

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture

Name of Author/Co-Author(s): Marco van der Meijden, Msc.

Company: TNO-FEL

Address: Oude Waalsdorperweg 63, P.O. Box 96864
2509 JG THE NETHERLANDS

Telephone: +31 (0)6 51293739

Fax: +31 (0)70 3740652

Email: vandermeijden@fel.tno.nl; c2sc.mip@rnla.mindef.nl

ABSTRACT

The Royal Netherlands Army (RNLA) has been working on its second generation C2 Information System since 2000. Its first generation Integrated Staff Information System (ISIS) that is currently deployed is based on ATCCIS specifications and uses the ATCCIS Replication Mechanism for database to database replication. For the next generation system, the C2 Workstation, the goal was to incorporate NCW/NCO concepts in its architecture. This was translated into the requirement for a scaleable and robust Distributed Information Exchange Architecture that makes maximum use of COTS technology and runs on top of the IP based RNLA network Infrastructure TITAAN (Theatre Independent Tactical Army and Air force Network). Implementation of the new architecture is on track for an initial fielding by Q1 2004.

In this paper, several key characteristics of the C2WS architecture are described divided into four architectural layers: the Operational Layer, the Data Replication Layer, the Middleware Layer and the IP (Multicast) Layer. The RNLA has developed a custom object database and database replication mechanism that is expected to scale to 10.000 or more databases and allows selective distribution and concurrent manipulation of data. The replication mechanism runs on top of a COTS Publish/Subscribe messaging middleware product. The middleware makes use of a reliable IP multicast protocol to allow efficient point-to-multipoint communication.

Our conclusion is that the C2WS architecture meets many of the goals of Network Centric Operations. In the realisation of our architecture, we have identified several areas where COTS technology imposes limitations, for example in the area of security in Publish/Subscribe based communication and achieving fine grained selectivity of data distribution down to the IP-level. At the end of this paper we give an overview of technologies that we are evaluating to address these limitation.

1 INTRODUCTION

The current RNLA operational C2 system that supports Brigade and Division level units is ISIS 2.5 (Integrated Staff Information System). Its development started in 1995 and its architecture was based on the ATCCIS specification for interoperable systems. ATCCIS¹ defined both a data model [3], and the ATCCIS Replication Mechanism [1] [2] for the exchange of information between national systems. The RNLA was the first nation to base its own national system on this specification in order to facilitate interoperability with other nations.

¹ The ATCISS program was merged into MIP [4] and the data model is currently referred to as the C2IEDM (C2 Information Exchange Data Model).

Paper presented at the RTO SCI Symposium on "Architecture for Network-Centric Operations", held in Athens, Greece, 20-22 October 2003, and published in RTO-MP-SCI-137.

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture

In 2000 the development of the architecture of the next generation C2 system (which is internally referred to as the C2 Workstation (C2WS), but will be fielded under the name of ISIS 3.0) started and is now expected to be fielded Q1 2004 for the 11 Air Mobile Brigade. In parallel to the C2 Workstation, the IP based network infrastructure TITAAN (Theatre Independent Tactical Army and Air force Network) [9] is being developed.

For the next generation architecture there were three main goals; COTS based, improve on ISIS 2.5 and incorporate Network Centric Operation concepts if possible. The RNLA had adopted the concepts of NCW and wanted to have them incorporated into future C2 systems.

2 FROM CONCEPTS TO ARCHITECTURE

In this paper, we wanted to focus on several of the key architectural issues we encountered during the project instead of covering the whole architecture at a high level of abstraction. In order to come to such a scoping in a coherent way, we will structure the system description into a number of architectural layers:

- Operational layer
- Data replication layer
- Middleware layer
- IP Layer

2.1 Operational layer

For the C2 Workstation architecture, we focussed initially on the basic C2 support from Battalion to Division in the area of Situational Awareness. In our experience, information management was likely to be the most difficult aspect of the Data Replication Layer that needed to be addressed in new architecture. In the ISIS 2.5 system, all information is replicated everywhere, which is undesirable from both a scalability and security point of view. Experiences with selective replication based on point-to-point contracts and filtering within ATCISS were not very successful, yielding poor performance of filter implementations and contract management overhead and scalability issues, so we were looking for alternative approaches.

In all contemporary C2 systems we investigated, information management was either not performed at all and information was disseminated everywhere like the ISIS 2.5 system, or carried out by the operator creating the appropriate ADatP-3 messages according to doctrine and Standard Operating Procedure (SOP). We envisioned that apart from the doctrinal based information flows that are currently used to generate a Common Operational Picture (COP), collaboration (for example a Collaborative Operational Picture [23]), would become important in the future. Also, in the shift to network centric operations the reach and richness of information exchange would likely increase beyond the current doctrinal rules.

What we needed prior to designing the technical architecture was an operational level concept of information dissemination and information management that was both technically feasible to implement and simple enough to be understood and employed by operators.

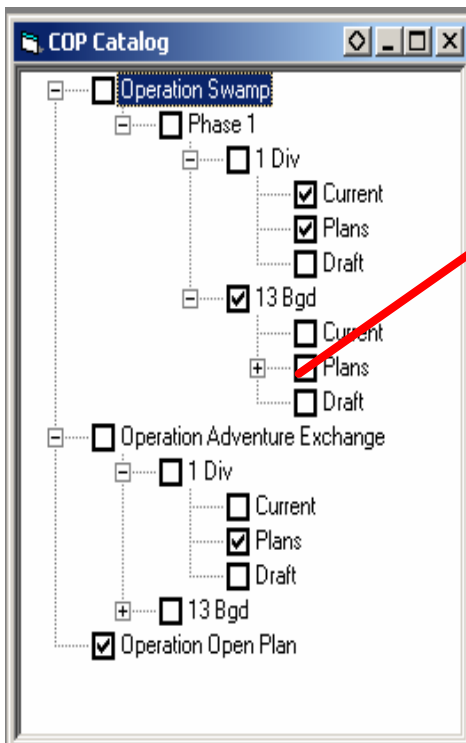
Our solution to the problem was as follows:

- All information is organized into **contexts**.
- A context is a set of information that gives meaning to the individual pieces of information it contains. An example of a context is the Recognised Enemy Picture created by a unit.
- Multiple users may work collaboratively within the same context.
- A context is defined as representing an unambiguous set of information to third parties. Any ambiguities concerning data must be resolved explicitly by using multiple contexts. For example by having a context that represents an alternative Enemy Picture.
- Contexts are hierarchically organized into the COP Catalog. Users can view this catalog and select the contexts they wish to access.
- The user is free to organise the COP Catalog prior and during each mission according to SOPs, there is not a statically defined set of contexts or workflow that is enforced by the system.

On top of these operational concepts we built the following technical concepts:

- Each context is a unit of
 - data distribution
 - security access

By defining the use of contexts in this way, we have effectively defined a Publish/Subscribe based operational architecture, where the existence of available information is published to all users who then subsequently can subscribe to the information that is relevant to them.



Subscribe to
relevant
information

Example of the COP Catalog. Users can select which information is relevant to them and subscribe to it.

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture

2.2 Data replication layer

The architecture of the data replication layer was first of all driven by requirements derived from the context concept. Based on these concepts the following design was derived:

- A user should be able to selectively subscribe to a context. Once subscribed to a context, only that context would be replicated and kept synchronised at the local system..
- Selective replication of the COP Catalog. It was very evident from the beginning that a single Catalog of information (i.e. the available contexts that a user can subscribe to) could become a bottleneck by itself if a very large numbers of contexts would have to be managed. Therefore the COP Catalog itself would have to be divided into multiple sections (that are Contexts themselves) in analogy to a folder structure. While traversing the structure, new sections would be subscribed-to on the fly.
- Support for collaborative manipulation of shared data.
- Each system should be able to work stand-alone and contain a local database.

Other requirements were based on the lessons learned from the ISIS 2.5 ATCISS based replication mechanism:

- Synchronisation should be possible on a per context basis. Instead of synchronising the whole database at once, it should be possible to synchronise each context independently. This would allow priority base sequencing of the order in which contexts are synchronised when network connectivity is restored.
- Synchronisation of the current data of a context only, with the option to synchronise historical data up to a period back in time. Experiences with ISIS 2.5 showed that after weeks of exercise the total database would become so large that complete synchronisation of new or failed nodes becomes a serious bottleneck. Therefore the current information should be synchronised first, and synchronisation of historical data should occur incrementally up to some relevant point back in time.
- Synchronisation should be possible in parallel with receiving and processing updates.
- Each system should be able to work stand alone or off-line and not be dependant on a server that could become a single point of failure. Each individual (client) machine would therefore have its own database and replication service. Up to 10.000 replicated databases should be supported.
- Switches from active to reserve command posts should be supported in both controlled and uncontrolled fashion. In general, network (re)-partitioning should be supported. In such situations, it is possible that during resynchronisation conflicts become apparent in data that has been manipulated concurrently [17]. This particular issue is described later on in more detail.
- The replication layer should be as plug and play as possible and be able to deal with changes in network topology.
- Able to deal with bandwidth as low as 32 Kbit/s on a WAN². Initially aimed at Battalion level units and upwards, current RNLA digital communication infrastructure could limit the bandwidth to 64 Kbit/s (bandwidth that is shared with other applications).

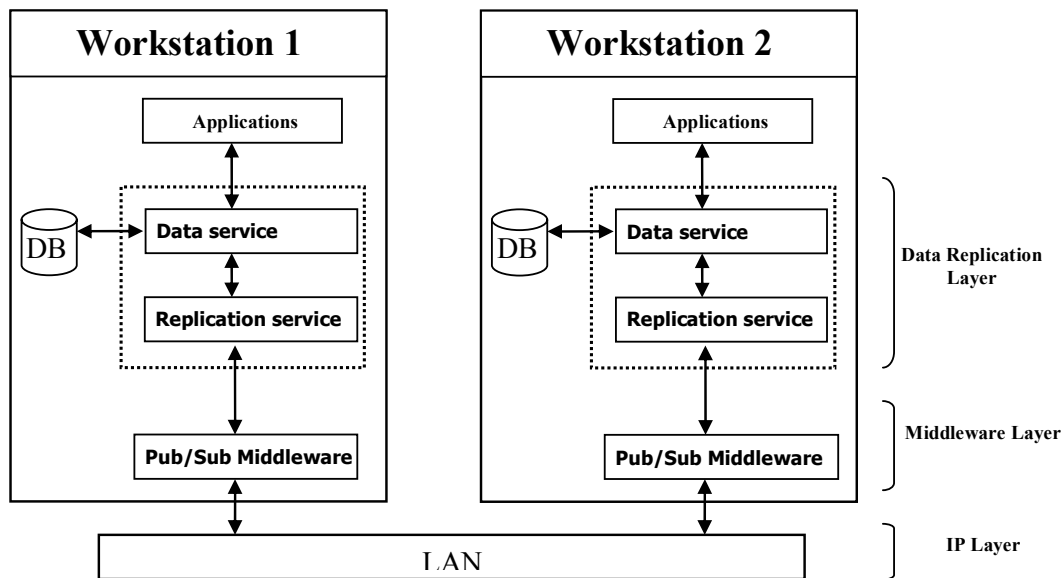
² Of course bandwidth consumption is dependant on the number of transactions that are created within the overall system. For typical RNLA use, a maximum of 400 user generated transactions per hour has been observed in exercises.

These requirements were not met by any commercial product at the time and by our knowledge this is still the case. The main deficiencies of commercial database replication products³ are:

- Designed for a static infrastructure and not designed to accommodate changes to network and replication topology. Generally replication topology is configured manually and can not be changed without stopping the replication services.
- No selective replication and distinction between current and historic information resynchronisation.
- Limited support for collaboration and conflict resolution.

The product that came closest to meeting these requirements was MS Active Directory, an object database primarily aimed at providing distributed directory services. It provides efficient replication of current data and efficient conflict resolution, but lacks the functionality to deal with dynamic network topologies.

Therefore the RNLA decided to develop its own data replication layer that was able to meet all of the above requirements. Its basic architecture is shown in the Figure below. The Data Replication Layer consists of two parts: the Data service that provides a distributed object database to applications and the Replication service that ensures that each local data store is replicated and kept in sync. A normal relational database is used for persistency of the object data.



Basic Architecture of the C2 Workstation

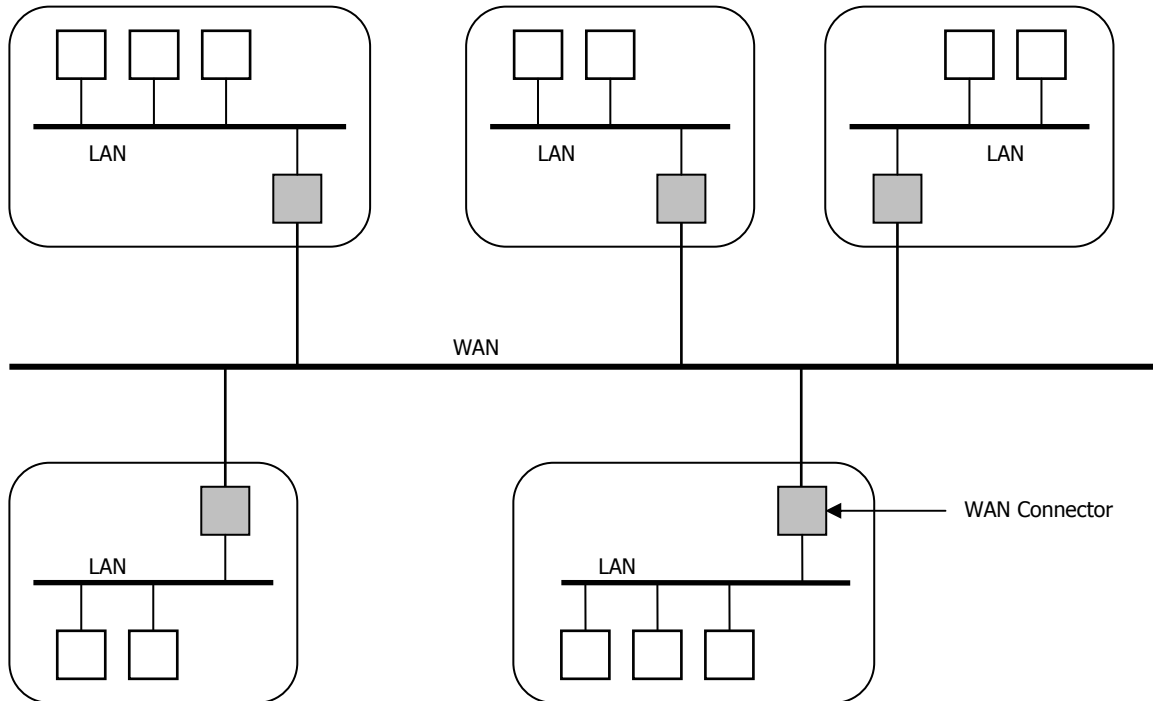
The C2WS Data Replication Layer consists essentially of a distributed object database with a dedicated replication mechanism [6]. The replication service works as a bus architecture, and communicates using Publish/Subscribe Middleware. In our configuration, each Command Post has a LAN that connects all workstations to a single bus. The Publish/Subscribe Middleware (described in more detail in the next section) effectively hides the underlying network. The synchronisation protocol is designed to use Publish and Subscribe operations only. In its simplest configuration, the Middleware could be setup to reach the whole network, effectively creating a single information bus that spans the whole system. This however is

³ A detailed overview of the respective features of individual products is out of scope for the article. Products investigated included the ATCISS ARM specification, Oracle 9i Replication and MS Active Directory.

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture

not scaleable and wastes bandwidth for example as all replication protocol messages may reach the whole network.

Therefore we have set-up a hierarchical network where the WAN forms a second bus to which all LANs are connected via a WAN Connector. This provides scalability as synchronization between workstations of different LANs never occurs directly but always through WAN Connectors. Replication protocol messages on the LAN bus are thus separated from the WAN bus. On the LAN and WAN, the same data replication protocol is used. WAN Connectors forward updates and synchronise data from LAN to WAN and visa-versa.



Each LAN is connected to the WAN via a WAN Connector.

The data replication protocol can be characterised as 'Multi-Master with loose consistency and convergence'. Multi-master means that data can be manipulated at each systems data store. Loose consistency means that not all data stores need to have the same data at the same time. There is thus no locking protocol like two-phase commit. The use of a locking protocol would mean that the liveliness [15] becomes a problem when connections are lost and locks have to time-out before data can be updated. Convergence means that the replication protocol strives continuously to get all systems synchronized to all systems and each system's local data store will converge to contain the same information. Such a convergence may only occur when updates are no longer created and the system as a whole reaches a stable state. As locking or transaction serialisation is not performed, conflict resolution becomes very important.

We have addressed conflict resolution in two ways:

- The actual or current value of a record is determined by a version number that increases whenever the record is updated. If records are updated concurrently, the value with the highest version number wins. The heuristic is thus that the value that has been updated the most is considered to be the most actual value. We have considered using the latest transaction time, but relying on time synchronisation is very dangerous⁴.
- During synchronisation of concurrently updated data, inconsistencies may appear, such as violations of referential integrity or business rules. We deal with this in the following ways:
 - Each application is designed to be able to handle inconsistent data. If it detects inconsistent data, it immediately makes it consistent for displaying purposes. If the user has write permission, the user is asked to review and save the corrected data.
 - When a large group of users has write permission, there is chance of concurrent correction. We do not yet do anything specific to handle this situation other than using a ‘delete only’ policy in the correction of data. This means that corrections should be performed primarily by deleting the data that causes the inconsistency. This may in the worst case lead to the complete deletion of the inconsistent data due to different concurrent corrections after which the user must re-enter new and consistent data. We estimate that actual conflicts will occur very rarely and that the current conflict resolution strategy will be acceptable in practise.

As can be seen, we have made the trade-off between consistency and liveliness in favour of liveliness. In case of a C2 system that has to support near-real time COP in a very dynamic network environment this seems a valid decision.

2.3 Middleware layer


The replication protocol is designed to use Publish/Subscribe operations only. This functionally can be provided by Middleware that provides Publish/Subscribe based messaging. An extensive product evaluation was performed to select a COTS product [5]. This evaluation included products such as Talarian Smartsockets⁵, Tibco Rendezvous and MSMQ. Most products offer two alternative messaging implementations (see also Figures below):

- **Message server based.** By configuring point-to-point connections between message servers an overlay network is created. Efficient point-to-multi-point publication of messages is generally supported as well as keeping track of subject interest to avoid sending messages for which there are no subscribers. None of the evaluated products was able to accommodate changing networks topologies automatically and needed manual reconfiguration [8]. Routing of messages between servers is done through routing algorithms that only take static weights into account. If the network load or available bandwidth changes, selection of non-optimal routes may lead to a waste of bandwidth.
- **IP Multicast based.** If IP Multicast is supported by the WAN, it effectively provides a Plug&Play messaging solution that adapts its routing automatically to changes in the network. We decided to choose IP Multicast mainly based on the fact that no configuration other than the Multicast addresses was required.


⁴ If a systems clock is back in time, its updates will not show up. If it's ahead of time, the data it creates will remain current and cannot be updated by other systems until the time on those systems catches up.

⁵ Talarian has been taken over by TIBCO recently who is continuing their product line.

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture



Publish/Subscribe messaging using message servers.



Publish/Subscribe messaging using IP Multicast

The replication protocol maps all the messages used to replicate each context to a unique subject⁶. This allows the Messaging Middleware to only route replication messages to systems that are subscribed to that context. This works automatically when using message servers. When using Multicast, selectivity can only be achieved by mapping each context to a unique IP Multicast address. The automatic mapping of subjects to Multicast addresses is not supported however⁷. In our default configuration we use a unique Multicast Address for each LAN and only one multicast address for all WAN communication. This severely limits the selectivity of information dissemination because each WAN Connector will receive the information from all contexts that is published on the WAN Multicast address.

To achieve finer grained selectivity of information distribution in the future, each context (or a set of contexts) will be mapped to a separate Multicast address and the multicast address will be added to the COP Catalog. Addresses can either be handed out by pre-configuring a pool of address at each WAN Connector or choosing one at random from a specific range⁸. As each multicast address that is used causes state information on the multicast distribution tree to be maintained in the routers and additional protocol traffic is generated to maintain this state, the ideal number of multicast addresses to be used is still subject of study [16].

2.4 IP Multicast

The COTS products investigated all implement Pragmatic General Multicast (PGM). PGM is an experimental IETF RFC [19] that specifies a reliable multicast protocol that addresses packet loss which may be caused by congestion or bit errors that are likely to occur on a tactical WAN. The use of PGM has the advantage above proprietary protocols that it can be supported by routers, thus enhancing the scalability. We are not using router support however as it consumes a lot of router resource and is currently only supported by a single version of the router operating system. The only disadvantage of PGM is that although a specification for congestions control exist [22], this is not implemented in the products we evaluated. This means that the transmit rate of a publisher may (temporarily) exceed the bandwidth of WAN links congestion in routers. Because there is no feedback mechanism that slows down the transmit rate of the senders like TCP/IP, the traffic may potentially overload WAN links leading to congestion. To avoid the congestion problem, we are using Weighted Fair Queuing with Micro Flowing for IP multicast. The Micro Flowing will detect the multicast traffic as a flow and assign a separate queue for it. This effectively prevents the IP Multicast traffic to overload the WAN links as it isolates the congestion to the IP Multicast traffic only. However, once congestion occurs on Multicast traffic, the congestion will remain until the Publishers stop sending. A partial fix to the problem is to statically limit the transmit rate of a single Publisher so that it remains below 50% of the available bandwidth. This becomes less effective however when the number of publishers grows.

PGM is an IP Multicast transport protocol. For multicast routing Protocol Independent Multicast (PIM) is used. We have succeeded in creating a configuration for PIM Sparse Mode⁹ that is able to handle network (re-)partitioning and router failure.

⁶ Publish/Subscribe Middleware generally uses subjects such as 'global.stocks.ms' to tag messages and specify receiver interest.

⁷ Some Internet drafts exist, such as MADCAP, but this requires a central sever that issues addresses.

⁸ IPv6 will make selecting an address at random much more attractive as the address range is greatly expanded reducing the chance of collisions.

⁹ PIM has two modes, Sparse and Dense. Sparse mode is generally recommended but requires configuration.

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture

3. SECURITY

Security in Publish/Subscribe Middleware is much more complex than for point-to-point protocols [18]. The main problem is the management of the keys that are used for the encryption of messages. All publishers and subscribers on a context need the same (symmetrical) key in order to communicate. In case IP Multicast is used, no Middleware product offered what we considered a viable solution. The solutions that were provided required a single system to do key management and access control. As this would introduce a single point of failure into our architecture we currently rely on link encryption and 'cocooning' of the infrastructure.

4. CONCLUSION

In our point-of-view the C2WS data replication architecture meets many goals of NCO. It is implemented on top of COTS Middleware and IP networks. So far it has been successfully evaluated on various networks, including satellite links and digital radio. Many trade-offs were made in order to realise an architecture that delivers a near real time COP in a highly dynamic network environment. For other types of applications, other trade-offs may be appropriate and we don't foresee that a single data replication mechanism will be able to meet all requirements. The RNLA for example also deploys a separate Tactical Messaging System (TMS).

COTS data replication products will most likely not be able to offer equivalent functionality in the near future. RNLA will therefore continue to develop its data replication architecture in the next years in a number of areas:

- The context concept requires that all data is partitioned into disjunctive contexts. This is not always desirable. We are expanding our replication architecture to be able to share data between contexts.
- The data replication layer will be extended to be able to handle low bandwidth radio networks of 19.200 Kbit/s. We expect this to be mainly an issue of timer optimisation and compression of the size of messages. As the basic replication protocol will remain the same this means we will have single data distribution mechanism from the vehicle level upwards that is only different in its configuration for the available bandwidth and Middleware that is used. As IP Multicast is most likely not an option on such low bandwidth links, we expect to make direct use of the messaging interface of the radio.
- Emerging peer-to-peer technology may become an alternative to the current Publish/Subscribe Middleware [7] [14]. For example Bayeux [11] or Scribe [13] offer Publish/Subscribe messaging on top of a self organising peer-to-peer substrate such as Tapestry [10] or Pastry [12]. These peer-to-peer technologies generally only work well¹⁰ in very large networks such as the internet. The advantage however is that because communication between peers is done through TCP/IP, no Multicast infrastructure is required. An experiment to port the current replication service to a peer-to-peer product is currently underway.
- Many Peer-to-peer products base their security architecture on PKI. We are researching if the same concepts can be used to implement security in our replication protocol independently of the underlying Middleware.
- The lack of congestion control in PGM implementations remains a problem. We are hoping for implementations of the PGM congestion control [22] in the near future. In the longer term, reliable multicast specifications in the IETF such as NORM [20] and associated congestion control specification [19] may become an alternative to PGM.

¹⁰ Routing is suboptimal in smaller networks.

References

- [1] ATCCIS Working paper 14-3, ATCCIS Replication Mechanism (ARM) Consolidated Specification, Edition 3.0, 5 June 2000
- [2] ATCCIS permanent working group, "ATCCIS Replication Mechanism" requirements, 1995
- [3] ATCCIS permanent working group, "AWP ST2 The Land C2 Information Exchange Data Model" Edition 2.0, 2000 (Proposed STANAG AdatP-32)
- [4] MIP Multilateral Interoperability Program briefing notes (MBN), March 2003. http://www.mip-site.org/Public_documents/MBN-SH-PMG-Edition1.1.pdf.
- [5] Tactical information distribution with Message Oriented Middleware, Julian de Wit, TNO report FEL-01-S234, 2001
- [6] Tactical Information Distribution using XML, M.G. van der Meijden, TNO report FEL-01-A124, 2001
- [7] Peer-to-peer technology in a C2 environment: a preliminary study, M. van Lent, TNO-FEL Traineeship report 20030716, 2003
- [8] Automatic reconfiguration of Tibco Rendezvous Routing Daemons in dynamic networks, D.J. Noort, TNO report FEL-03-S075, April 2003.
- [9] TITAAN Phase 1 (HRF HQ): Detailed Technical Design, J. Vytopil. Royal Netherlands Army, C2SC, 2002. Version 1.01.
- [10] Tapestry: A resilient global-scale overlay for service deployment, B.Y. Zhao, L. Huang, J. Stribling, S.C. Rhea, A.D. Joseph, and J.D. Kubiatoiwicz. IEEE Journal on Selected Areas in Communications, To appear, 2003. URL http://www.cs.berkeley.edu/~ravenben/publications/pdf/tapestry_jsac.pdf.
- [11] Bayeux: An architecture for scalable and fault-tolerant wide-area data dissemination, S.Q. Zhuang, B.Y. Zhao, A.D. Joseph, R.H. Katz, and J.D. Kubiatoiwicz In Proceedings of the 11th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV), Port Jefferson, NY, June 2001.
- [12] (Pastry) Topology-aware routing in structured peer-to-peer overlay networks M. Castro, P. Druschel, Y.C. Hu, and A. Rowstron.. Technical Report MSR-TR-2002-82, Microsoft Research, 2002. <http://www.ovmj.org/GNUnet/papers/location.pdf>
- [13] Scribe: A large-scale and decentralized application-level multicast infrastructure M. Castro, P. Druschel, A.M. Kermarrec, and A. Rowstron.. IEEE Journal on Selected Areas in Communications, 20(8), October 2002.
- [14] An evaluation of scalable application-level multicast built using peer-to-peer overlays, M. Castro, M.B. Jones, A.M. Kermarrec, A. Rowstron, M. Theimer, H. Wang, and A. Wolman. IEEE Infocom'03, April 2003.
- [15] Mastering agreement problems in distributed systems, M. Raynal and M. Singal, IEEE Software July/August 2001

Network Centric Warfare Concepts in the Royal Netherlands Army C2 Architecture

- [16] Consideration of receiver Interest for IP Multicast Delivery, B.N. Levine, J. Crowcroft, C. Diot, J.J. Garcia-Luna-Aceves, J.F. Kurose
- [17] Data Replication in Low Bandwidth Military Environments – State of the Art Review, A. Gibb and S. Chamberlain, TTCP, C3I Group Technical Panel 10
- [18] Security Issues and Requirements for Internet-Scale Publish-Subscribe Systems, C. Wang, A. Carzaniga, D. Evans, A.L. Wolf
- [19] TCP-Friendly Multicast Congestion Control (TFMCC):Protocol Specification, <http://www.ietf.org/internet-drafts/draft-ietf-rmt-bb-tfmcc-02.txt>, work in progress
- [20] NACK-Oriented Reliable Multicast Protocol (NORM), <http://www.ietf.org/internet-drafts/draft-ietf-rmt-pi-norm-07.txt>, work in progress
- [21] PGM Reliable Transport Protocol Specification, <http://www.ietf.org/rfc/rfc3208.txt?number=3208>
- [22] PGMCC single rate multicast congestion control: Protocol Specification <http://www.ietf.org/proceedings/01aug/I-D/draft-ietf-rmt-bb-pgmcc-00.txt>, work in progress
- [23] DARPA's Command Post of the Future Program (CPOF), Laurie B. WAISEL