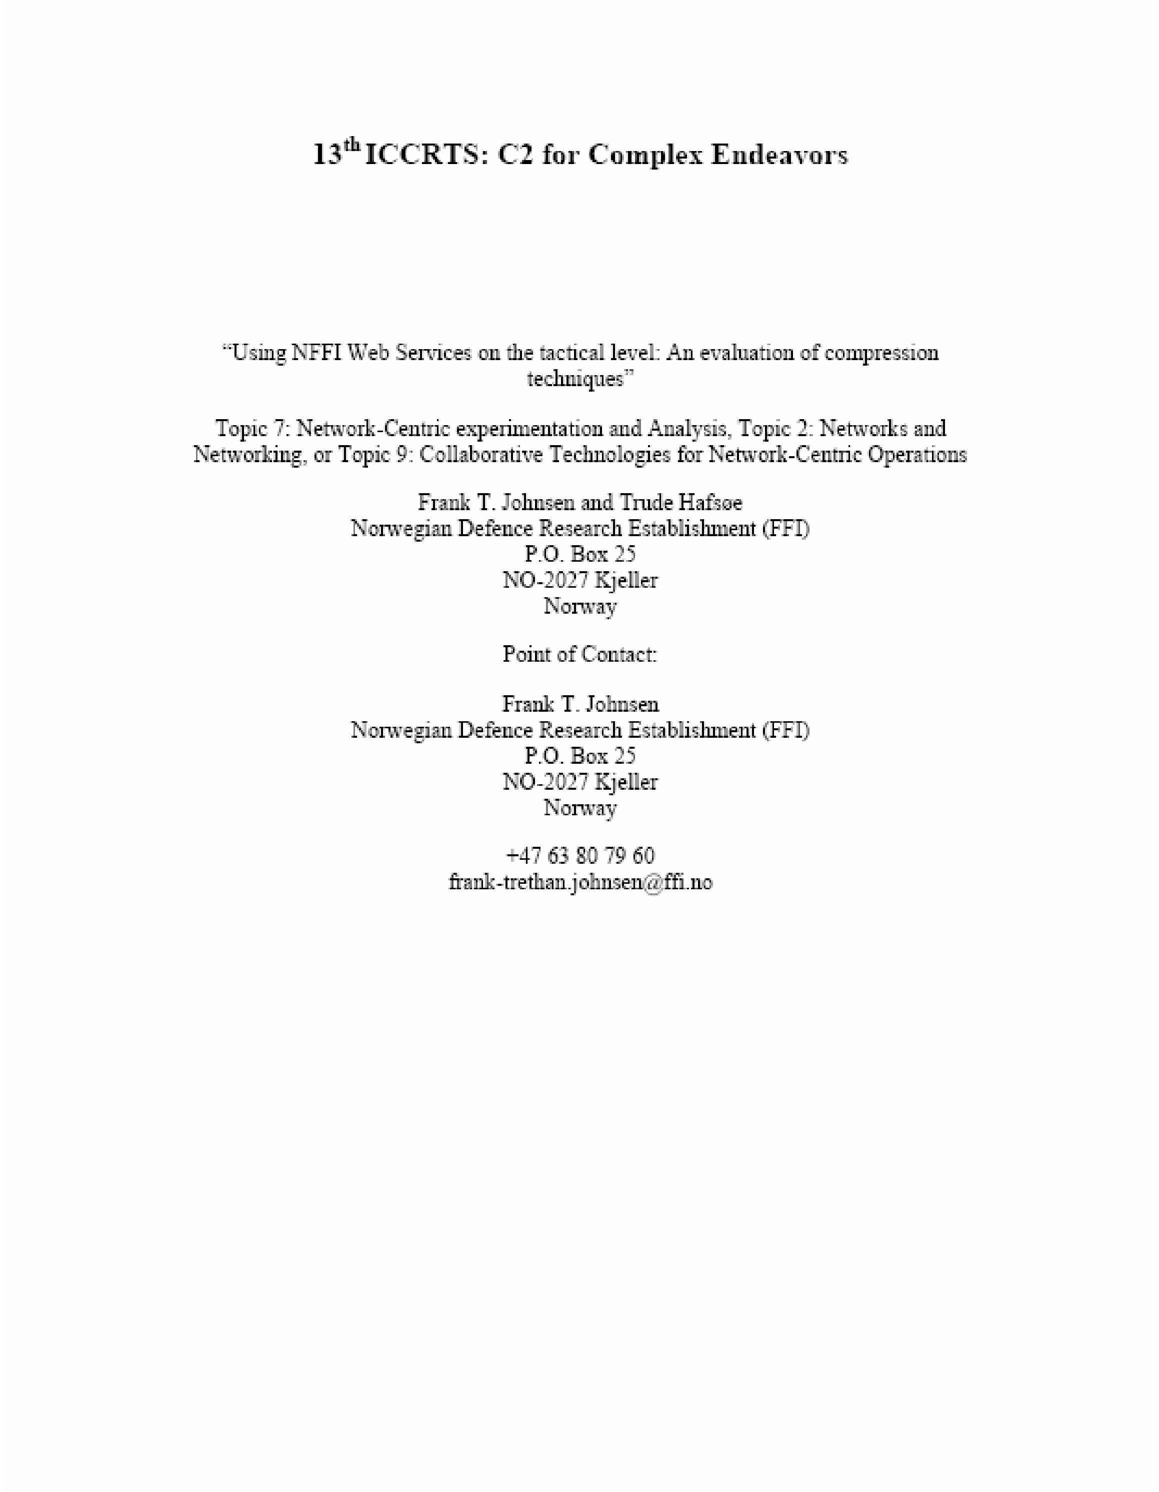

13th ICCRTS: C2 for Complex Endeavors

"Using NFFI Web Services on the tactical level: An evaluation of compression techniques"

Topic 7: Network-Centric experimentation and Analysis, Topic 2: Networks and Networking, or Topic 9: Collaborative Technologies for Network-Centric Operations

Frank T. Johnsen and Trude Hafsøe
Norwegian Defence Research Establishment (FFI)
P.O. Box 25
NO-2027 Kjeller
Norway

Point of Contact:

Frank T. Johnsen
Norwegian Defence Research Establishment (FFI)
P.O. Box 25
NO-2027 Kjeller
Norway

+47 63 80 79 60
frank-trethan.johnsen@ffi.no

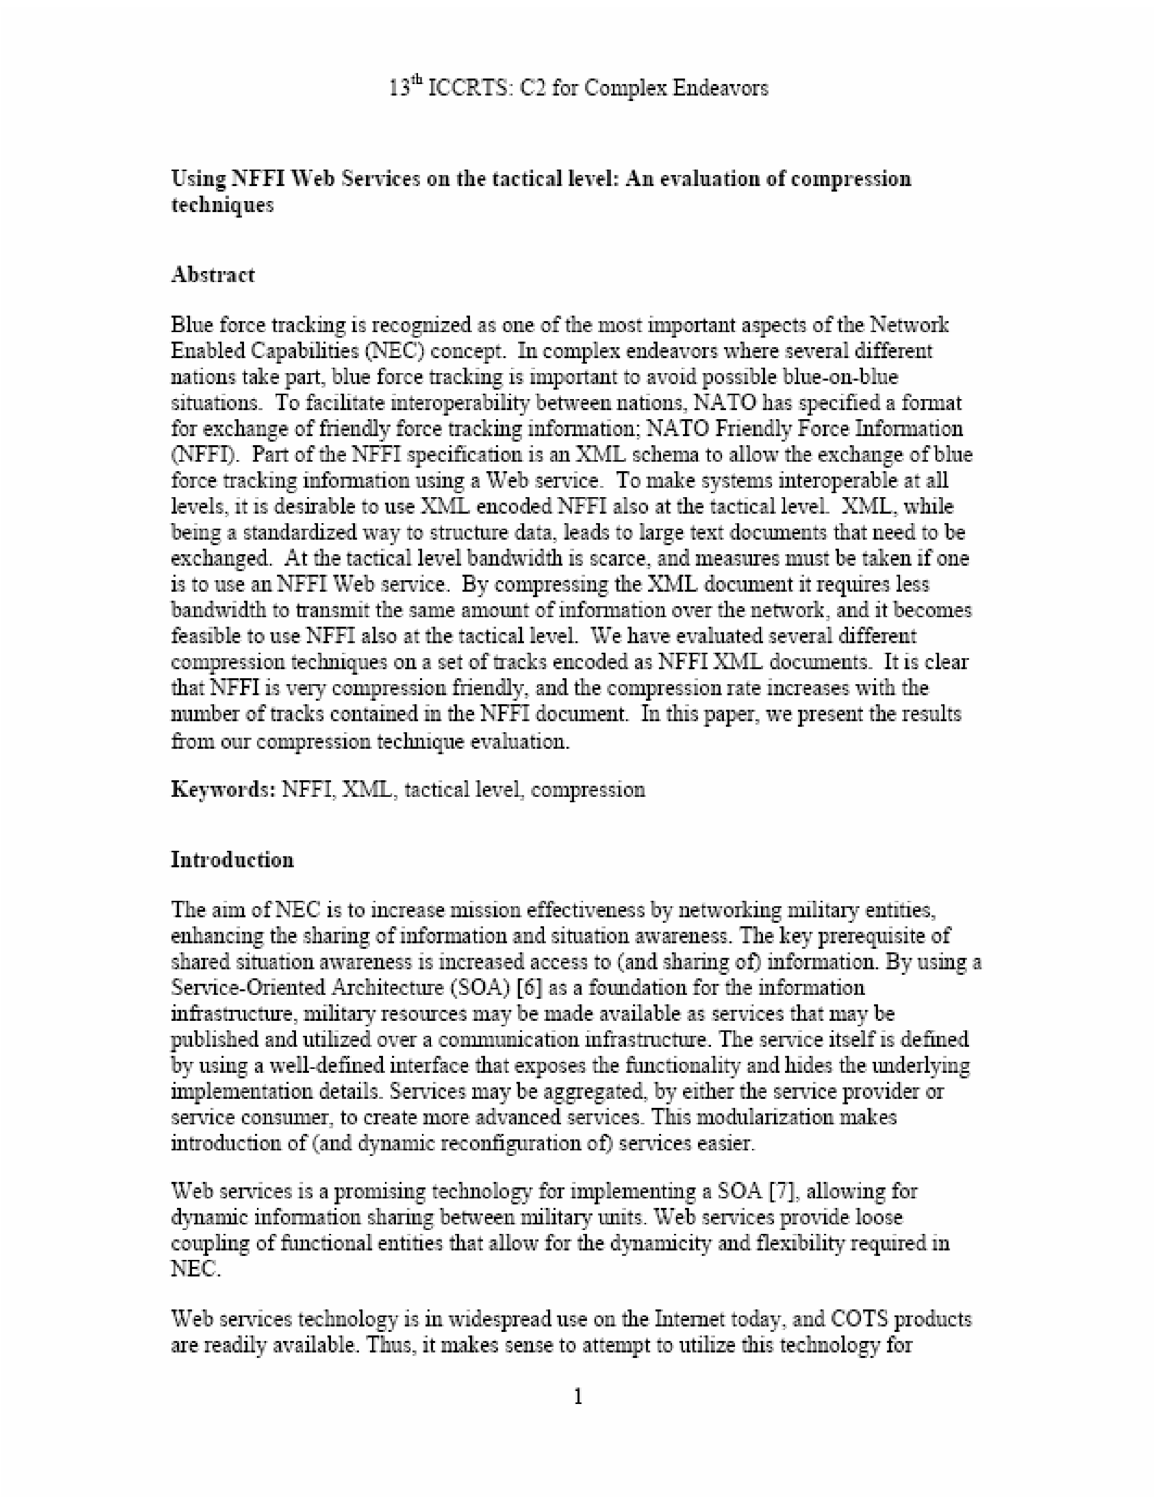

Using NFFI Web Services on the tactical level: An evaluation of compression techniques

Abstract

Blue force tracking is recognized as one of the most important aspects of the Network Enabled Capabilities (NEC) concept. In complex endeavors where several different nations take part, blue force tracking is important to avoid possible blue-on-blue situations. To facilitate interoperability between nations, NATO has specified a format for exchange of friendly force tracking information; NATO Friendly Force Information (NFFI). Part of the NFFI specification is an XML schema to allow the exchange of blue force tracking information using a Web service. To make systems interoperable at all levels, it is desirable to use XML encoded NFFI also at the tactical level. XML, while being a standardized way to structure data, leads to large text documents that need to be exchanged. At the tactical level bandwidth is scarce, and measures must be taken if one is to use an NFFI Web service. By compressing the XML document it requires less bandwidth to transmit the same amount of information over the network, and it becomes feasible to use NFFI also at the tactical level. We have evaluated several different compression techniques on a set of tracks encoded as NFFI XML documents. It is clear that NFFI is very compression friendly, and the compression rate increases with the number of tracks contained in the NFFI document. In this paper, we present the results from our compression technique evaluation.

Keywords: NFFI, XML, tactical level, compression

Introduction

The aim of NEC is to increase mission effectiveness by networking military entities, enhancing the sharing of information and situation awareness. The key prerequisite of shared situation awareness is increased access to (and sharing of) information. By using a Service-Oriented Architecture (SOA) [6] as a foundation for the information infrastructure, military resources may be made available as services that may be published and utilized over a communication infrastructure. The service itself is defined by using a well-defined interface that exposes the functionality and hides the underlying implementation details. Services may be aggregated, by either the service provider or service consumer, to create more advanced services. This modularization makes introduction of (and dynamic reconfiguration of) services easier.

Web services is a promising technology for implementing a SOA [7], allowing for dynamic information sharing between military units. Web services provide loose coupling of functional entities that allow for the dynamicity and flexibility required in NEC.

Web services technology is in widespread use on the Internet today, and COTS products are readily available. Thus, it makes sense to attempt to utilize this technology for

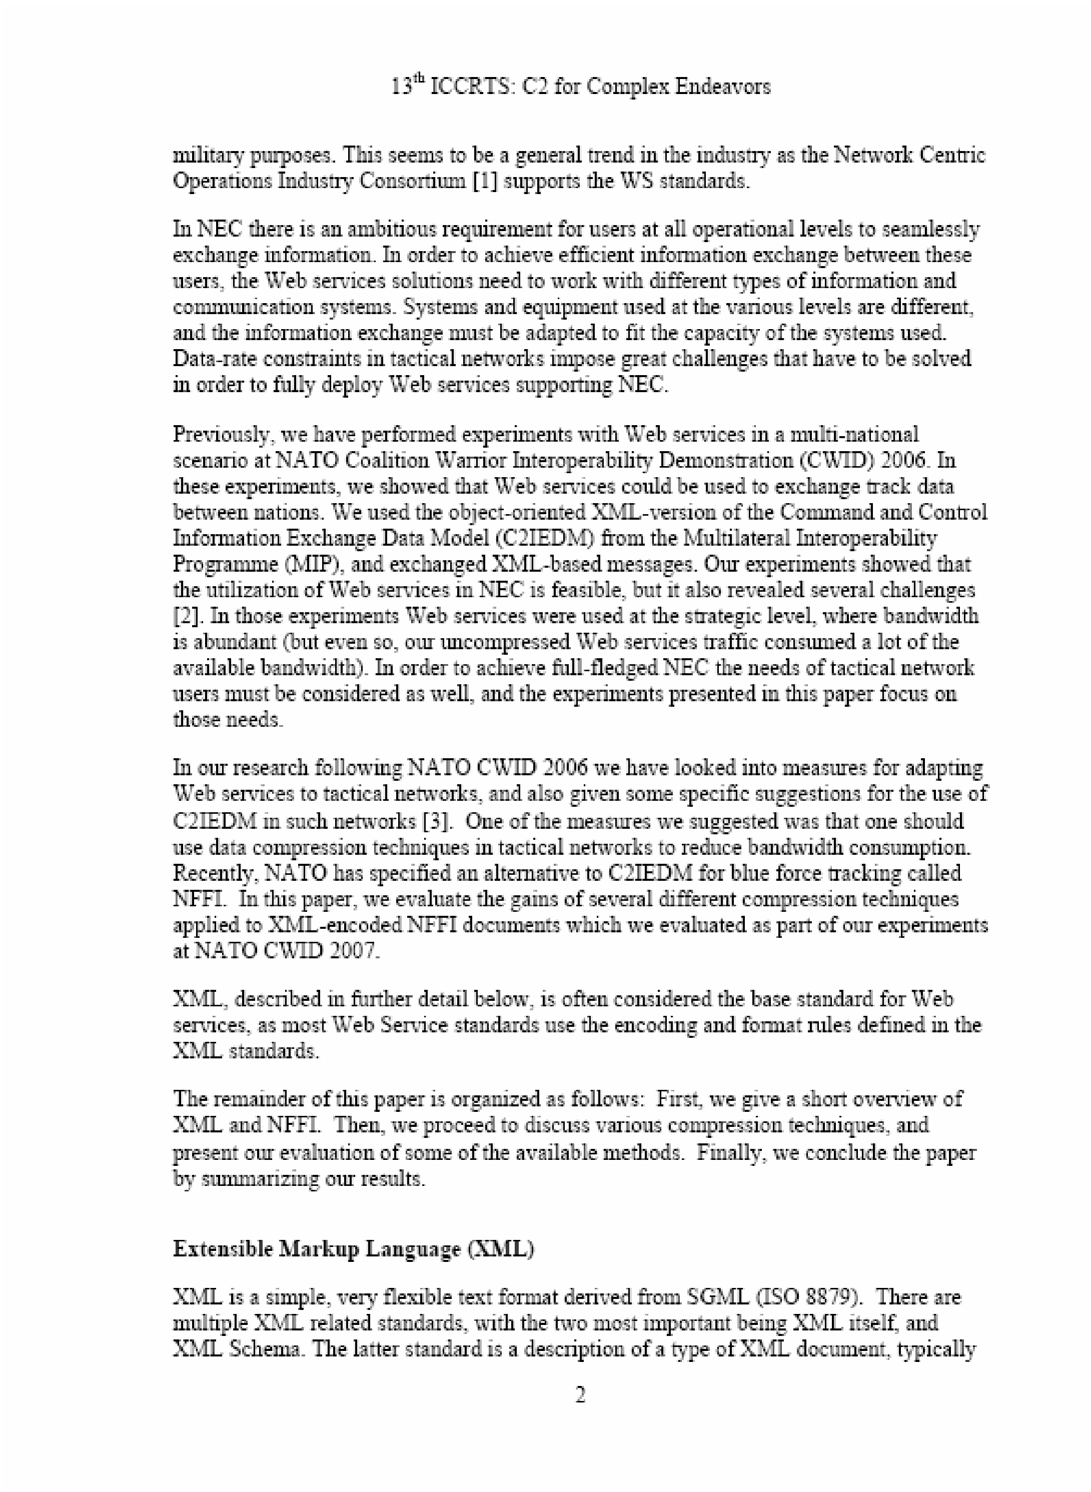

military purposes. This seems to be a general trend in the industry as the Network Centric Operations Industry Consortium [1] supports the WS standards.

In NEC there is an ambitious requirement for users at all operational levels to seamlessly exchange information. In order to achieve efficient information exchange between these users, the Web services solutions need to work with different types of information and communication systems. Systems and equipment used at the various levels are different, and the information exchange must be adapted to fit the capacity of the systems used. Data-rate constraints in tactical networks impose great challenges that have to be solved in order to fully deploy Web services supporting NEC.

Previously, we have performed experiments with Web services in a multi-national scenario at NATO Coalition Warrior Interoperability Demonstration (CWID) 2006. In these experiments, we showed that Web services could be used to exchange track data between nations. We used the object-oriented XML-version of the Command and Control Information Exchange Data Model (C2IEDM) from the Multilateral Interoperability Programme (MIP), and exchanged XML-based messages. Our experiments showed that the utilization of Web services in NEC is feasible, but it also revealed several challenges [2]. In those experiments Web services were used at the strategic level, where bandwidth is abundant (but even so, our uncompressed Web services traffic consumed a lot of the available bandwidth). In order to achieve full-fledged NEC the needs of tactical network users must be considered as well, and the experiments presented in this paper focus on those needs.

In our research following NATO CWID 2006 we have looked into measures for adapting Web services to tactical networks, and also given some specific suggestions for the use of C2IEDM in such networks [3]. One of the measures we suggested was that one should use data compression techniques in tactical networks to reduce bandwidth consumption. Recently, NATO has specified an alternative to C2IEDM for blue force tracking called NFFI. In this paper, we evaluate the gains of several different compression techniques applied to XML-encoded NFFI documents which we evaluated as part of our experiments at NATO CWID 2007.

XML, described in further detail below, is often considered the base standard for Web services, as most Web Service standards use the encoding and format rules defined in the XML standards.

The remainder of this paper is organized as follows: First, we give a short overview of XML and NFFI. Then, we proceed to discuss various compression techniques, and present our evaluation of some of the available methods. Finally, we conclude the paper by summarizing our results.

Extensible Markup Language (XML)

XML is a simple, very flexible text format derived from SGML (ISO 8879). There are multiple XML related standards, with the two most important being XML itself, and XML Schema. The latter standard is a description of a type of XML document, typically

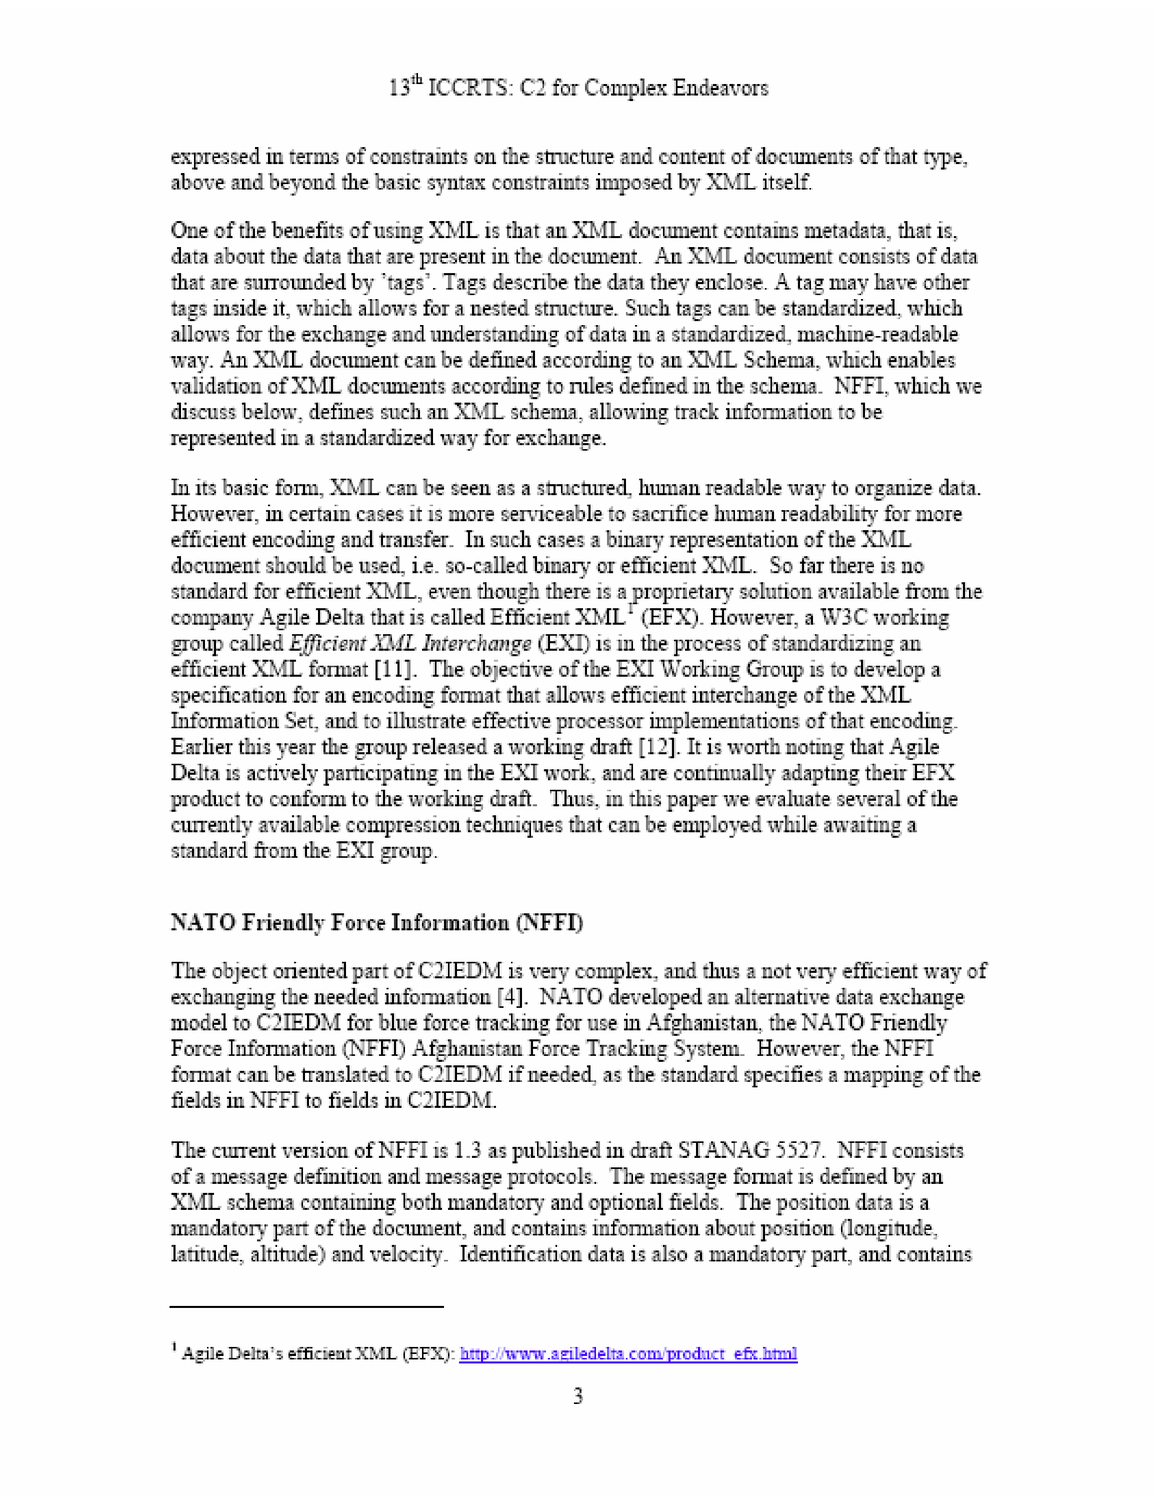

expressed in terms of constraints on the structure and content of documents of that type, above and beyond the basic syntax constraints imposed by XML itself.

One of the benefits of using XML is that an XML document contains metadata, that is, data about the data that are present in the document. An XML document consists of data that are surrounded by 'tags'. Tags describe the data they enclose. A tag may have other tags inside it, which allows for a nested structure. Such tags can be standardized, which allows for the exchange and understanding of data in a standardized, machine-readable way. An XML document can be defined according to an XML Schema, which enables validation of XML documents according to rules defined in the schema. NFFI, which we discuss below, defines such an XML schema, allowing track information to be represented in a standardized way for exchange.

In its basic form, XML can be seen as a structured, human readable way to organize data. However, in certain cases it is more serviceable to sacrifice human readability for more efficient encoding and transfer. In such cases a binary representation of the XML document should be used, i.e. so-called binary or efficient XML. So far there is no standard for efficient XML, even though there is a proprietary solution available from the company Agile Delta that is called Efficient XML[1] (EFX). However, a W3C working group called *Efficient XML Interchange* (EXI) is in the process of standardizing an efficient XML format [11]. The objective of the EXI Working Group is to develop a specification for an encoding format that allows efficient interchange of the XML Information Set, and to illustrate effective processor implementations of that encoding. Earlier this year the group released a working draft [12]. It is worth noting that Agile Delta is actively participating in the EXI work, and are continually adapting their EFX product to conform to the working draft. Thus, in this paper we evaluate several of the currently available compression techniques that can be employed while awaiting a standard from the EXI group.

NATO Friendly Force Information (NFFI)

The object oriented part of C2IEDM is very complex, and thus a not very efficient way of exchanging the needed information [4]. NATO developed an alternative data exchange model to C2IEDM for blue force tracking for use in Afghanistan, the NATO Friendly Force Information (NFFI) Afghanistan Force Tracking System. However, the NFFI format can be translated to C2IEDM if needed, as the standard specifies a mapping of the fields in NFFI to fields in C2IEDM.

The current version of NFFI is 1.3 as published in draft STANAG 5527. NFFI consists of a message definition and message protocols. The message format is defined by an XML schema containing both mandatory and optional fields. The position data is a mandatory part of the document, and contains information about position (longitude, latitude, altitude) and velocity. Identification data is also a mandatory part, and contains

[1] Agile Delta's efficient XML (EFX): http://www.agiledelta.com/product_efx.html

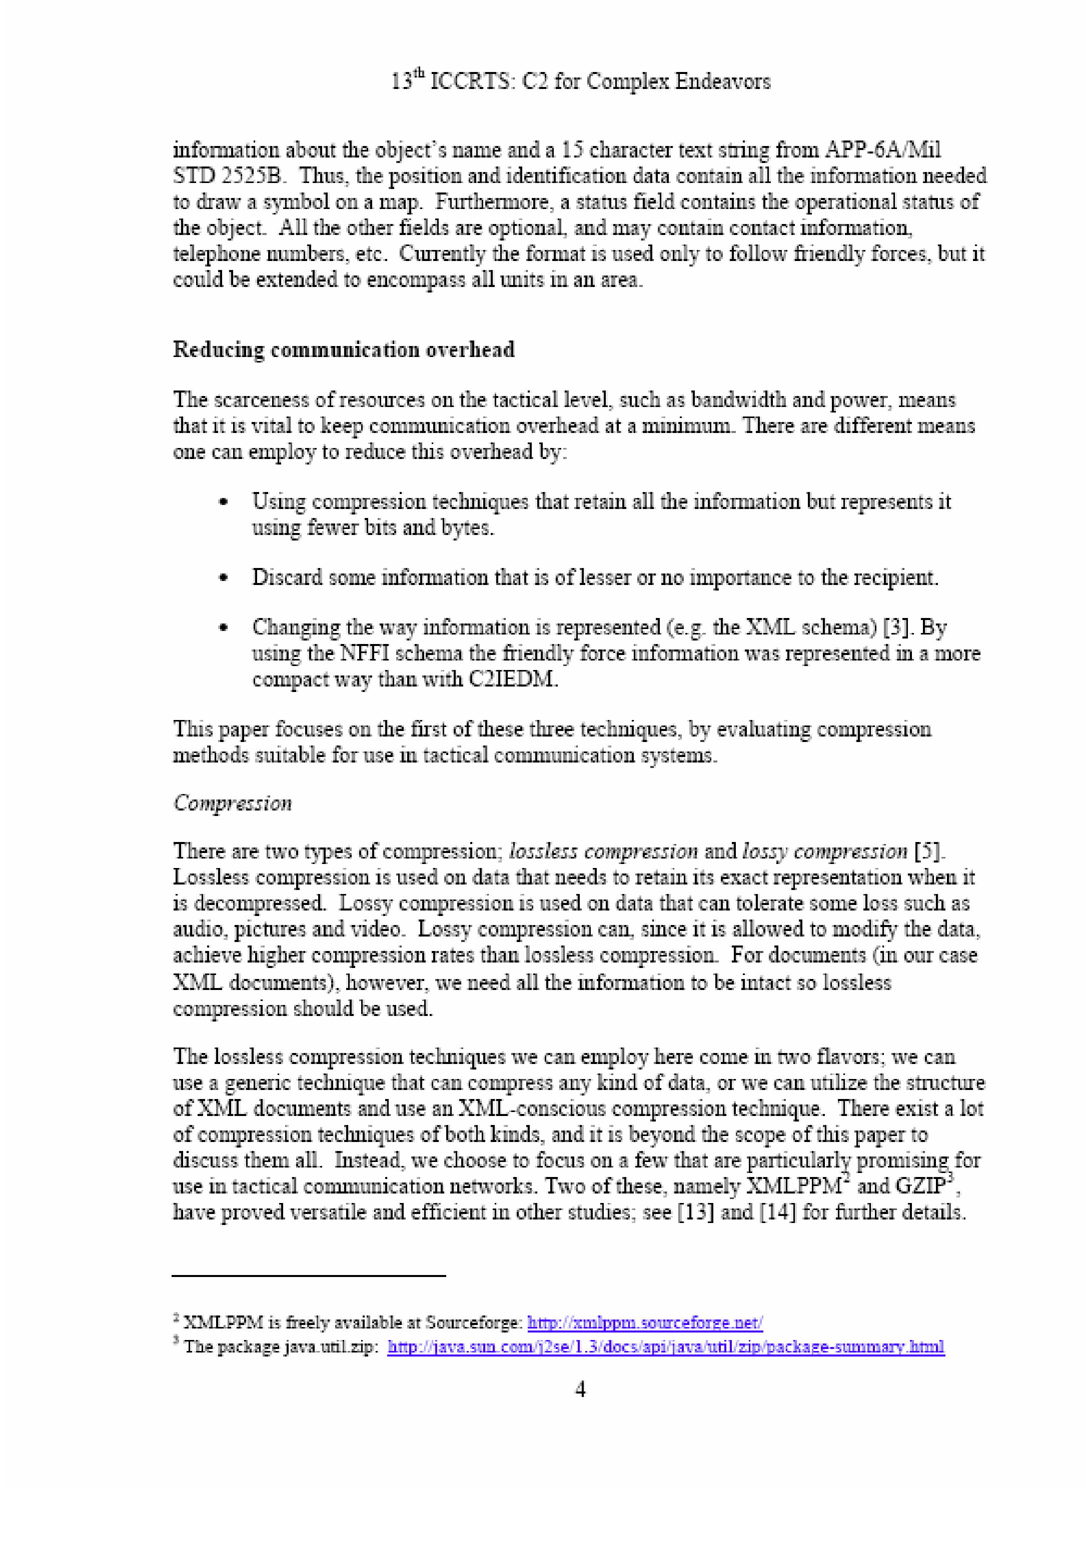

information about the object's name and a 15 character text string from APP-6A/Mil STD 2525B. Thus, the position and identification data contain all the information needed to draw a symbol on a map. Furthermore, a status field contains the operational status of the object. All the other fields are optional, and may contain contact information, telephone numbers, etc. Currently the format is used only to follow friendly forces, but it could be extended to encompass all units in an area.

Reducing communication overhead

The scarceness of resources on the tactical level, such as bandwidth and power, means that it is vital to keep communication overhead at a minimum. There are different means one can employ to reduce this overhead by:

- Using compression techniques that retain all the information but represents it using fewer bits and bytes.
- Discard some information that is of lesser or no importance to the recipient.
- Changing the way information is represented (e.g. the XML schema) [3]. By using the NFFI schema the friendly force information was represented in a more compact way than with C2IEDM.

This paper focuses on the first of these three techniques, by evaluating compression methods suitable for use in tactical communication systems.

Compression

There are two types of compression; *lossless compression* and *lossy compression* [5]. Lossless compression is used on data that needs to retain its exact representation when it is decompressed. Lossy compression is used on data that can tolerate some loss such as audio, pictures and video. Lossy compression can, since it is allowed to modify the data, achieve higher compression rates than lossless compression. For documents (in our case XML documents), however, we need all the information to be intact so lossless compression should be used.

The lossless compression techniques we can employ here come in two flavors; we can use a generic technique that can compress any kind of data, or we can utilize the structure of XML documents and use an XML-conscious compression technique. There exist a lot of compression techniques of both kinds, and it is beyond the scope of this paper to discuss them all. Instead, we choose to focus on a few that are particularly promising for use in tactical communication networks. Two of these, namely XMLPPM[2] and GZIP[3], have proved versatile and efficient in other studies; see [13] and [14] for further details.

[2] XMLPPM is freely available at Sourceforge: http://xmlppm.sourceforge.net/

[3] The package java.util.zip: http://java.sun.com/j2se/1.3/docs/api/java/util/zip/package-summary.html

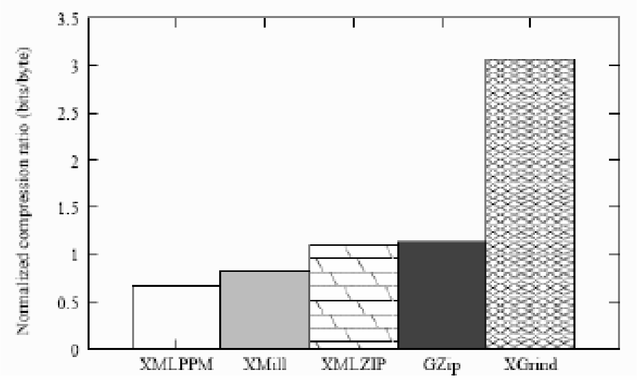

Figure 1 Average Compression Ratio(fetched from [13])

Figure 1 shows how a number of compression techniques performed in one of these studies. XMLPPM gave the best average compression ratios of the XML-conscious compression techniques, while GZIP was the best of the generic compression algorithms tested. Because these two techniques have been shown to be the best of their respective types, we chose to test these further in our evaluation.

Interoperability is a key challenge in NEC, so a standard based compression method is preferable. Since, as mentioned above, EFX is continually adapted to conform to the working drafts released by the EXI, we found it important to investigate EFX in the context of NFFI compression. The studies in [13] and [14] did not investigate EFX, since it is a rather new technique, but [14] mentioned its existence and that it should be evaluated in future work.

EFX can be used in one of two modes of operation; *generic* and *schema specific* compression. The generic option can compress any valid XML document without knowledge of the schema. The schema specific option needs to have access to the XML schema when it performs compression and decompression, thereby sacrificing generality for a very slight increase in compression rate. We used the generic option in our experiments enabling us to compare EFX directly to XMLPPM (which provides only non-schema specific XML compression). When evaluating the efficiency of the algorithms we focused on compression results and not resource usage during compression (memory and CPU usage). The reason for this is that for our intended use, i.e. in tactical networks, the bandwidth is the most limiting resource. Power consumption is also an issue when using battery powered communication equipment. However, in an earlier study [4], we have shown that the difference in computation time between the various techniques is in the millisecond range. Transmission of data also requires power, and by using a compression technique with a high compression ratio, we can reduce the transmission time by several seconds. It is reasonable to assume that the reduction in power consumption caused by reduced transmission time greatly outweighs the benefits of saving a few milliseconds when performing compression and decompression.

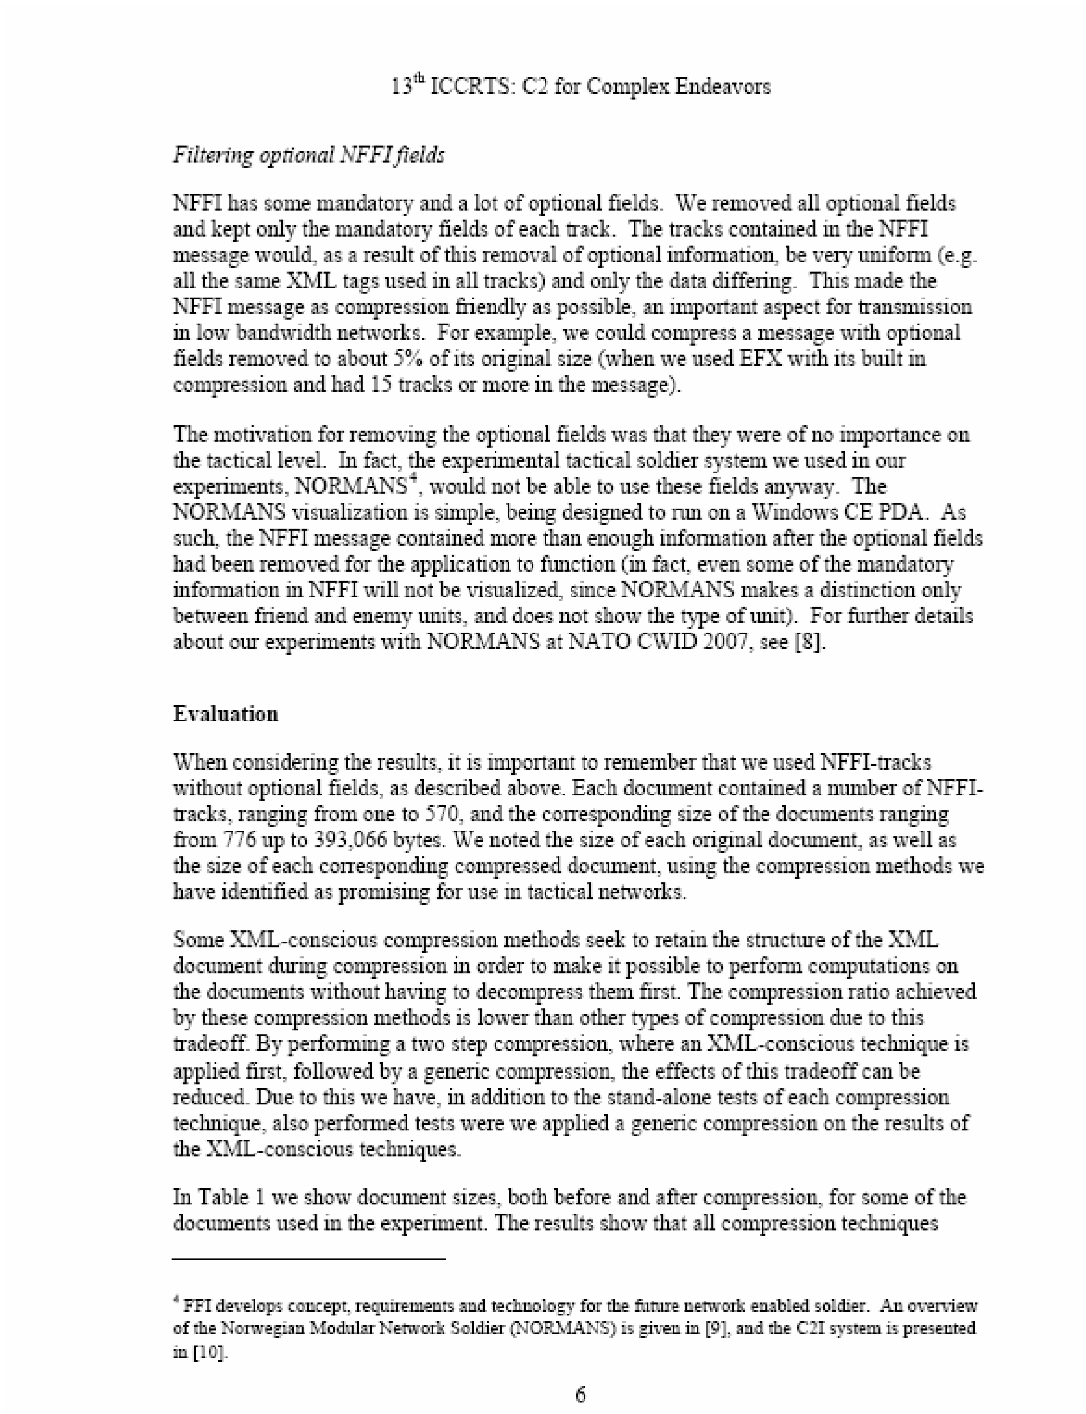

Filtering optional NFFI fields

NFFI has some mandatory and a lot of optional fields. We removed all optional fields and kept only the mandatory fields of each track. The tracks contained in the NFFI message would, as a result of this removal of optional information, be very uniform (e.g. all the same XML tags used in all tracks) and only the data differing. This made the NFFI message as compression friendly as possible, an important aspect for transmission in low bandwidth networks. For example, we could compress a message with optional fields removed to about 5% of its original size (when we used EFX with its built in compression and had 15 tracks or more in the message).

The motivation for removing the optional fields was that they were of no importance on the tactical level. In fact, the experimental tactical soldier system we used in our experiments, NORMANS[4], would not be able to use these fields anyway. The NORMANS visualization is simple, being designed to run on a Windows CE PDA. As such, the NFFI message contained more than enough information after the optional fields had been removed for the application to function (in fact, even some of the mandatory information in NFFI will not be visualized, since NORMANS makes a distinction only between friend and enemy units, and does not show the type of unit). For further details about our experiments with NORMANS at NATO CWID 2007, see [8].

Evaluation

When considering the results, it is important to remember that we used NFFI-tracks without optional fields, as described above. Each document contained a number of NFFI-tracks, ranging from one to 570, and the corresponding size of the documents ranging from 776 up to 393,066 bytes. We noted the size of each original document, as well as the size of each corresponding compressed document, using the compression methods we have identified as promising for use in tactical networks.

Some XML-conscious compression methods seek to retain the structure of the XML document during compression in order to make it possible to perform computations on the documents without having to decompress them first. The compression ratio achieved by these compression methods is lower than other types of compression due to this tradeoff. By performing a two step compression, where an XML-conscious technique is applied first, followed by a generic compression, the effects of this tradeoff can be reduced. Due to this we have, in addition to the stand-alone tests of each compression technique, also performed tests were we applied a generic compression on the results of the XML-conscious techniques.

In Table 1 we show document sizes, both before and after compression, for some of the documents used in the experiment. The results show that all compression techniques

[4] FFI develops concept, requirements and technology for the future network enabled soldier. An overview of the Norwegian Modular Network Soldier (NORMANS) is given in [9], and the C2I system is presented in [10].

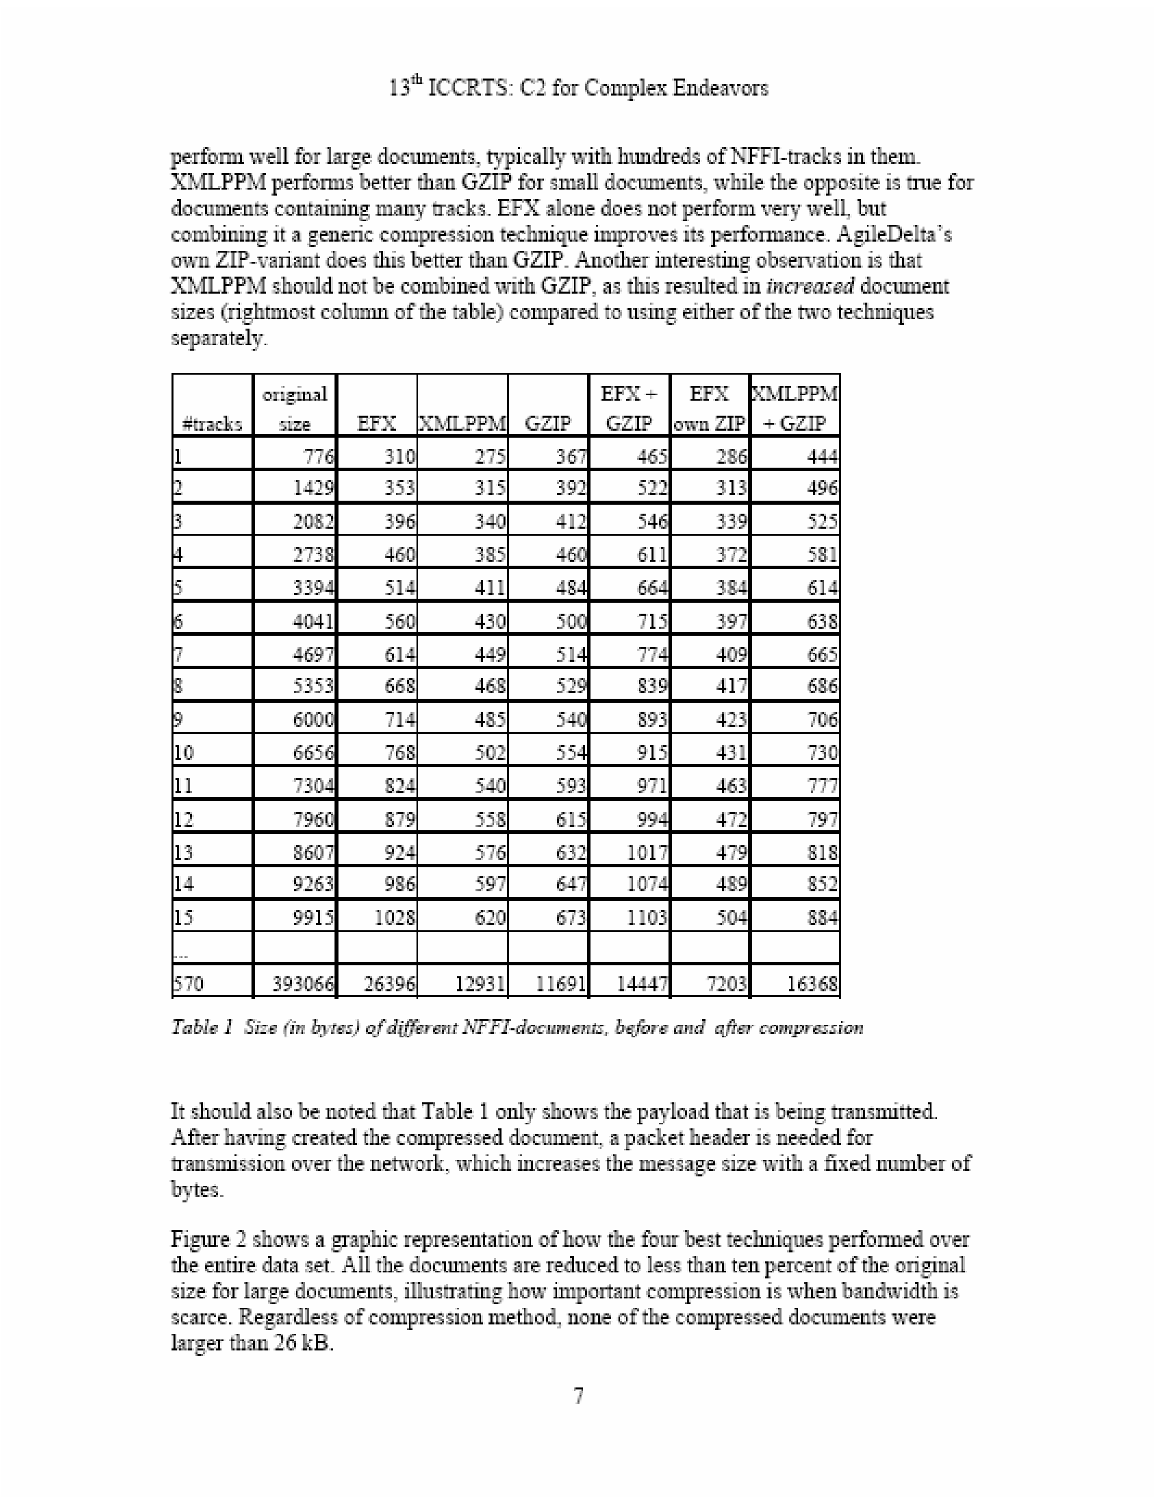

perform well for large documents, typically with hundreds of NFFI-tracks in them. XMLPPM performs better than GZIP for small documents, while the opposite is true for documents containing many tracks. EFX alone does not perform very well, but combining it a generic compression technique improves its performance. AgileDelta's own ZIP-variant does this better than GZIP. Another interesting observation is that XMLPPM should not be combined with GZIP, as this resulted in *increased* document sizes (rightmost column of the table) compared to using either of the two techniques separately.

#tracks	original size	EFX	XMLPPM	GZIP	EFX + GZIP	EFX own ZIP	XMLPPM + GZIP
1	776	310	275	367	465	286	444
2	1429	353	315	392	522	313	496
3	2082	396	340	412	546	339	525
4	2738	460	385	460	611	372	581
5	3394	514	411	484	664	384	614
6	4041	560	430	500	715	397	638
7	4697	614	449	514	774	409	665
8	5353	668	468	529	839	417	686
9	6000	714	485	540	893	423	706
10	6656	768	502	554	915	431	730
11	7304	824	540	593	971	463	777
12	7960	879	558	615	994	472	797
13	8607	924	576	632	1017	479	818
14	9263	986	597	647	1074	489	852
15	9915	1028	620	673	1103	504	884
...							
570	393066	26396	12931	11691	14447	7203	16368

Table 1 Size (in bytes) of different NFFI-documents, before and after compression

It should also be noted that Table 1 only shows the payload that is being transmitted. After having created the compressed document, a packet header is needed for transmission over the network, which increases the message size with a fixed number of bytes.

Figure 2 shows a graphic representation of how the four best techniques performed over the entire data set. All the documents are reduced to less than ten percent of the original size for large documents, illustrating how important compression is when bandwidth is scarce. Regardless of compression method, none of the compressed documents were larger than 26 kB.

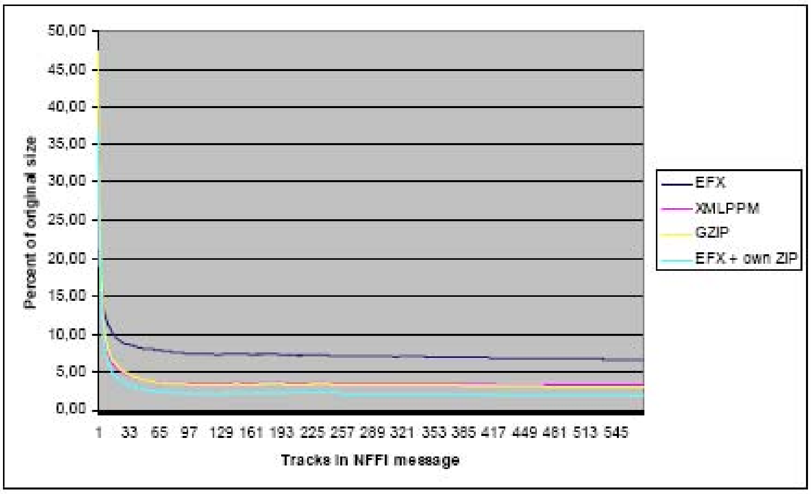

Figure 2 Percent of original size of the XML document after compression compared with the number of tracks in the NFFI message

Figure 2 also shows that when the number of tracks in the documents reaches a certain level, EFX combined with its own built-in ZIP always performs best. However, when there are fewer tracks to report, the differences between the compression techniques are larger. In a disadvantaged grid, every byte saved can have a noticeable impact on the transmission time, and it is therefore important to take the individual differences between the techniques into account. Figure 3 shows the results for the 15 smallest documents. For the smallest documents with only one or two tracks, XMLPPM does the best job. EFX alone does not perform very well, but combining it with Agile Delta's own ZIP-variant makes it out-perform all the other techniques as soon as the documents start growing in size.

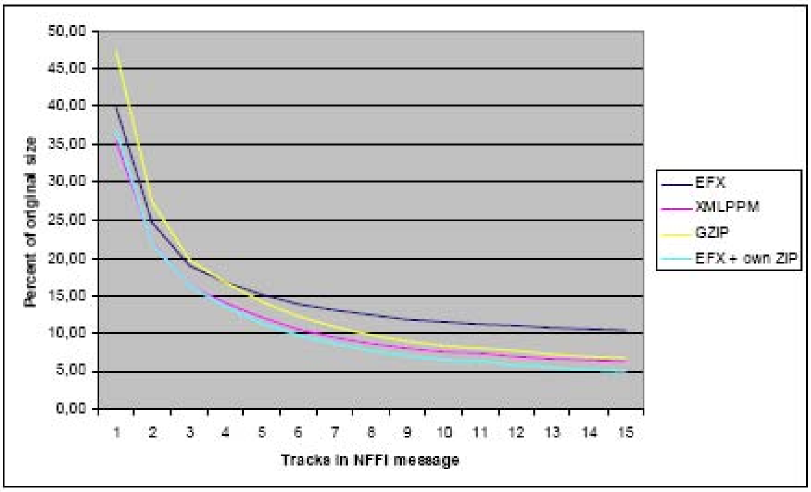

Figure 3 Percent of original size of the XML document after compression compared with the number of tracks in the NFFI message

For all four compression methods, the compression ratio continues to increase with increasing number of NFFI-tracks in the documents. The best ratio we measured was using EFX with Agile Delta's ZIP on the largest document with 570 NFFI-tracks, where we achieved a compressed document size of less than 2% of the original size, as is shown in Table 2.

#tracks	EFX	XMLPPM	GZIP	EFX + GZIP	EFX own ZIP	XMLPPM + GZIP
1	39.95	35.44	47.29	59.92	36.86	57.22
10	11.54	7.54	8.32	13.75	6.48	10.97
50	8.11	3.98	4.01	6.30	2.69	5.31
100	7.42	3.52	3.34	4.64	2.19	4.58
200	7.24	3.56	3.38	4.73	2.19	4.55
300	7.10	3.52	3.36	4.46	2.11	4.48
400	6.92	3.43	3.22	4.07	2.00	4.36
500	6.74	3.34	3.04	3.84	1.89	4.24
570	6.72	3.29	2.97	3.68	1.83	4.16

Table 2 Size (in percent of the original, uncompressed XML document) of different NFFI-documents after compression

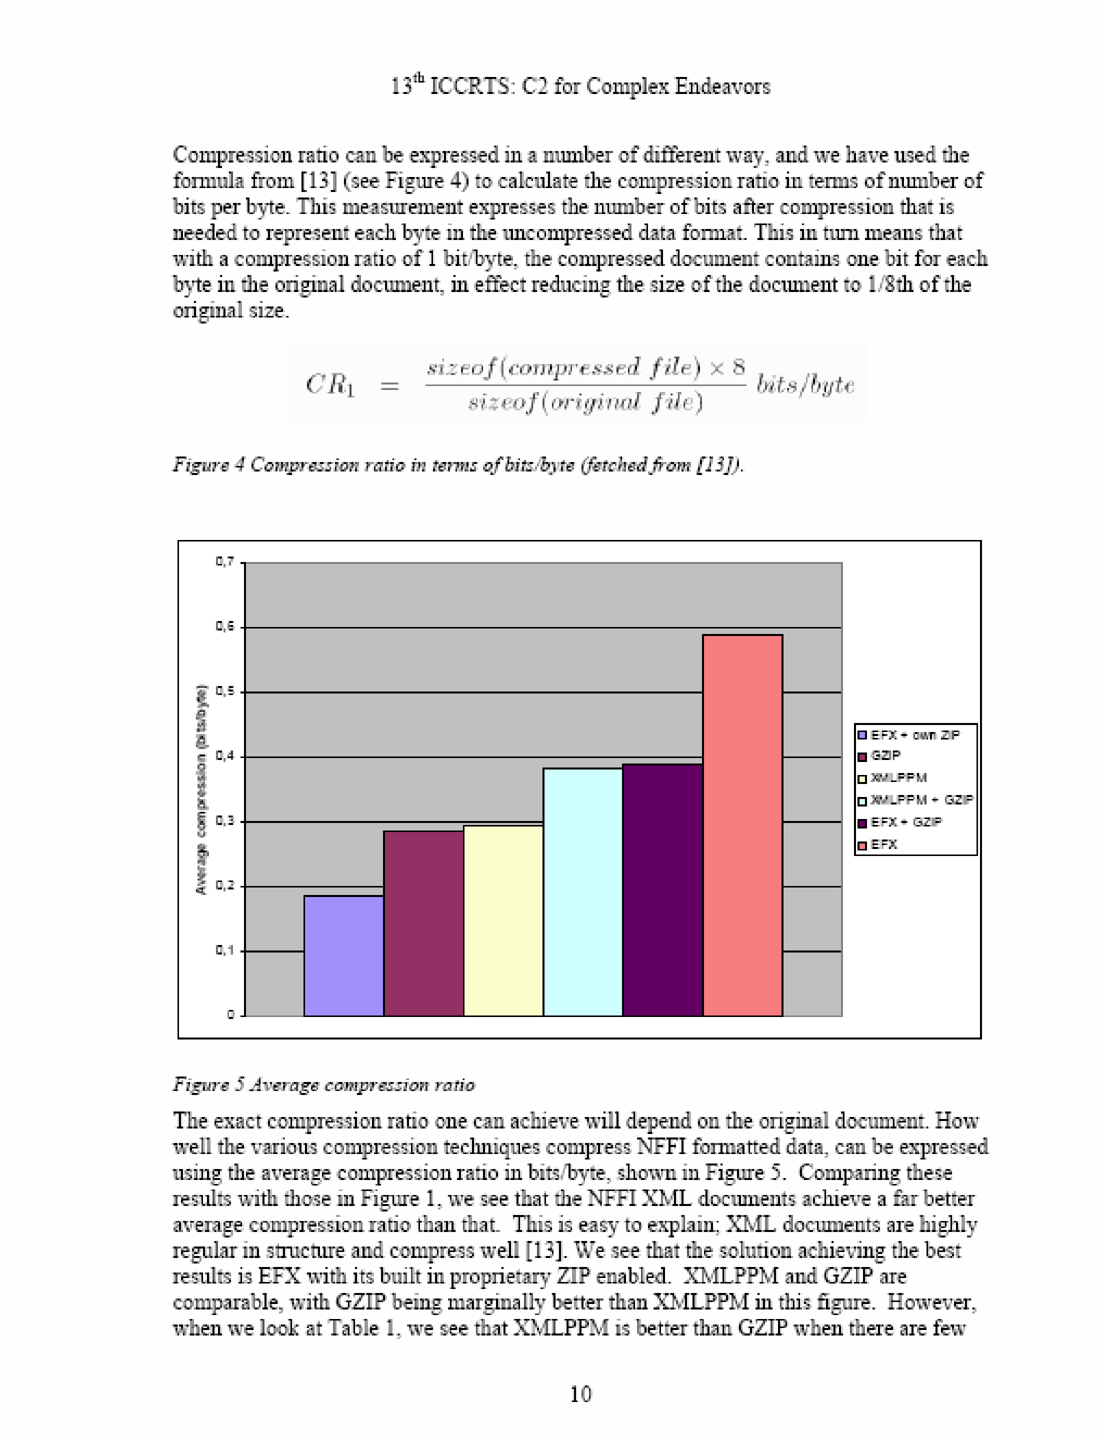

Compression ratio can be expressed in a number of different way, and we have used the formula from [13] (see Figure 4) to calculate the compression ratio in terms of number of bits per byte. This measurement expresses the number of bits after compression that is needed to represent each byte in the uncompressed data format. This in turn means that with a compression ratio of 1 bit/byte, the compressed document contains one bit for each byte in the original document, in effect reducing the size of the document to 1/8th of the original size.

$$CR_1 = \frac{sizeof(compressed\ file) \times 8}{sizeof(original\ file)}\ bits/byte$$

Figure 4 Compression ratio in terms of bits/byte (fetched from [13]).

Figure 5 Average compression ratio

The exact compression ratio one can achieve will depend on the original document. How well the various compression techniques compress NFFI formatted data, can be expressed using the average compression ratio in bits/byte, shown in Figure 5. Comparing these results with those in Figure 1, we see that the NFFI XML documents achieve a far better average compression ratio than that. This is easy to explain; XML documents are highly regular in structure and compress well [13]. We see that the solution achieving the best results is EFX with its built in proprietary ZIP enabled. XMLPPM and GZIP are comparable, with GZIP being marginally better than XMLPPM in this figure. However, when we look at Table 1, we see that XMLPPM is better than GZIP when there are few

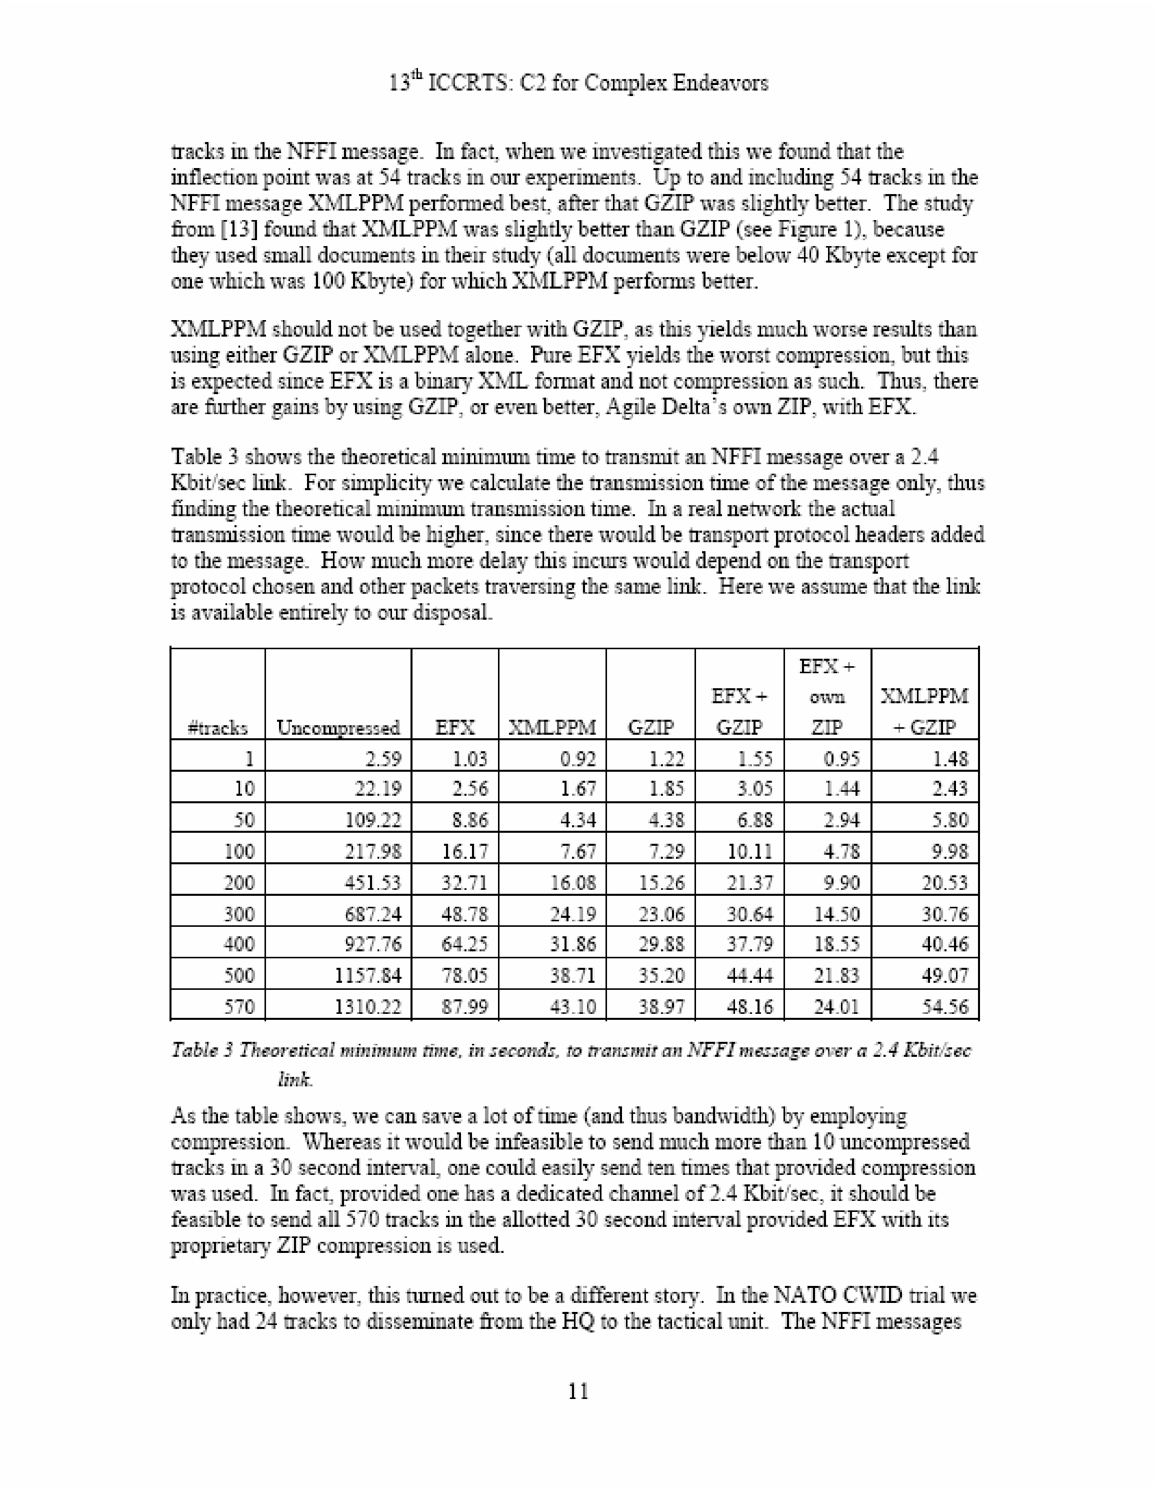

tracks in the NFFI message. In fact, when we investigated this we found that the inflection point was at 54 tracks in our experiments. Up to and including 54 tracks in the NFFI message XMLPPM performed best, after that GZIP was slightly better. The study from [13] found that XMLPPM was slightly better than GZIP (see Figure 1), because they used small documents in their study (all documents were below 40 Kbyte except for one which was 100 Kbyte) for which XMLPPM performs better.

XMLPPM should not be used together with GZIP, as this yields much worse results than using either GZIP or XMLPPM alone. Pure EFX yields the worst compression, but this is expected since EFX is a binary XML format and not compression as such. Thus, there are further gains by using GZIP, or even better, Agile Delta's own ZIP, with EFX.

Table 3 shows the theoretical minimum time to transmit an NFFI message over a 2.4 Kbit/sec link. For simplicity we calculate the transmission time of the message only, thus finding the theoretical minimum transmission time. In a real network the actual transmission time would be higher, since there would be transport protocol headers added to the message. How much more delay this incurs would depend on the transport protocol chosen and other packets traversing the same link. Here we assume that the link is available entirely to our disposal.

#tracks	Uncompressed	EFX	XMLPPM	GZIP	EFX + GZIP	EFX + own ZIP	XMLPPM + GZIP
1	2.59	1.03	0.92	1.22	1.55	0.95	1.48
10	22.19	2.56	1.67	1.85	3.05	1.44	2.43
50	109.22	8.86	4.34	4.38	6.88	2.94	5.80
100	217.98	16.17	7.67	7.29	10.11	4.78	9.98
200	451.53	32.71	16.08	15.26	21.37	9.90	20.53
300	687.24	48.78	24.19	23.06	30.64	14.50	30.76
400	927.76	64.25	31.86	29.88	37.79	18.55	40.46
500	1157.84	78.05	38.71	35.20	44.44	21.83	49.07
570	1310.22	87.99	43.10	38.97	48.16	24.01	54.56

Table 3 Theoretical minimum time, in seconds, to transmit an NFFI message over a 2.4 Kbit/sec link.

As the table shows, we can save a lot of time (and thus bandwidth) by employing compression. Whereas it would be infeasible to send much more than 10 uncompressed tracks in a 30 second interval, one could easily send ten times that provided compression was used. In fact, provided one has a dedicated channel of 2.4 Kbit/sec, it should be feasible to send all 570 tracks in the allotted 30 second interval provided EFX with its proprietary ZIP compression is used.

In practice, however, this turned out to be a different story. In the NATO CWID trial we only had 24 tracks to disseminate from the HQ to the tactical unit. The NFFI messages

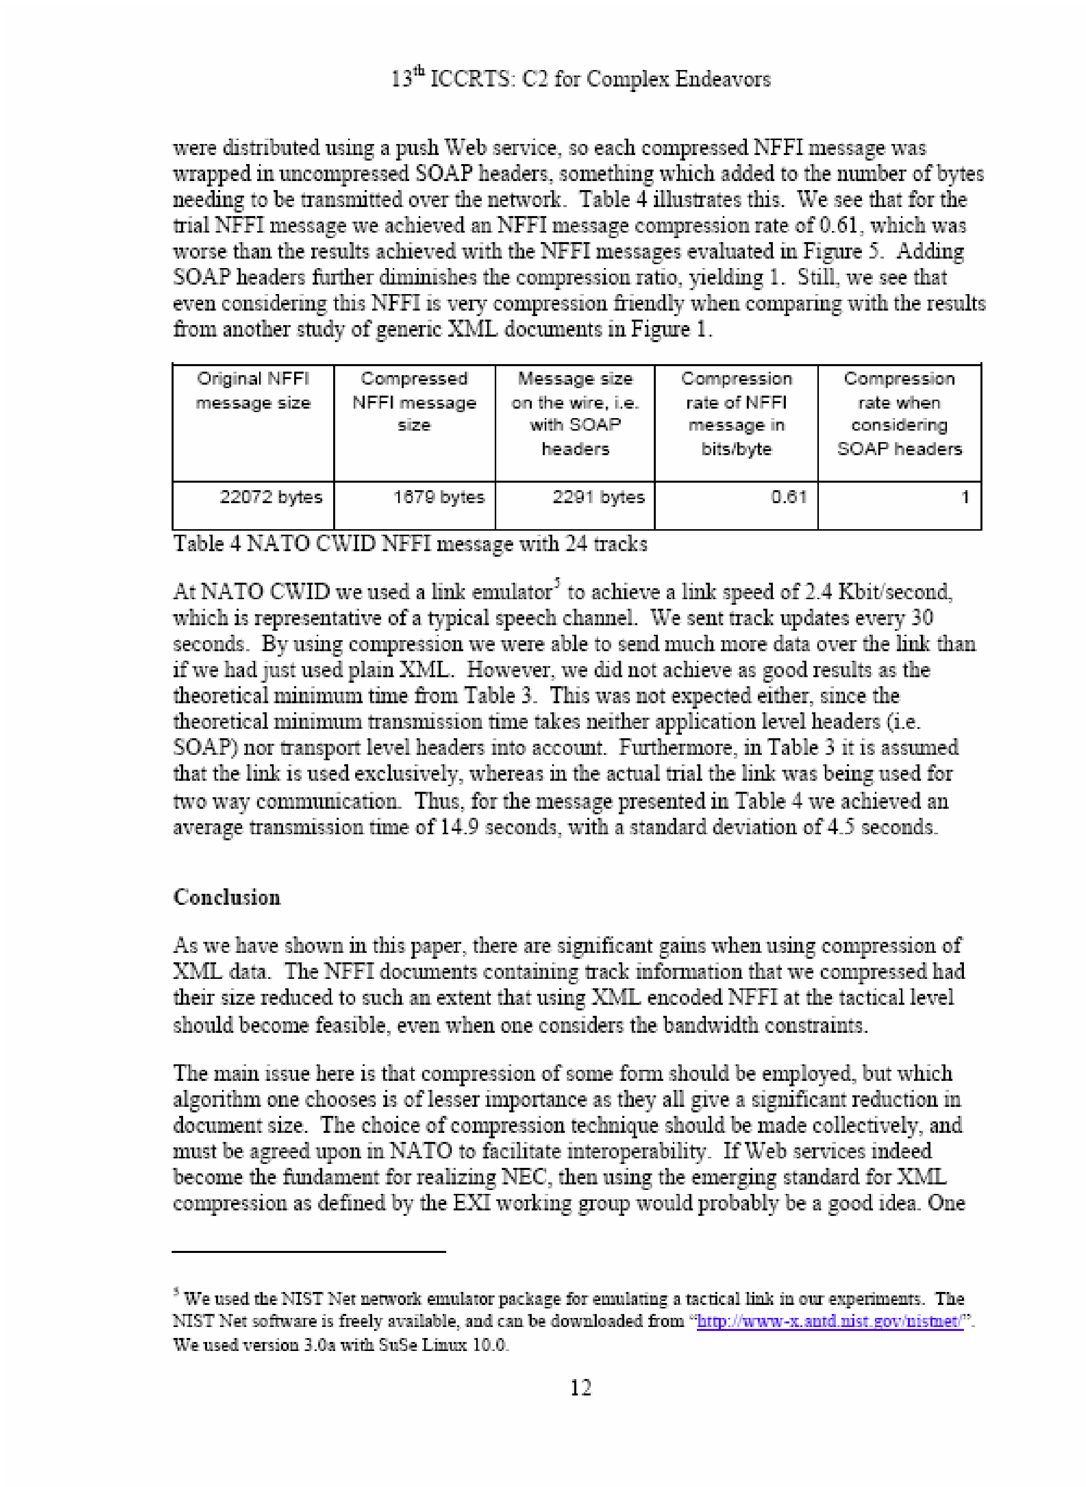

were distributed using a push Web service, so each compressed NFFI message was wrapped in uncompressed SOAP headers, something which added to the number of bytes needing to be transmitted over the network. Table 4 illustrates this. We see that for the trial NFFI message we achieved an NFFI message compression rate of 0.61, which was worse than the results achieved with the NFFI messages evaluated in Figure 5. Adding SOAP headers further diminishes the compression ratio, yielding 1. Still, we see that even considering this NFFI is very compression friendly when comparing with the results from another study of generic XML documents in Figure 1.

Original NFFI message size	Compressed NFFI message size	Message size on the wire, i.e. with SOAP headers	Compression rate of NFFI message in bits/byte	Compression rate when considering SOAP headers
22072 bytes	1679 bytes	2291 bytes	0.61	1

Table 4 NATO CWID NFFI message with 24 tracks

At NATO CWID we used a link emulator[5] to achieve a link speed of 2.4 Kbit/second, which is representative of a typical speech channel. We sent track updates every 30 seconds. By using compression we were able to send much more data over the link than if we had just used plain XML. However, we did not achieve as good results as the theoretical minimum time from Table 3. This was not expected either, since the theoretical minimum transmission time takes neither application level headers (i.e. SOAP) nor transport level headers into account. Furthermore, in Table 3 it is assumed that the link is used exclusively, whereas in the actual trial the link was being used for two way communication. Thus, for the message presented in Table 4 we achieved an average transmission time of 14.9 seconds, with a standard deviation of 4.5 seconds.

Conclusion

As we have shown in this paper, there are significant gains when using compression of XML data. The NFFI documents containing track information that we compressed had their size reduced to such an extent that using XML encoded NFFI at the tactical level should become feasible, even when one considers the bandwidth constraints.

The main issue here is that compression of some form should be employed, but which algorithm one chooses is of lesser importance as they all give a significant reduction in document size. The choice of compression technique should be made collectively, and must be agreed upon in NATO to facilitate interoperability. If Web services indeed become the fundament for realizing NEC, then using the emerging standard for XML compression as defined by the EXI working group would probably be a good idea. One

[5] We used the NIST Net network emulator package for emulating a tactical link in our experiments. The NIST Net software is freely available, and can be downloaded from "http://www-x.antd.nist.gov/nistnet/". We used version 3.0a with SuSe Linux 10.0.

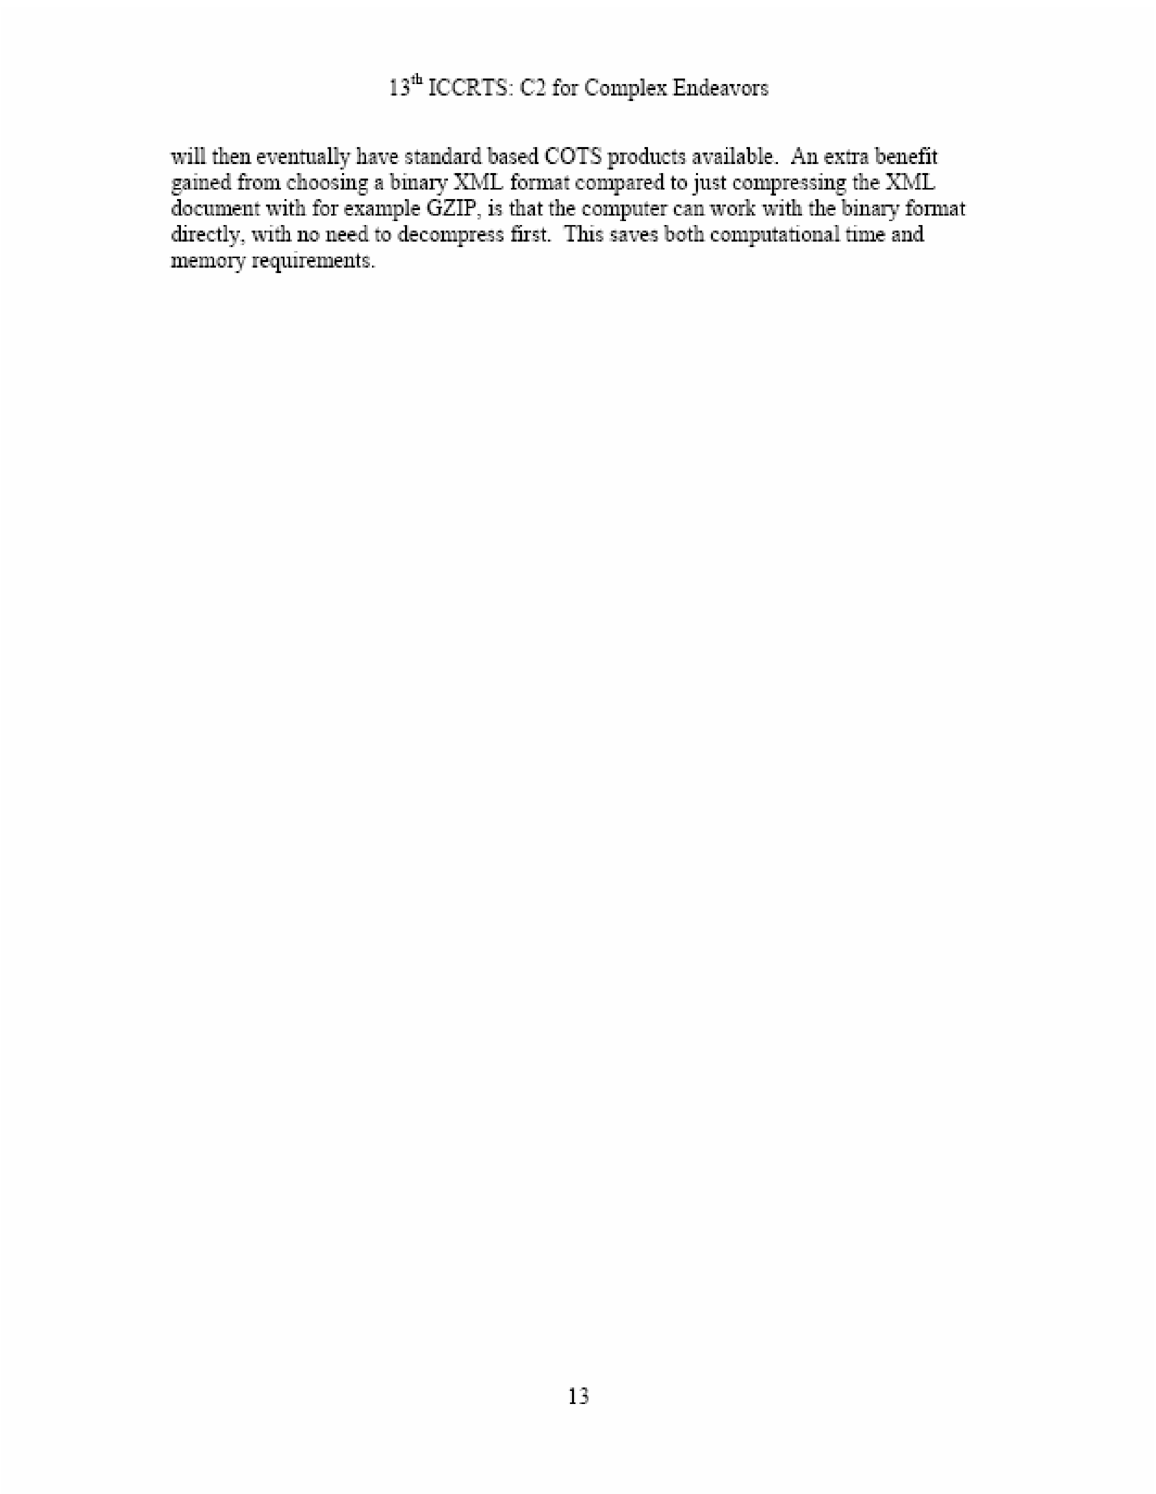

will then eventually have standard based COTS products available. An extra benefit gained from choosing a binary XML format compared to just compressing the XML document with for example GZIP, is that the computer can work with the binary format directly, with no need to decompress first. This saves both computational time and memory requirements.

References

[1] Network Centric Operations Industry Consortium (NCOIC)
http://www.ncoic.org/home

[2] Haakseth, R., Hadzic D., Lund, K., Eggen, A., Rasmussen, R. E., "Experiences from implementing dynamic and secure Web services", 11th Coalition Command and Control in the Networked Era (ICCRTS), September 2006, Cambridge, UK.

[3] Hafsøe, Trude, Johnsen, Frank T., Lund, Ketil, and Eggen, Anders
"Adapting Web Services for Limited Bandwidth Tactical Networks"
12th International Command and Control Research and Technology Symposium, (ICCRTS), June 2007, Newport, RI, USA.

[4] Hadzic, D. et al, "Web Services in networks with limited data rate" (in Norwegian), FFI-Report 2006/03886, ISBN 978-82-464-1049-4

[5] Salomon, David. Springer 2000. "Data Compression – The Complete Reference, 2nd edition", ISBN 0-387-95045-1

[6] Erl, Thomas. Prentice hall 2005. "Service-Oriented Architecture – Concepts, Technology, and Design", ISBN 0-13-185858-0

[7] Erl, Thomas. Prentice hall 2004. "Service-Oriented Architecture – A Field Guide to Integrating XML and Web Services", ISBN 0-13-142898-5

[8] Johnsen, Frank T., Hafsøe Trude, and Lund, Ketil, "NATO CWID 2007 disadvantaged grids experiments", FFI-Notat 2007/02063

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

[10] J. Flathagen and L. Olsen, "NORMANS KKI" (in Norwegian), FFI Facts, November 2006, http://www.mil.no/multimedia/archive/00086/Faktark-NORMANS-KKI-_86445a.pdf

[11] Efficient XML Interchange Working Group
http://www.w3.org/XML/EXI/

[12] Efficient XML Interchange (EXI) Format 1.0, W3C Working Draft July 2007
http://www.w3.org/TR/2007/WD-exi-20070716/

[13] W. Ng, W.-Y. Lam, and J. Cheng, "Comparative Analysis of XML Compression Technologies", World Wide Web 9(1), pages 5-33, Kluwer Academic Publishers, March 2006.

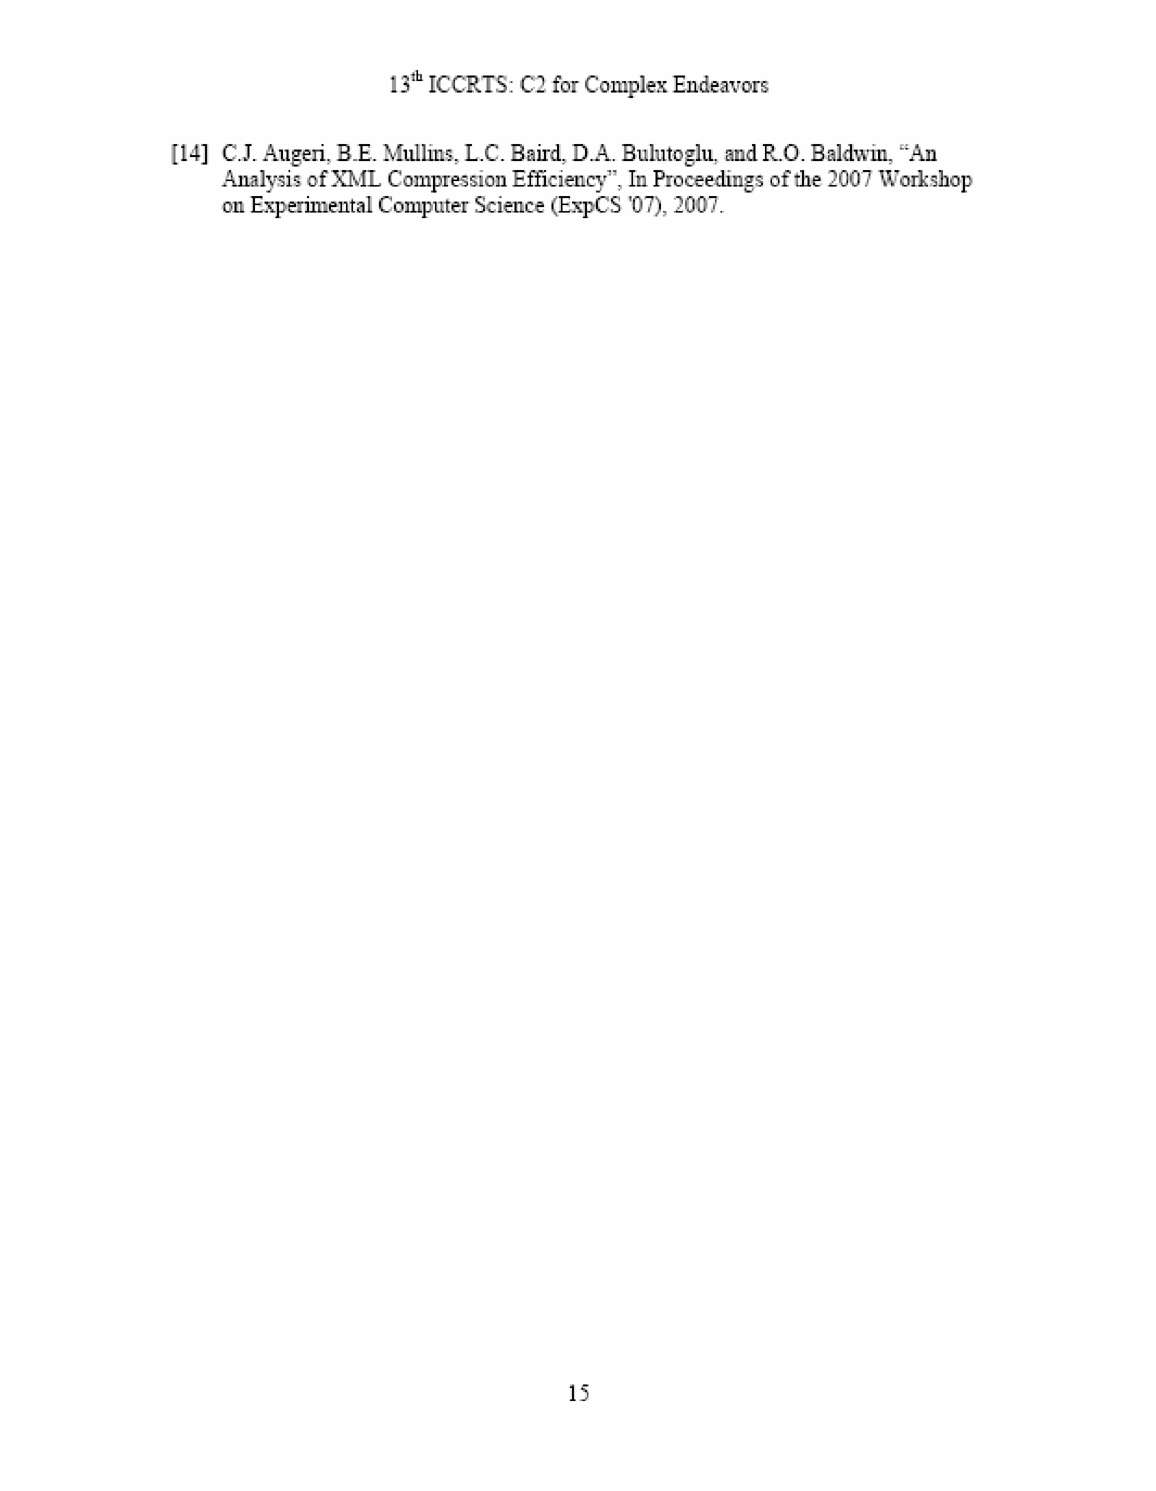

[14] C.J. Augeri, B.E. Mullins, L.C. Baird, D.A. Bulutoglu, and R.O. Baldwin, “An Analysis of XML Compression Efficiency”, In Proceedings of the 2007 Workshop on Experimental Computer Science (ExpCS '07), 2007.