Chapter 1 – INTRODUCTION

Task Group 12 on 'Information Management over Disadvantaged Grids' was formed in January 2001 under the auspices of the Information Systems Technology Panel of the NATO Research and Technology Organisation. The Task Group consisted of four countries: Canada, Germany, Poland and United States, with the Chairman being provided by Canada. The objective of the Task Group was the following:

> Investigation of adaptive information management schemes, implemented in the nodes of tactical command and control systems, to mitigate the effects of low bandwidth, variable throughput, unreliable connectivity and energy-constrained nodes imposed by the mobile wireless communications grid that links the command and control nodes.

The Task Group limited the scope of its study to the tactical wireless domain for a Land Force operating in a "national" context (i.e., issues related to multinational coalition interoperability were not addressed).

As part of its programme of work, the Task Group undertook to organize three workshops, each addressing a key issue in information management at a different level of C2IS architecture (application, middleware, communications network). The first workshop was held at DRDC – Valcartier in Canada in September 2002 on the topic of 'Data Replication over Disadvantaged Tactical Communication Links'. The second workshop was held at FGAN/FKIE in Germany in August 2003 on the topic of 'Role of Middleware in Systems Functioning over Mobile Wireless Networks'. The third workshop was held at Naval Research Laboratory in the US in June 2004 on the topic of 'Cross-Layer Issues in the Design of Tactical Mobile Ad Hoc Wireless Networks: Integration of Communication and Networking Functions to Support Optimal Information Management'.

This report discusses the first workshop in the series, which took place September 11-12, 2002 in Québec City, Canada, and was hosted by the System of Systems Section of DRDC Valcartier. The technical programme of this workshop is provided at Annex A.

The objective of the workshop on 'Data Replication over Disadvantaged Tactical Communication Links' was to address the problem of data replication among distributed databases occurring over disadvantaged (unreliable, low bandwidth or energy-constrained) mobile wireless military communication networks. This document provides a technical evaluation of the workshop by the Task Group. Section 1.1 discusses the operational military requirement for data replication in the tactical wireless domain. Chapter 2 describes the military command and control and communications environment in which data must be replicated. Chapter 3 discusses different data replication models and identifies the model that is most appropriate for the tactical wireless domain. Chapter 4 discusses desirable characteristics of a data replication service for the tactical wireless domain. Chapter 5 provides a summary and conclusions.

1.1 BACKGROUND

Digitization of Command and Control information seeks to facilitate increased operational tempo by reducing the need to slow down, or stop, in order to regain situational understanding. In order to realize this operational objective, information must continue to flow as the forces move throughout the battlespace. Since commanders may need to receive and disseminate situational awareness and execution information from any location in the area of operations, all of the sources and destinations of information should remain available to the commander at all times. The objective of managed information exchange is to support the commander's ability to execute command and control by providing a timely flow of accurate, relevant information.

Computers assist human beings by maintaining and manipulating large amounts of data that represent meaningful information to the human beings that use the computers. For information to be managed and manipulated by a computer, however, it must be highly structured. This structure typically takes the form of a structured data schema implemented in a database. The structure is important to permit the computer to deal with the data. However, the structure also serves to preserve a context for the data. Preservation of context (structure) for data is critical as the data are manipulated and exchanged.

The operational benefits of digital information technology employed in support of military command and control are derived mainly from the speed with which computers can access, retrieve, process and display structured data that have been entered into them. However, the operational benefits are also highly dependent upon the ability to exchange structured data of operational importance between computers with the speed necessary to support operational tempo.

The term 'replication' refers to the systematic propagation and maintenance of copies of data between datastores within a distributed computing environment. The field of replication has two domains; file replication and data replication. Most commercial products address one of these two but not both. File replication is concerned with the copying of complete files while data replication operates at the level of database transactions. This report deals with the domain of data replication.

In database systems, a transaction is a set of database operations that must succeed or fail as a single unit. A transaction can be considered a logical unit of work that transforms the database from one consistent state to another consistent state. Data replication middleware is used to copy database transactions from a source database to multiple replicate databases. In the military context of the present workshop, the databases in question are contained in computers located in stationary or mobile command headquarters, in military vehicles, or on dismounted soldiers. From the communications perspective, each of these entities is considered a node on a military command and control network. The following section presents an overview of the tactical Army command, control and communications environment.

Chapter 2 – ARMY TACTICAL COMMAND, CONTROL AND COMMUNICATIONS ENVIRONMENT

2.1 MILITARY COMMAND AND CONTROL SYSTEM STRUCTURE

Command and control systems must support three types of relationships: command, support and proximity. The following extract is taken from [1]:

> "Command relationships exist whenever one unit or formation commander is a direct subordinate of another…Command requires a rich bi-directional exchange of information between the higher headquarters and the subordinate headquarters. The purpose of this exchange is to pass command information (plans, orders, task organization, battlefield geometry, alerts, warnings and status) between the two headquarters. This type of information exchange follows the parent-child relationship. The superior (parent) headquarters supplies directive information to the subordinate (child) headquarters – higher to lower; the subordinate (child) headquarters provides status information to the superior (parent) headquarters – lower to higher.
>
> Support relationships are a particular type of command relationship. Support relationships are established when one organization must aid, protect, complement or sustain another force. In the context of Command and Control (but not, for example, fire control) these organizations have the same requirements for information exchange as command relationship. Support relationships are of two types: Direct Support and General Support.
>
> Proximity relationships exist when units with no direct command or support relationship are operating in proximity to each other and must exchange non-command information in order to establish and maintain situational awareness. Examples of this type of relationship could be the flank coordination of adjacent tank and infantry battalions or the forward passage of lines of an armoured regiment through a mechanized infantry battle group. In proximity relationships, information flows horizontally between the headquarters of the units involved as peers, not parents/children. Units involved in proximity relationships may be subordinate to different higher headquarters."

The military command and control structure is hierarchical (Figure 1). A headquarters at a certain command level will be parent of one or more subordinate headquarters, and will itself be subordinate (child) to a higher headquarters. The communication infrastructure and flow of information over that infrastructure reflects this command hierarchy. A commander at a given level will generally be required to maintain information from one level up and two levels down in the command hierarchy, as well as from flanking formations with which he has a proximity relationship.

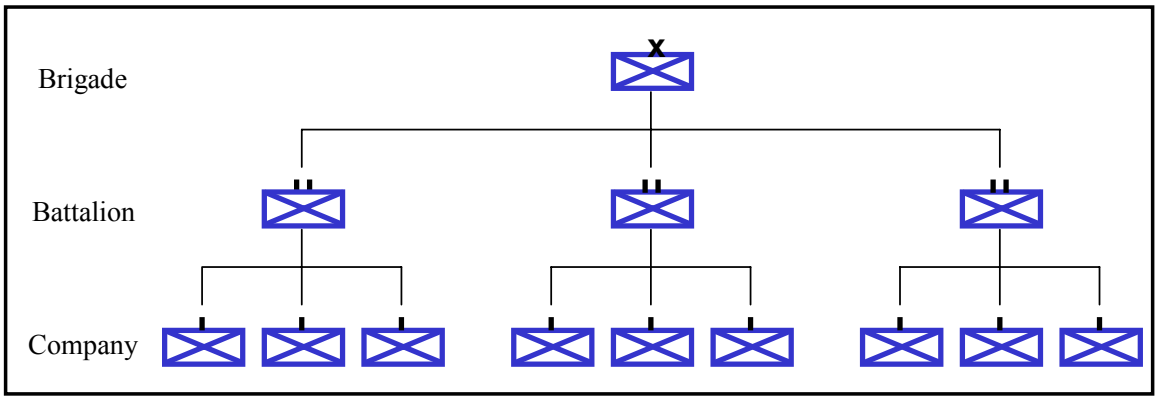

Figure 1: A Typical Hierarchy of Command Headquarters.

2.2 COMMAND AND CONTROL COMMUNICATIONS INFRASTRUCTURE

The command and control communications infrastructure is generally organized as a series of hierarchical subnetworks. A brigade command subnet will include a node at each of the three battalion headquarters under the brigade's command. A battalion command subnet may have a dozen or more nodes, but will include a node at each of the three company command posts under the battalion's command. Links from battalion headquarters to brigade or higher headquarters are typically provided by relatively reliable and high bandwidth satellite links or dial trunk systems employing wire or wideband microwave links. The workshop focused its attention on the Army battlefield environment forward of a battle group or battalion[1] headquarters in which all communication occurs between mobile nodes equipped with a combat net radio operating in either the Very High Frequency (VHF) or Ultra High Frequency (UHF) band. A radio subnetwork consists of a set of radios tuned to a common assigned frequency. The subnet is linked to an adjacent subnet through a relay or gateway node that is common to both subnets. The relay node contains two or more radios, each tuned to a different subnet frequency. A verbal message received at the relay node from a sending node on one subnet is recorded in writing by a human operator, and then re-transmitted verbally on the appropriate subnet to reach the destination addressee(s) on that subnet, if required. A gateway performs the same function for a data transmission, except that caching of the received transmission, and re-transmission on the target subnet(s), are handled automatically by the gateway.

2.3 COMBAT NET RADIO COMMUNICATIONS ENVIRONMENT

Johnson [2] provided an overview of the tactical Army communications environment. Forward of battalion, communication occurs over line-of-sight (LOS) radios operating in the Very High Frequency (VHF) or Ultra High Frequency (UHF) bands. Most of these radios are vehicle-mounted, but man portable versions are also employed. At the lowest echelon, such as an infantry section on foot engaged in urban warfare, soldiers may use short-range radios operating in the High Band UHF. These radios have ranges of 400 meters or less. UHF radios are used for medium-range LOS wireless communication from 400 meters to 15 kilometres. VHF Radios must be used where non-LOS communication beyond 15 kilometres is required.

[1] Battalions are either mechanized battalions (two armoured companies, one infantry company) or infantry battalions (two infantry companies, one armoured company). A battle group is a battalion augmented with other assets (usually engineer or artillery assets).

The vast majority of VHF combat net radios have base bandwidths of 16 kilobits per second (Kbps) half duplex[2]. UHF radios can have base bandwidths as high as 288 Kbps full duplex. High Bandwidth UHF radios can have base bandwidths as high as 11 Mbps. In spite of their limited bandwidth, combat net radios operating in the VHF band still predominate because of their greater range and beyond-LOS capability.

The values quoted above apply to the physical and data link layers of the network, the two lowest layers of the seven layer OSI model (Figure 2). The reality is that the useable throughput at the application layer is a fraction of this base rate. Factors such as forward error correction (FEC), encryption overhead (e.g., crypto synchronization sequence), acknowledgement request and retransmission mechanisms at data link and transport layer, and media access control mechanisms, are responsible for this reduction. Moreover, the effective throughput at any time may fall well below this maximum value due to variations in the performance of the physical radio channel caused by real-world factors such as terrain interference, atmospheric interference, multi-path (reflections) and prolonged fading. Effective throughput can become zero for periods of time for certain links, or for the entire subnet (in the case of imposed radio silence).

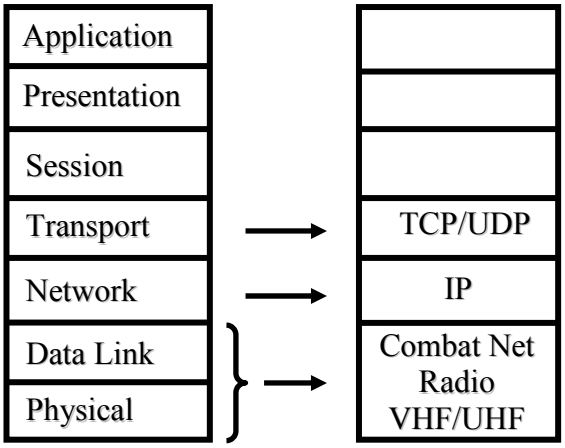

Figure 2: Seven-Layer ISO Network Reference Model.

A VHF combat net radio with a base rate of 16 Kbps may have a maximum effective throughput as low as 1 Kbps due to the factors quoted above. If ten users share the radio subnet, the maximum effective throughput per user will be only 100 bps. This figure assumes a data-only network with no voice contention for the channel. The reality is that the residual bit-error rate (BER) (i.e., the BER after error correction coding) for these radio channels can be as high as 10^{-5}.

One known UHF combat net radio can have a base rate as high as 288 Kbps. However, UHF radios have shorter ranges than VHF radios since the higher frequency UHF waves are more susceptible to the real-world factors quoted above. For the case considered, field measurements suggest that the maximum effective throughput at the top of the data link layer in a tactical environment would be approximately 16 Kbps for a link with a 22 Kbps base rate and approximately 80 Kbps for a link with a 100 Kbps base rate. For a subnet with ten users, maximum effective throughput per user would be 1.6 Kbps and 8.0 Kbps respectively. Residual BER are similar to rates experienced with VHF radios.

[2] Half duplex means that the radio can either send or receive on the same channel, but cannot simultaneously send and receive. Full duplex means that the radio can simultaneously send and receive on the same channel.

In the Internet world, the transport protocol of choice for most applications is the Transport Control Protocol (TCP). TCP provides reliable packet delivery using a sequencing and positive acknowledgement scheme and is a connection-oriented protocol as it establishes a connection with each recipient. Although point-to-multipoint addressing schemes are being developed for TCP, in the radio domain the vast majority of systems still use a point-to-point addressing scheme for TCP. Therefore, in this domain, if a transmission is intended for N recipients, it must be sent N times. In a highly bandwidth-constrained wireless environment, the communication overhead associated with the use of a connection-oriented transport protocol like TCP is generally unacceptable. As well, the congestion-control mechanisms used by TCP were developed for the wired domain. TCP interprets latency and packet loss as evidence of congestion, to which its reaction is to throttle back the offered load to the network. This results in a significantly lower TCP throughput than the network can actually provide. In low bandwidth, high latency, and relatively high error networks like those found in the tactical radio domain, many TCP connections would be terminated due to these effects. User Datagram Protocol (UDP) is a connectionless alternative to TCP. However, UDP is unreliable since it does not employ sequencing or acknowledgement. Custom middleware operating just above UDP may be required to optimise reliability of packet transmissions in a connectionless, high BER tactical wireless environment. The workshop labelled this custom middleware the Replication Transport Layer (RTL). The characteristics required of the Replication Transport Layer are discussed in Section 4.6.

Chapter 3 – DATA REPLICATION IN A BANDWIDTH-CONSTRAINED WIRELESS ENVIRONMENT

3.1 VOICE VERSUS DATA NETWORKS

Tactical radio networks forward of battalion have, for decades, consisted of voice-only radios. As part of the revolution in military affairs and the move toward network-centric warfare, armies are basing their new command and control systems on digital technology. As a consequence, data traffic is increasingly replacing voice traffic. Forward of battalion, networks based on mixed-mode voice/data radios or data-only radios are increasingly replacing voice-only radio networks.

3.2 INFORMATION EXCHANGE MODELS FOR A DIGITAL ENVIRONMENT

Over voice radio, the standard means of communication for decades has been verbal transmission of structured military messages (example: "0 this is 6, Contact Grid 247653, six enemy tanks using routes into the built up area and concentrating in that location oriented South. Continuing to observe. Time of contact 1430 hours. Over."). In the conversion from voice to digital technology, one option is to preserve the structured message paradigm in digital form. It has the advantage that it is a paradigm with which armies are already familiar. However, in a bandwidth-constrained data communication environment, use of structured military messages has two important disadvantages. First, a structured message must, by definition, be semantically complete (i.e., in a form sufficiently complete to be fully understood by the recipient without resort to external references). Second, the message structure itself creates a certain amount of overhead (header, field separators, etc.). The need for semantic completeness can create a tendency on the part of users to recall a previously sent message, modify the fields whose contents have changed, and resend the entire message (including fields whose values have not changed). Sending of redundant information and the message structure itself combine to create a significant communication overhead that can limit the throughput of operationally important new data in a bandwidth-challenged radio environment.

An alternative to the use of structured messages is to perform information exchange via direct database-to-database exchange. Using this approach, one or more database transactions modify attributes in relational tables in a local database using some combination of insert, update and delete operations. The local database transactions are packaged for transmission and copied using data replication middleware to recipient nodes. At the receiving end, the database transactions are unpackaged and applied directly to the databases in the recipient nodes. Such database to database exchanges can be more bandwidth-efficient than structured message exchanges because they can have minimal overhead and transmit only data values that have changed.

3.3 SYNCHRONOUS VERSUS ASYNCHRONOUS REPLICATION

Data replication middleware is used to copy database transactions from a source database to multiple replicate databases. In a commercial application (e.g., a financial system), the transactions are applied synchronously (i.e., simultaneously), using a protocol known as two-phase commit. The protocol ensures that either all participating databases commit the transaction at the same instant, or none do. Use of such a protocol produces what is termed 'tight' consistency of database content. Tight consistency means that data values are consistent (i.e., the same) across source and replicate databases at all times.

The two-phase commit protocol is a two-step protocol in which each replicate database sends a positive acknowledgement back to the source database for each step. In a bandwidth-constrained environment, use of

such a protocol generates such a significant communication overhead that it can seriously impede throughput of operational data. In this case, asynchronous replication must be employed. Under asynchronous replication the transaction(s) are applied to the source database before being replicated. Replicated transactions are applied to the replicate databases sometime after they are applied to the source database. Use of asynchronous replication produces what is termed 'loose' consistency of database content[3]. Loose consistency means that modified data values always become consistent across source and replicate databases, but only after a certain time delay (latency).

In the very low bandwidth regime associated with combat net radio, where the data throughput of the radio network can be chronically less than the load offered to the network, it is necessary to relax traditional consistency expectations. In this situation, at any point in time, a percentage of data values in the source and replicate databases will be inconsistent, producing a condition that has been dubbed 'lazy consistency'. Since one can achieve neither tight nor loose consistency of full database content in this situation, the challenge shifts to preservation of consistency for those attributes judged to be of highest operational importance. This problem is inherently more challenging than the problem of achieving tight or loose consistency. In the latter two cases, the end state is inherently stable, i.e., the database content will always evolve toward a state of consistency, albeit with some delay. In the case of chronically inadequate and variable bandwidth, the end state is uncertain and unstable. Data delivery is on a 'best effort' basis and there is no assured 'audit trail' for the database changes that are propagated. A continual deterioration in the consistency of database content is likely, unless steps are taken to manage information flow in a proactive manner.

In this low bandwidth environment, the data replication and transport mechanisms must include intelligent information management protocols capable of adapting to the variable throughput of the radio network and to changing battlefield priorities.

[3] For this to be true, it is essential that the asynchronous replication mechanism preserve transactional semantics. Database-to-database replication of transaction events as they occur in the source database generally preserves transactional semantics whereas a complete or incremental periodic refresh (snapshot) mechanism does not.

Chapter 4 – DESIRABLE CHARACTERISTICS OF THE DATA REPLICATION SERVICE

In his keynote address, Chamberlain [3] discussed characteristics that a data replication service should have to function well in the tactical radio domain. He noted that the variation in bandwidth is as important as the actual bandwidth values in determining data throughput. Zero is a valid value for throughput in this domain. A command and control node must continue to function even when the data throughput equals zero. He proposed the use of local predictive algorithms to estimate changes in data values (e.g., vehicle position) during these periods.

To be useful, the information exchange service must support rapid selection and dissemination of information and hands-off operation. According to Chamberlain, the service should have three characteristics:

1) Automatic (hands-off, context-sensitive);
2) Adaptive (context-based, responsive to changing bandwidth); and
3) Affordable (treated as part of a Distributed Computing Environment).

The goal is to balance or 'tune' information management to available network resources.

Chamberlain's proposal for what he terms 'Model-Based Battle Command' is based upon the following four principles:

1) Use data abstractions as the medium of exchange (data replication).
2) Control exchange by active database triggers (provides context-sensitivity since triggers link the decision to replicate to the state (i.e., value) of certain database elements).
3) Allow reasonably different perceptions of battlefield (accept 'lazy consistency' of database content).
4) Treat database synchronization as the realistic control of differing perceptions (management of 'lazy consistency' rather than enforcement of loose/tight consistency of database content).

4.1 NETWORK AWARENESS

A key element in the above approach is the ability to monitor or measure bandwidth so that synchronization efforts can be adapted to current communication resources. To achieve this objective, according to Chamberlain a communications model, incorporating information such as network performance and connectivity data, should be part of the data schema. As well, network performance statistics should be passively collected (e.g., channel access delay, round trip time, queue delay, number of failed versus successful transmissions), analyzed, and results made available to the application through the communications model in the data schema. Chamberlain also proposes that innovative protocol mechanisms be considered to permit more direct performance feedback to the application and more efficient data transmission. Possible protocol mechanisms include: stack cognizance (sharing of information between ISO layers, e.g., transport and data link layer), 'just-in-time' packet construction (packing several small (highest-priority) application Protocol Data Units (PDU) into a Maximum Transmission Unit (MTU)), and overhearing (the direct passage between the network and transport layer to the host of packets rather than addresses).

4.2 'ALL-INFORMED' DATA DISTRIBUTION MODEL

St-Jacques [4] described a Canadian prototype of a digital command and control system for a close-support artillery regiment, Artillery Regimental Data System Advanced Development Model (ARDS ADM). This prototype system employed a model-based approach (database to database exchange) and an asynchronous data replication mechanism, based upon triggers and stored procedures and a negative acknowledgement scheme, designed for the combat net radio environment, to maximize the consistency of database content across databases on the same subnet ('all-informed' model). The presentation identified the operational advantages of an 'all-informed' data distribution model for the tactical domain. They were:

a) Eliminates single point of failure;

b) Permits data recovery from any node;

c) Permits a node to assume role of neighbour node without substantial one-time data transfer; and

d) Takes maximum advantage of shared medium (radio).

4.3 DATA OWNERSHIP

The ARDS ADM project explored three models for data ownership: single, dynamic (i.e., transferable) and shared. 'Ownership' of a given data element refers to the right to modify the value of that data element. In a single ownership model, the originator of the data retains ownership throughout the lifetime of the data element. In a dynamic ownership model, any given data element has a single owner at any point in time, but the ownership is transferable; the present owner of the data element may not be the originator of that data element. In a shared ownership model, there is no single owner for the data element. Any participating database may initiate a change in the value of the data element. The project concluded that a shared ownership model should be avoided because it maximizes the possibility of data conflicts, and resolution of data conflicts can generate an undesirable increase in network traffic. Furthermore, it is important that the system track the identity of the owner of the data element and the ownership model that applies to that data element.

4.4 DATA RECOVERY

The ARDS ADM project also identified the need for data recovery to be intelligently managed. At the packet level, this means tying the number of attempted retransmissions to the assigned transmission priority for the packet and to the known current performance of the network. At the level of bulk recovery (e.g., by a node that has been disconnected from the network for an extended period), this means minimizing the size of the recovery package by being selective as to content (for example, by ignoring time-sensitive data whose operational value ages quickly). The latter approach will guarantee that the content of the recovering database will never be fully consistent with that of the other databases. However, a bulk recovery may tie up the radio sub-net for minutes or tens of minutes. By limiting the content of the recovery package, one is trading off consistency of stale data in one database in favour of consistency of fresh data in all databases. Such a trade-off is an example of the 'realistic control of differing perceptions' in this communications environment referred to by Chamberlain.

4.5 SELECTIVE DATA DISTRIBUTION MODEL

The all-informed data distribution model referred to in Section 4.2 is based on the assumption that there is value in maintaining synchronized data content across all participating nodes on a subnetwork, at least for an

agreed subset of entities within the data model. However, an alternative model, selective data distribution, makes the opposite assumption, i.e., that only data of interest to a particular node should be sent to that node. The underlying assumption in such a data distribution strategy is that a node should only receive data from external sources that are important to its assigned role (or that are permitted by a security policy). One of the most important examples of this approach, the ATCCIS Replication Mechanism, is described in Section 4.7 and Annex B. A comparison of the advantages and disadvantages for the tactical wireless domain of a selective data distribution model versus an all-informed data distribution model can be found in Section 4.8.

4.6 FUNCTIONAL REQUIREMENTS – DATA REPLICATION SERVICE

One of the discussion groups at the workshop addressed the question:

> "What key concepts and design principles should drive the design of data replication/transport mechanisms to function optimally over disadvantaged tactical communication links?"

The discussion group responded to the question by identifying a list of key functional requirements that should be satisfied by the replication mechanism and transport mechanism. This section presents the results of that analysis.

The Replication Mechanism (RM) is responsible to determine when replication of data from a local node to other network nodes is to occur and what information is to be replicated. It also packages and unpackages data to be replicated into/from a unit that the group called a Replication Protocol Data Unit (PDU)[4]. A PDU is information that is delivered as a unit among peer entities of a network that may contain control information, address information, or data. The Replication Transport Layer (RTL) sits beneath the Replication Mechanism and provides a transport mechanism for the purposes of passing replication data, in the form of Replication PDUs, over the communications bearer.

The Combat Net Radio environment, when compared to wired networks, is characterized by:

- Low bandwidth (e.g., 1 Kbps – 80 Kbps);
- High and variable latency (e.g., 0.5 sec); and
- Intermittent connectivity.

For optimal performance in the tactical wireless domain, the communication protocol in the Replication Transport Layer needs to be matched to both the characteristics of the communications bearer and the information exchange requirements of the Replication Mechanism.

The discussion group identified a list of key functional requirements that should be satisfied by the replication mechanism and transport mechanism. The list was organized into three parts, namely:

1) Functions to be implemented in the Replication Mechanism;
2) Functions to be implemented through some combination of the Replication Mechanism, the Replication Transport Layer, and, perhaps, an additional application; and
3) Functions to be implemented in the Replication Transport Layer.

[4] In some analyses, the Replication Mechanism is further decomposed into a Replication Agent (the component that determines when replication is to occur and what information is to be replicated) and a Replication Server (the component that packages and unpackages data to be replicated into/from Replication PDUs).

4.6.1 Functional Requirements

A) Replication Mechanism:

- Intelligence to determine when replication is to occur. This function must be context sensitive.
- Intelligence to determine what data is to be replicated once a decision to replicate has been taken.
- Assembly of the replication protocol data unit (PDU). If information is available on what replications a sender or senders already possess, it is possible to determine if the PDU is semantically complete.
- Structure of the trigger and stored procedure languages should support implementation and application of the intelligence to individual data elements or to data aggregates within the local databases, and support assembly of the replication PDUs.

B) Combination of Replication Mechanism and Replication Transport Layer or Other Application:

- Functionality must be provided to enforce and mediate dependencies of RM and RTL on other system architecture components. Examples: dependence of RM on characteristics of a particular DBMS or other application, dependence of RTL on characteristics of a particular Data Model, RM, etc.
- Data Ownership: RM and RTL must be able to identify the owner of any given piece of data at any given time. This is tied to the policy for database key management and requires decisions about the implementation level and authority structure for such management. These decisions can depend upon how the data ownership information is being used (e.g., enforcing transactional integrity, resolving data conflicts, providing traceability of ownership over time).
- Intelligence to determine the level of effort devoted to transmitting a Replication PDU across the network based on the importance of the information contained in the PDU. For example: (1) how many retransmission attempts should be made by the RTL before it stops trying to transmit the PDU; and (2) if the RTL has more than one class of transport service (e.g., 'guaranteed' and 'best effort' transport services), which one should be used for each Replication PDU? Assuming the RTL has several mechanisms to distinguish classes of service, one of the significant issues is defining the criteria for assigning a level of 'importance' to each PDU.

C) Replication Transport Layer:

- RTL must support an Acknowledgement Scheme (negative acknowledgement preferred; positive acknowledgement scheme is very bandwidth-intensive).
- RTL must supply a retransmission protocol that takes account of time-varying communications bearer performance.
- Priority Scheme: RTL must support a prioritization scheme at the PDU level that takes account of time-varying bearer performance.
- RTL must enforce a degree of fault tolerance and be sensitive to time-varying communications bearer performance.
- RTL must support fragmentation and defragmentation of PDUs, as appropriate, determined by the characteristics of the network.
- Addressing Scheme: RTL should support an addressing scheme. Choices are a broadcast communications protocol or an addressing scheme that would allow point-to-point or point-to-multipoint communications to occur between RTL peers.

- Cognition of Network Structure: RTL should be cognizant of the network structure (who is on the radio net and the status of communications transmitted from each node). This information, if available, could be used to determine what information transmitted by a remote node has not been received by the local node and vice-versa. This capability would, however, introduce overhead to the RTL communications protocol.

4.6.2 Non-Functional Requirements

The RM and RTL must operate effectively in both the non-constrained and constrained communication environments.

4.7 ATCCIS REPLICATION MECHANISM (ARM)

One of the most important examples of selective data distribution in the NATO context is the ATCCIS Replication Mechanism [5]. The present section contains a brief summary of the ATCCIS replication mechanism (a fuller description can be found in Annex B). Some advantages and disadvantages of the selective data distribution model and all-informed data distribution model for the tactical wireless domain are discussed in Section 4.8.

ATCCIS (Army Tactical Command and Control Information System) was an international programme consisting of NATO nations (although not a formal NATO effort) aimed at identifying the minimum set of specifications to be included within C2ISs to allow the automatic transfer of selected command and control (C2) data. Their objective was to develop a specification for a hardware/software/vendor-independent interoperability solution. The ATCCIS programme ran from 1982 to 2002.

The Multinational Interoperability Programme is an international programme consisting of NATO nations (also not a NATO effort) whose focus is the fielding of an interoperability solution for multinational C2ISs. In 2002, ATCCIS merged with MIP. MIP adopted the products of the ATCCIS work as the basis for direct database-to-database exchange. However, MIP also maintains a structured message exchange mechanism.

The ATCCIS concept of interoperability is based upon the automatic transfer of standardized data elements that utilize agreed and common data identifiers based upon a common data interchange model. The common data model is called the C2 Information Exchange Data Model[5].

Additionally, the ATCCIS programme developed the specification for a mechanism that will permit interoperability of automated C2ISs through the partial replication of database content. The ATCCIS Replication mechanism (ARM) is selective in (a) the data to be exchanged, (b) the recipients of the data and (c) the transfer facility to be used.

Under the ATCCIS concept, nations use the common data model to preserve the meaning and relationships of the information exchanged between C2IS across national boundaries. Nations are free to develop differently structured C2 databases for national use. The ARM is used to manage the exchange of information between databases of C2ISs across national boundaries based on the common data model.

The ARM employs the concept of contracts and filters. A *replication contract* is the means for controlling (selective) replication of database changes. A contract is established between a pair of replication nodes,

[5] The Land Command and Control Information Exchange Data Model (LC2IEDM) developed under ATCCIS has been extended under MIP to serve Joint Force requirements. The extended data model is referred to as 'C2IEDM', without the qualifier 'Land'.

designated as Data Provider (DP) and Data Receiver (DR). In the contract, the DP and DR agree that the DP will provide the DR with all data that satisfies the conditions of the contract. A contract specifies a *filter* and parameter values used to set filter conditions, as well as a DP and a DR. A filter is a set of criteria applied to the instances of a database in order to reduce the total set of data selected to a subset. Examples of filter types include geographical area, time, and order of battle (organizational). The contracts enforce a 'push' model in which the only data pushed to recipient nodes are those negotiated with the recipient node under the pre-agreed contract. To modify the set of data pushed to a particular data recipient by a data provider, a filter must be applied or the contract must be modified.

4.8 SELECTIVE DATA DISTRIBUTION VERSUS ALL-INFORMED DATA DISTRIBUTION MODEL

As noted in Section 4.5, the all-informed data distribution model is based on the assumption that there is value in maintaining synchronized data content across all participating nodes on a subnetwork, at least for an agreed subset of entities within the data model. However, a selective data distribution model makes the opposite assumption, i.e., that only data of interest to a particular node should be sent to that node. The underlying assumption in such a data distribution strategy is that a node should only receive data from external sources that are important to its assigned role (or that are permitted by a security policy). Under this distribution model, database content will, by design, not be consistent across nodes on a subnet. An important example of this approach, the ATCCIS Replication Mechanism, was described in the preceding section.

A selective data distribution model ensures that a user of data is not forced to deal with data that are of no interest or importance to his role, and that, in the event of an unsuccessful retransmission, no effort is expended re-transmitting the data to nodes that have no identified need for it. As well, it limits the quantity of received data that must be stored at the receiver node. Finally, if a point-to-point addressing scheme is used, it can support a security policy based on selective dissemination since transmission of a data payload can be limited to specific receiver nodes. However, in the tactical wireless domain, there can be serious operational disadvantages associated with selective data exchange caused by the very limited bandwidth and variable connectivity of the wireless network. All the advantages of an all-informed data distribution model listed in Section 4.2 become disadvantages for a selective data distribution model. The disadvantages are:

a) Does not take advantage of shared medium (radio)

   Most of the time, nodes on a wireless data subnet overhear data transmissions that are addressed to other nodes on the subnet. When data replication is used as the exchange mechanism, the database changes contained in those transmissions can be applied to the database in the overhearing node, even though they are not addressed to, or directly pertinent to the role of, the overhearing node. An all-informed data exchange model exploits these overheard data transmissions to the maximum extent possible while a selective data exchange model routinely discards this 'free' information.

b) May introduce single point(s) of failure

   Under selective data distribution, each node is the sole provider of data which it 'owns', and subsets of its owned dataset are shared with a set of data receiver nodes. Each receiver node is subscribed, in principle, to an unique subset of the provider's dataset (the subsets for different receivers may overlap, but are, in general, not identical). Under this scheme, a given receiver node will only have knowledge of data being shared with it, not with the other receiver nodes, since the exchange agreements are between node-pairs. If the provider node is disabled, due to enemy action or other reason, for all practical purposes each receiver node will be limited to the narrow slice of the provider's

dataset received from the provider node, and the set of receiver nodes will have no simple means of recovering or reconstructing all or part of the total dataset of the provider node, if that proves necessary (certainly no mechanism efficient enough to work acceptably over a tactical radio system). Under an all-informed data distribution model, nodes on the subnet overhear and capture data transmissions, even those for which they were not explicitly identified as a data receiver. If a provider node is disabled, each node on the subnet will have a copy of the content of all (or almost all) data transmissions emanating from the provider node to that point in time recorded in its database, thereby largely avoiding the problems noted above.

c) Requires data recovery from multiple nodes

In the event that a node leaves the subnetwork for a period of time, rejoins the network, and wants to recover data that it has missed, it will have no choice but to request missing data from each data provider node with which it has an exchange contract. If there are n such provider nodes, that will require n separate data transmissions. If three nodes need to recover data at the same time, and they have exchange contracts with n_1, n_2 and n_3 provider nodes respectively, that will require $n_1 + n_2 + n_3$ distinct data transmissions. This recovery process increases the traffic level on the radio channel and wastes bandwidth. By contrast, in an all-informed data exchange model, the node(s) can recover missing information in a single data transmission from any neighbour node, since the database content is, in principle, the same across all nodes.

d) Node cannot readily assume role of neighbour node without substantial one-time data transfer from that neighbour node

If a node is required to assume the role of a neighbour node, under a selective data exchange model it would be required to download all, or a substantial portion of, the information specific to the neighbour's role from the database of the neighbour node, since it would not have information specific to the neighbour's role in its own database. If the neighbour node database is not available, due to enemy action or other reason, it would have to request recovery of those role-specific data from the n data provider nodes with which the original node had an exchange contract. This would require n distinct (potentially large) data transmissions. In an all-informed data exchange model, no such large data transmissions are required because, in principle, the replacing node has the same data in its database as the replaced node.

An all-informed distribution model can still enforce a security policy of selective dissemination. Tactical radio transmissions are routinely encrypted for transmission and then decrypted at the receiving end. Thus the received decrypted transmission can be assumed to have originated from an authentic source. Selective dissemination can be enforced at the software level by implementing control software in the receiving nodes to ignore a (decrypted) data transmission if a security caveat in the received transmission instructs it to do so. With this model, data will be shared on an all-informed basis across all nodes on a subnet except for instances when the security policy specifically excludes this option.

If a point-to-multipoint addressing scheme is employed, it is possible to combine an all-informed distribution model for certain types of data (e.g., reports of friendly vehicle positions) with a selective distribution model for other types of data (e.g., passage of a fragmentary order from battalion HQ to company HQs). This hybrid distribution model will permit full synchronization of database content for certain types of data, and partial synchronization of database content for other types of data. Such a hybrid model may well provide the best match between operational requirements and bandwidth utilization in the tactical wireless domain.

Chapter 5 – SUMMARY AND CONCLUSIONS

Driven by the revolution in military affairs and the move toward network-centric warfare, armies are basing their new command and control systems on digital technology. The operational benefits of digital information technology employed in support of military command and control are derived mainly from the speed with which computers can access, retrieve, process and display structured data that has been entered into them. The operational benefits are also highly dependent upon the ability to exchange structured data of operational importance between computers with the speed necessary to support operational tempo. In an operational military environment, two forms of data exchange are employed: structured message exchange and exchange of database transactions. The latter form can be more bandwidth-efficient because it propagates data changes only. It is implemented via a data replication mechanism.

The term 'replication' refers to the systematic propagation and maintenance of copies of data between datastores within a distributed computing environment. Data replication middleware is used to copy database transactions from a source database to multiple replicate databases. In the military context studied in this workshop, the databases in question are contained in computers located in stationary or mobile command headquarters, in military vehicles, or on dismounted soldiers, all operating forward of battalion. From the communications perspective, each of these entities is considered a node on a military command and control network. The nodes communicate via a combat net radio, operating in either the VHF or UHF band. The combat net radio environment, when compared to wired networks, is characterized by:

- Low bandwidth (e.g., 1 Kbps – 80 Kbps);
- High and variable latency (e.g., 0.5 sec); and
- Intermittent connectivity.

A data replication service can be considered to consist of two basic components, a Replication Mechanism and a Replication Transport Layer. The Replication Mechanism is responsible to determine when replication of data from a local node to other network nodes is to occur and what information is to be replicated. It also packages data to be replicated and unpackages data that has been replicated. The Replication Transport Layer sits beneath the replication mechanism and provides a transport mechanism for the purposes of passing replication data over the communications bearer.

To be effective in the low bandwidth tactical wireless environment, the data replication and transport mechanisms must include intelligent information management protocols capable of adapting to the variable throughput of the radio network and to changing battlefield priorities. This adaptivity requires that the mechanisms must have dynamic awareness of the network state and battlefield state. For optimal performance in the tactical wireless domain, the communication protocol in the Replication Transport Layer needs to be matched to both the characteristics of the communications bearer and the information exchange requirements of the replication mechanism.

Two data distribution models were identified at the workshop. The all-informed data distribution model is based on the assumption that there is value in maintaining synchronized data content across all participating nodes on a subnetwork, for an agreed subset of entities within the data model. A selective data distribution model makes the opposite assumption, i.e., that only data of interest to, or relevant to, a particular node should be sent to that node. The ATCCIS Replication Mechanism is an example of a mechanism based on selective data distribution. All-informed distribution takes maximum advantage of the shared radio medium and may have some important operational advantages in the bandwidth-constrained tactical radio environment.

Chapter 6 – REFERENCES

(Task Group workshop presentations referred to below can be accessed through workshop programme in Annex A.)

[1] Multilateral Interoperability Programme Tactical C2IS Interoperability Requirement, MIP Operational Working Group document MTIR-OWG, draft for Edition 2.0, iteration 3, dated 27 September 2003, p. 17.

[2] Johnson, T., *The Harsh Reality of a Harsh Communications Environment*, NATO IST-030/RTG-012 Workshop on 'Data Replication over Disadvantaged Tactical Communication Links', Québec City, Canada, 11-12 September 2002.

[3] Chamberlain, S., *Design Concepts for Resilient Database Replication in Tenuous Communication Environments*, NATO IST-030/RTG-012 Workshop on 'Data Replication over Disadvantaged Tactical Communication Links', Québec City, Canada, 11-12 September 2002.

[4] St-Jacques, J.-C., *Artillery Regimental Data System Advanced Development Model – Replication Issues*, NATO IST-030/RTG-012 Workshop on 'Data Replication over Disadvantaged Tactical Communication Links', Québec City, Canada, 11-12 September 2002.

[5] Angel, P., *ATCCIS Replication Mechanism (ARM) Fundamental Concepts*, NATO IST-030/RTG-012 Workshop on 'Data Replication over Disadvantaged Tactical Communication Links', Québec City, Canada, 11-12 September 2002.