Chapter 7 – CANADIAN EXPERIMENTS USING LOW BANDWIDTH TEST BED

This chapter presents some results of simulations conducted under a technology demonstration project of Defence R&D Canada entitled 'High Capacity Tactical Communications Networks' using a simulation environment called the Low Bandwidth Test Bed (LBTB) [17]. The simulations were conducted to examine the impact of application-layer information management techniques on the throughput of operationally important information over a congested tactical radio subnet. The results illustrate the positive impact that some techniques discussed in Chapters 4 and 5 of this report can have on information flow over disadvantaged tactical communication grids.

7.1 DESCRIPTION OF LOW BANDWIDTH TEST BED

The Low Bandwidth Test Bed is a discrete event simulator that couples a performance model of a tactical combat net radio system to simulated tactical nodes (STNs) containing real databases. It supports two modes of operation: distributed and integrated mode.

In distributed mode, up to 32 computers can be employed. One of the computers acts as the controller for the simulation while a number of STNs can be defined on each of the remaining computers. There is no limit to the number of STNs per machine. In integrated mode, the entire simulation environment (control element and all defined STNs) resides on, and is executed on, a single machine. This option permits execution of multiple simulations simultaneously, but each simulation executes more slowly.

To provide the capability to predefine the identity of the participating nodes, as well as the timing and nature of the communication events emanating from each node in a master communication event script, a Scenario Script Management Tool was integrated into the test bed. At the time of simulation initialization, this script is propagated to each simulated node. When the simulation begins, the scripts are executed simultaneously on each node under the direction of the control node, which advances scenario time in a synchronized fashion at one-second intervals. Each communication event in the script mimics the creation of a tactical message report on one of the local STNs and its replication to the other STNs. When a script event is executed, the template maps the pre-defined data fields for that event to the local database tables. Templates for ten tactical message types have been implemented, but for the experiments described in this report, the only message type employed was the Own Station Position Report. The mechanisms for exchanging this information with other nodes are discussed below. At any time, the simulation can be paused, the databases locked, and the database tables on each node interrogated to populate the measurement logs. Once execution of the script has been completed, the data in the measurement log tables are analyzed to correlate the quality of the shared tactical picture (as represented by the database content on each STN) with the time-varying behaviour of the simulated tactical network.

Information exchange in the Low Bandwidth Test Bed is accomplished via a custom data replication mechanism that replicates data between databases on participating network nodes. Replicated payload content for a particular script event is based upon the information content of the tactical message type associated with that event in the script. Two custom replication mechanisms are available:

a) Rule-Based Replication Implemented using Triggers and Stored Procedures

   When a particular event instance is encountered in the event script, a stored procedure is called to prepare the replication payload. At the same time, a separate stored procedure is called containing the

control logic that decides whether to permit or suppress replication of the replication payload associated with that event instance. This mechanism is event-driven and permits the use of information management rules that provide fine-grained context-sensitive control over the replication process based upon knowledge of battlefield state and network state.

b) Scheduled Replication

In contrast to the first mechanism, the second mechanism is time-driven. This mechanism approximates closely the incremental update employed by the ATCCIS Replication Mechanism (see Annex D). The mechanism employs a selective negotiated 'push' of information on a periodic basis from a provider node to one or more receiver nodes. The information content to be pushed is specified in a 'contract'. The contract specification consists of the set of database tables from which data will be extracted on the provider node, one or more filter criteria to specify which rows from those tables will be selected, the identity of the receiver nodes, and the wait period between successive contract fulfilments. This mechanism does not provide the context-sensitive control of replication available with rule-based replication.

For the experiments reported in this chapter, only the rule-based replication mechanism was employed.

Underpinning each of these mechanisms is a Replication Transport Layer (RTL) with features designed for the bandwidth-constrained wireless domain (Figure 7-1). These features include a system of priority output queues and advanced load-balancing mechanisms. As well, it uses a multicast protocol, implements a packet recovery mechanism based on negative acknowledgement, and can manage transmissions, requests for retransmission, and retransmissions. The RTL interfaces to the transport layer through a User Datagram Protocol (UDP), which acts as a pass-through to the network layer. UDP is used in place of TCP because, in an environment with bit errors and long latency, TCP's congestion controls and timeouts seriously degrade throughput. The RTL provides some functionality such as packet sequencing normally provided by TCP.

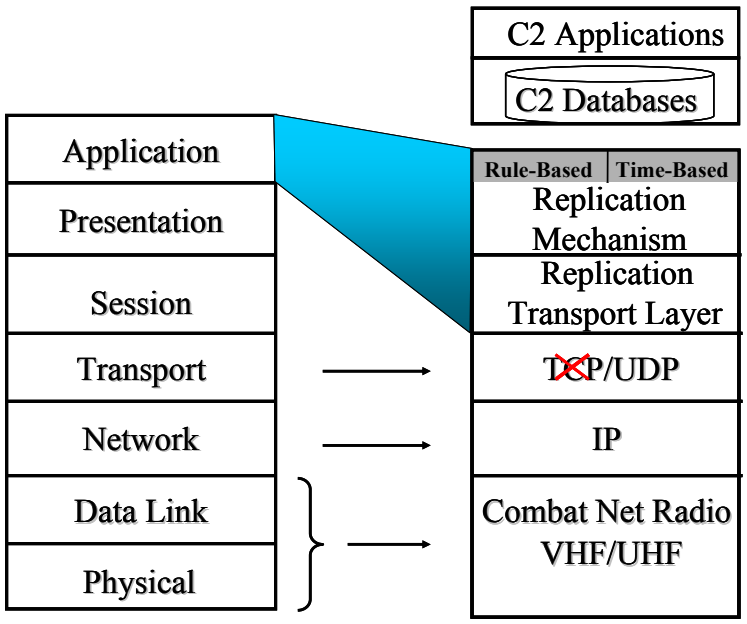

Figure 7-1: Position of Custom Replication Mechanisms in Network Protocol Stack.

The Low Bandwidth Test Bed simulates in software a single-hop subnet of combat net radios (i.e., all radios on a common frequency; multi-hop routing between nodes does not occur). The simulator implements a performance model of the network layer and the data link layer. The term 'performance model' refers to the fact that, for certain functionality, the functionality itself is not fully implemented, but the effect that that functionality would have if it were present is accurately modelled. The Media Access Control functionality for the data link layer is based on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA).

Forward error correction capability modelled in the data link layer includes a majority vote detection (MVD) scheme[1] and Golay coding[2]. Simulations can be executed with either or both error correction mechanisms activated or deactivated. Application of the majority vote detect algorithm effectively multiplies the size of the transmitted protocol data unit by a factor *n*, while the application of Golay coding effectively doubles the size of the transmitted protocol data unit.

At the radio physical layer, the radio channel model uses a probability function to generate bit errors on the channel during transmission. One of several channel profiles can be selected. One profile permits perfect channel conditions ('all-pass') or extremely bad channel conditions ('all-fail') to be specified. A second profile permits the overall percentage of bit errors or packet errors to be specified. The model does not have the ability to model transmission on a link-by-link basis. Rather, the model assumes that, for a given receiving node, the channel profile that applies to received transmissions is independent of the identity of the sending node. The radio model also supports two modes of operation – 'data only' or 'mixed voice and data'. In mixed mode, voice transmissions can pre-empt data transmissions.

All experiments described in this report used the following characteristics:

- 'Data-only' mode of operation;
- Overall percentage of packet errors (0%, 10%, 20% and 30% packet loss);
- Majority Vote Detection with $n = 5$; and
- Golay FEC deactivated.

7.2 OVERVIEW OF EXPERIMENTS

The experiments described in this chapter are based on a scenario in which 15 tactical nodes participating on a single tactical subnet with a common assigned frequency exchange updates of their own position at an interval of 60 seconds, over a period of 60 minutes. Eleven of the nodes are moving at a constant speed of 40 km/hr while four of the nodes are moving very slowly (2 km/hr). The differential in speed was selected to serve the requirement of an information management rule, described below, which selectively suppressed replications of position updates based on vehicle displacement. One of the important measures described below, position error, is directly proportional to vehicle displacement, and so is sensitive to both vehicle speed and trajectory. To eliminate distortions due to varying vehicle speed and trajectory, a given node was assigned the same speed in each experiment, the vehicle trajectories were defined as parallel straight lines, and the vehicles were all moving in the same direction. The rule-based replication mechanism was employed for all experiments.

[1] In a majority vote detection scheme, each bit of real data is duplicated an odd number of times *n* in the transmission. At the receiving node, the values of the duplicate bits are compared. If, due to transmission error, not all *n* duplicate bit values agree, the value that is in the majority is assumed to be the correct value.

[2] Golay coding is a form of forward error correction coding in which, for a number *m* of data bits, *k* additional bits are added to the original *m* bits. These *k* bits are used to detect and correct errors in the original *m* bits on a receiving node.

The objective of the experiments was to demonstrate how information management protocols operating at the application layer can improve the throughput of operationally important information over a tactical radio subnet under conditions of network congestion (i.e., when the offered load to the network exceeds the network's capacity, resulting in reduced transmission success rate and high latency). Two approaches are demonstrated:

1) Dynamic reduction of the offered load through use of an information management rule to selectively suppress transmissions; and
2) Reduction in payload size to reduce transmission time on the radio channel.

7.2.1 Dynamic Reduction of Offered Load through Use of an Information Management Rule

The rule-based replication mechanism described in Section 7.1 causes a stored procedure to be called on a local node whenever a replication message is generated on that node. The stored procedure contains the information management rule that determines whether or not to suppress replication of that replication message. The particular information management rule employed for these experiments is predicated on the supposition that it may be advisable for stopped or slowly moving nodes to broadcast their position updates less frequently than the faster-moving nodes when the network is congested. Specifically, the rule implements the following logic:

> "If the cumulative distance travelled by the local node since the last permitted replication of a position update is less than *x* meters, suppress transmission of the present position update."

For the experiments, a value of $x = 600$ meters and a replication interval of one minute were selected. This means, that for the eleven nodes moving at 40 km/hr, the vehicle displacement satisfied the rule each time and replication of the position update was allowed. On the other hand, for each of the remaining four nodes that were moving at only 2 km/hr the cumulative displacement exceeded 600 meters only once every 18 minutes. Replication of position updates from these nodes was thus suppressed 17 out of 18 times.

The LBTB permits exactly the same scenario to be repeated under exactly the same simulated radio channel conditions first with, then without, the IM rule active. Analysis showed that the overall effect of using this IM rule in this particular scenario was to reduce the offered load to the network by 25%. The results of simulations with and without the IM rule active are reported in Section 7.3.2.2.

7.2.2 Reduction in Payload Size through Choice of Payload Format

The second set of experiments examined the impact of reduced payload size on the throughput of operationally important information. Two types of payload were examined:

a) ODB PDU

The format for this payload type conforms to the format of a Replication Manager Data PDU defined in the specification for the ATCCIS Replication Mechanism (now known as a 'MIP Data PDU' and used by the MIP Data Exchange Mechanism). The Canadian version of this Data PDU is termed an 'Operational Database (ODB) PDU'. Data exchange between ODBs exploits the 'ODB Data Model', which is based on the MIP C2IEDM data model with some Canadian extensions. The version of C2IEDM known as 'Generic Hub 4' is used in the LBTB.

b) Simple Format

This format was developed to support the LBTB design by the LBTB contractor. It is more compact than the ODB PDU format. The Scenario Script within the LBTB consists of a series of communication

events; each event corresponds to the generation of a tactical message report on a specific node, and the replication of that report to other nodes, at a specific scenario time. When a script event is encountered during a simulation, a template corresponding to the report type for that event takes the pre-defined data for the event stored in a file and inserts the data into the appropriate tables in the database on the local Simulated Tactical Node. The same template is then used to extract the just-inserted data from the database to compose the replication message. The Simple Format exploits the fact that, within a template for a particular report type, the relevant set of database tables, the referential relationships between those tables, and the order of data insertion, are well defined. If the same template also exists on the receiving nodes, this information can be exploited to parse and insert the received data into the database tables on the receiving node as well. Information about table identities, order of insertion, etc. does not have to form part of the replicated information because it is already part of the template. What need to be transmitted over the network are the data values in appropriate order and format, and a call to the stored procedure on the receiving node that contains the appropriate template.

Elimination of metadata about database structure permits creation of a much smaller transmission payload. The same position update information that requires 1134 bytes to transmit in ODB PDU format requires only 161 bytes in Simple Format. It should be noted that, in using the Simple Format, one is trading off generality to achieve reduced payload size. The contracting mechanism that is the basis of the ARM is, by design, general in nature. There are few restrictions on the identity of the nodes that establish a contract, and on contract content. This generality is deliberate. The ATCCIS/MIP efforts are focused on achieving interoperability in the context of a NATO coalition. However, the generality comes at a price. Table identities that cannot be predefined, or defined by reference, must be transmitted each time, resulting in larger payloads. In using the Simple Format, one accepts that data exchange will occur only through pre-defined templates with a well-defined structure (the same principle behind the use of structured messages). The loss of generality associated with use of the Simple Format may not be acceptable in all situations. However, in cases where the trade-off is acceptable, the results presented in Section 7.3.2.1 suggest that considerable performance gains may be achieved over tactical communication nets.

7.2.3 Reduction in Payload Size through Use of Data Compression

The effect of reducing payload size through use of a classical lossless data compression algorithm was also examined. Such algorithms achieve compression through detection and removal of redundancies in the data. zlib (compression level 6)[3], the lossless compression algorithm mandated for use by participants in MIP, was implemented in the LBTB and simulations were carried out employing ODB PDU or Simple Format, with and without zlib compression. The average compression factors achieved with this algorithm are listed in Table 7-1:

Table 7-1: Data Compression Achieved with zlib (Compression Level 6)

Payload Format	Payload Size (bytes)		
	Normal	Compressed	Compression Ratio
ODB PDU	1134	250	4.7
Simple	161	117	1.4

[3] Open-source lossless data compression algorithm (http://www.zlib.net). Nine levels of compression are available, offering different tradeoffs between compression level and execution time. Default is level 6.

7.3 ANALYSIS/INTERPRETATION OF RESULTS

This section describes the measures of performance employed and presents the experimental results.

7.3.1 Measures of Performance

7.3.1.1 Location Fidelity

The location fidelity measure is based upon the measurement of the difference between the last reported position of a neighbour node and its actual position within a particular node's database at a particular time. For analysis purposes, the 60-minute scenario duration was divided into one-minute time slices. For each receiving node the error in the position of each of the other (sending) nodes was calculated at one-minute intervals. These errors were then averaged over all sending nodes. For each time slice, an average over all receiving nodes was then performed to produce a network-averaged position error for that time slice. For statistical validity, each simulation was repeated ten times, each time with a different random number seed. The same analysis was carried out for each simulation. For each time-slice, the network-averaged position error from each of the ten simulations was statistically averaged and an expected error in the average value based on the standard deviation was calculated. Plots of network-averaged position error vs. scenario time revealed that position error value achieved a plateau. A time interval lying within this plateau region was selected, and the average value of network-averaged position error over this time interval was calculated. It is this time-averaged and statistically averaged value for network-averaged position error that is quoted in the charts that follow.

7.3.1.2 Currency

The currency measure is based upon the time elapsed since the last successful update of a given piece of information. From data logged during the simulation, an average currency value for position updates, averaged over all receiving nodes and all scenario time, was calculated for each sending node. A network-wide average for currency was then calculated by averaging over the sending nodes[4]. For statistical validity, each simulation was repeated ten times with a different random number seed. A network-averaged currency value was calculated for each simulation. These ten values were statistically averaged and an expected error in the average value based on the standard deviation was calculated. It is this time-averaged and statistically averaged value for network-averaged currency that is quoted in the charts that follow.

7.3.1.3 Latency

The latency measure is based upon a measurement of the actual transit time of a replication message through the radio network from a sending node to a receiving node. The measurement is defined as the difference between the time that a replication message is submitted to the network layer of the protocol stack on the sending node to the time that that same replication message is received, intact, by the network layer of the protocol stack on a receiving node, ready to be passed to the application layer on the receiving node. Latency measurements are recorded only for transmissions that are successfully received. From data logged during the simulation, an average latency value averaged over all receiving nodes and all scenario time was calculated for each sending node. A network-wide average for latency was then calculated by averaging over all sending

[4] For the experiments reported in this chapter, the Information Management rule employed had the effect of suppressing replication of position updates from four of the 15 nodes, thereby increasing the currency for position updates from these four nodes. The decrease in currency for the fast-moving nodes (due to improved network conditions) was masked by the rule-induced increase in currency for the slow-moving nodes. To eliminate this distortion, the currency averages were performed only over the 11 fast-moving sending nodes (all 15 nodes were still counted as receiving nodes).

nodes. For statistical validity, each simulation was repeated ten times with a different random number seed. A network-averaged latency value was calculated for each simulation. These ten values are statistically averaged and an expected error in the average value based on the standard deviation is calculated. This statistically averaged value for network-averaged latency is quoted in the charts that follow.

7.3.2 Experimental Results

For all of the simulations reported in this chapter, strict control of experimental variables was exercised. To compare the impact of a particular variable (e.g., payload format), only that variable was varied from one experiment to the next. The results with and without use of that variable were then compared. Each experiment was performed for four different sets of channel conditions – all pass (no lost packets) and three different levels of packet loss (10%, 20%, and 30%). The variable in question was considered to have a positive influence if a reduction in any or all of the measured quantities, position error, currency, and latency, was observed when that variable was employed. To minimize variability arising from the stochastic nature of the channel access protocol and the transmission process a given simulation was repeated ten times with a different random number seed, and the results averaged as described in Section 7.3.1.

7.3.2.1 Effect of Payload Format

Figure 7-2 compares the network-averaged values of position error, currency and latency obtained for four different packet loss rates with the ODB PDU (size 1134 bytes) and Simple Format (161 bytes). It is clear that, for all packet loss rates, the smaller payload size offered by the Simple Format results in a reduction of all three quantities. Figure 7-3 presents the reductions in position error, currency and latency due to use of the Simple Format, expressed as a percentage of the value obtained with the ODB PDU format. Percentage reductions are significant in all cases, lying in the range of 40% to 70%.

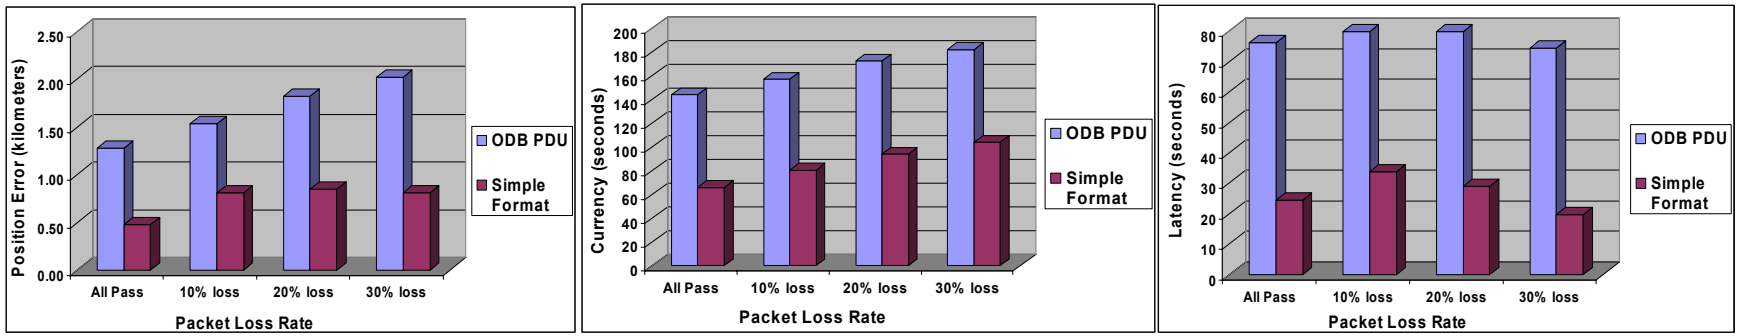

Figure 7-2: Network-Averaged Position Error, Currency and Latency for Different Payload Formats.

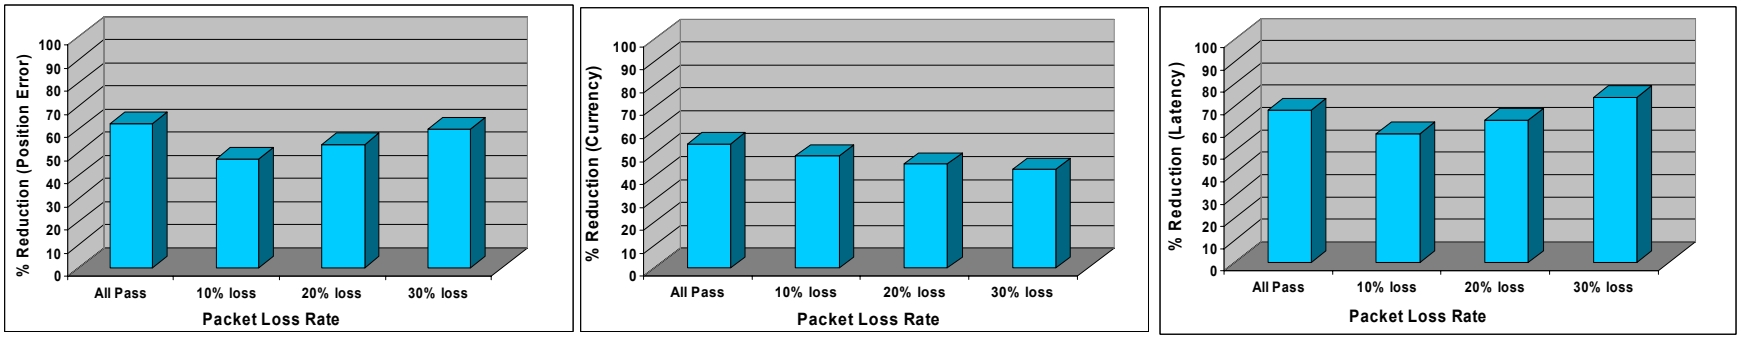

Figure 7-3: Percent Reduction in Network-Averaged Position Error, Currency and Latency Due to Use of Simple Payload Format.

7.3.2.2 Effect of Information Management Rule

Figure 7-4 compares the network-averaged values of position error, currency and latency obtained for four different packet-loss rates without, and with, the IM Rule described in Section 7.2.1 (use of the IM Rule has the overall effect of reducing network traffic by 25%). The first tier of charts corresponds to the ODB PDU payload format. The second tier of charts corresponds to the Simple Format. Figure 7-5 presents the reductions in position error, currency and latency due to use of the IM Rule, expressed as a percentage of the value obtained without use of the IM Rule. For the ODB PDU format, reductions in the range of 30 – 40% are obtained except for the case of 'All Pass', where the reductions are in the range of 60 – 95%. For the Simple Format, the reductions due to the smaller payload size (Section 7.3.2.1) are apparent (Figure 7-4). For non-zero packet loss rates, the IM Rule results in further reductions of 20 – 45% for position error and 25 – 60% for latency (Figure 7-5), with the size of reduction being smallest for the poorest channel conditions. For the 'All Pass' condition, marked reductions of 80% and 90% are observed for position error and latency respectively. Currency shows a more modest reduction of 10 – 20%, because its value is already close to the minimum (60 seconds).

7.3.2.3 Effect of Data Compression

Figure 7-6 compares the network-averaged values of position error, currency and latency obtained for four different packet loss rates without, and with, use of the zlib lossless data compression algorithm. The first tier of charts corresponds to the ODB PDU payload format. The second tier of charts corresponds to the Simple Format. Figure 7-7 presents the reductions in position error, currency and latency due to use of data compression, expressed as a percentage of the value obtained without use of data compression. For the ODB PDU format, reductions of 50 – 70% are observed for position error and latency, while reductions of 30 – 40% are observed for currency. For the Simple Format, the reductions due to the smaller payload size (Section 7.3.2.1) are apparent (Figure 7-6). When data compression is employed, small reductions, in the range of 0 – 15% for position error and latency, and 0% for currency, are observed (Figure 7-7). The small reductions can be attributed to the fact that data compression improves throughput by reducing payload size, but most of this advantage has already been achieved by the choice of the smaller payload format.

7.3.2.4 Combined Effect of Information Management Rule and Data Compression

Figure 7-8 compares the network-averaged values of position error, currency and latency obtained for four different packet loss rates without, and with, combined use of both the IM rule and zlib lossless data compression algorithm. The first tier of charts corresponds to the ODB PDU payload format. The second tier of charts corresponds to the Simple Format. Figure 7-9 presents the reductions in position error, currency and latency due to combined use of IM rule and data compression, expressed as a percentage of the value obtained when no IM rule or data compression are employed. For the ODB PDU format, reductions of 65% to 95% are observed for position error and latency, and 50% to 60% for currency. For the Simple Format, reductions due to smaller payload size (Section 7.3.2.1) are apparent (Figure 7-8). For non-zero packet loss rates, combined use of IM rule plus data compression results in further reductions of 20 – 50% for position error and 20 – 60% for latency (Figure 7-9), with the size of reduction being smallest for the poorest channel conditions. For the 'All Pass' condition, reductions of 80% and 90% are observed for position error and latency respectively. Currency shows a more modest reduction of 10 – 20%, because its value is already close to the minimum (60 seconds).

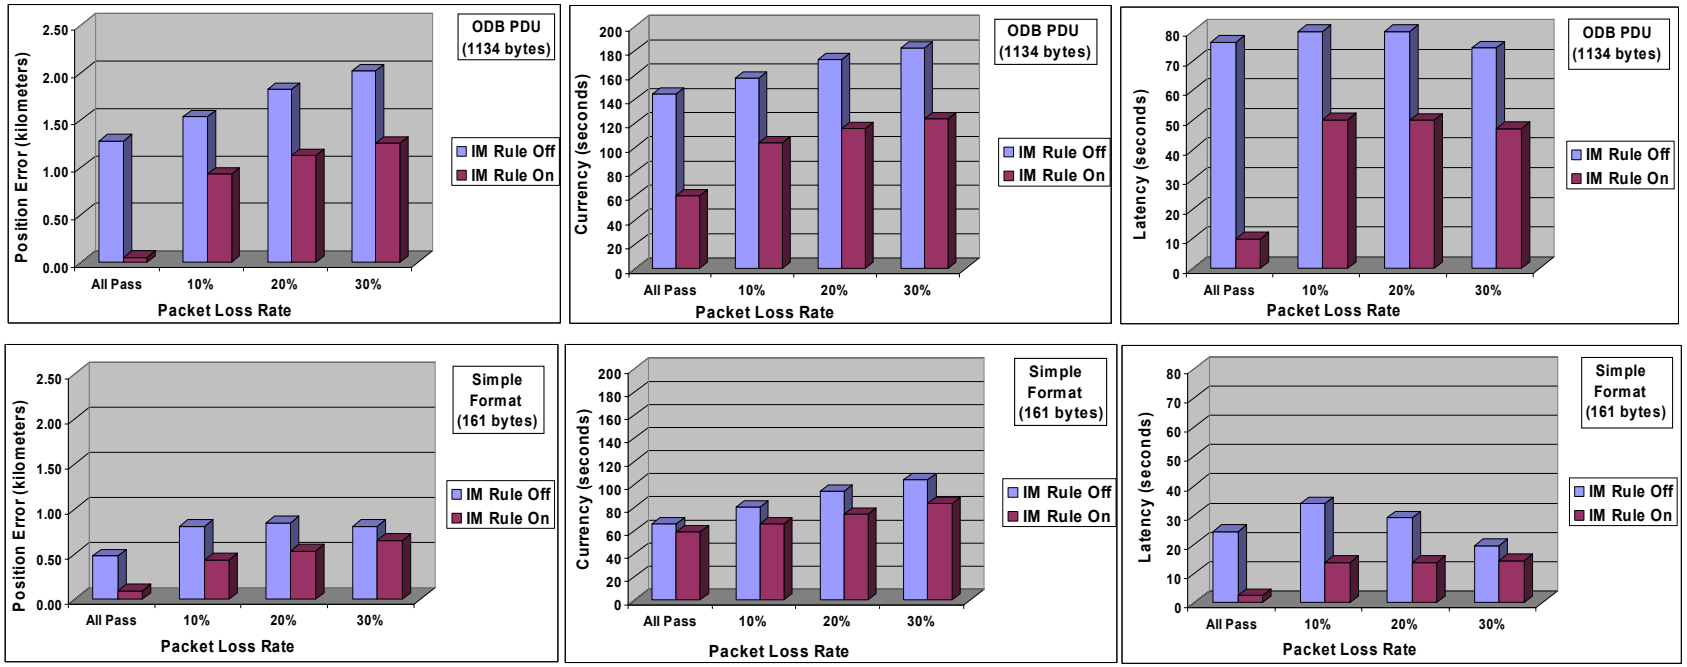

Figure 7-4: Effect of Information Management Rule on Network-Averaged Position Error, Currency and Latency.

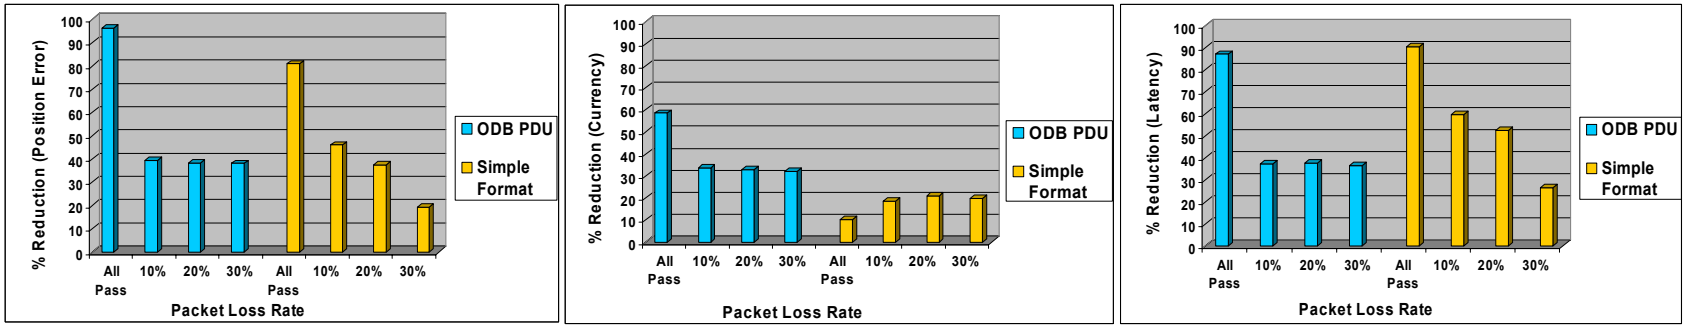

Figure 7-5: Percent Reduction in Position Error, Currency and Latency Due to Use of Information Management Rule.

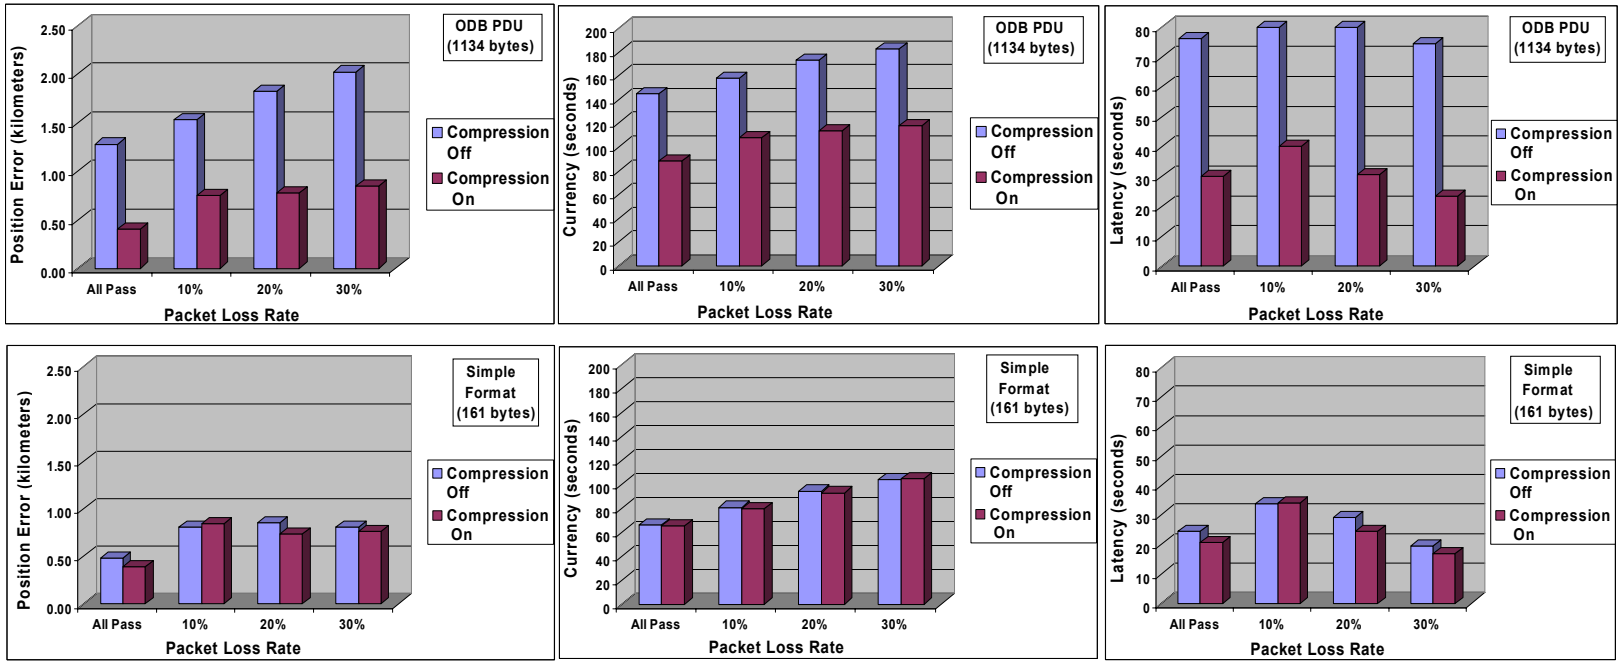

Figure 7-6: Effect of Data Compression on Network-Averaged Position Error, Currency and Latency.

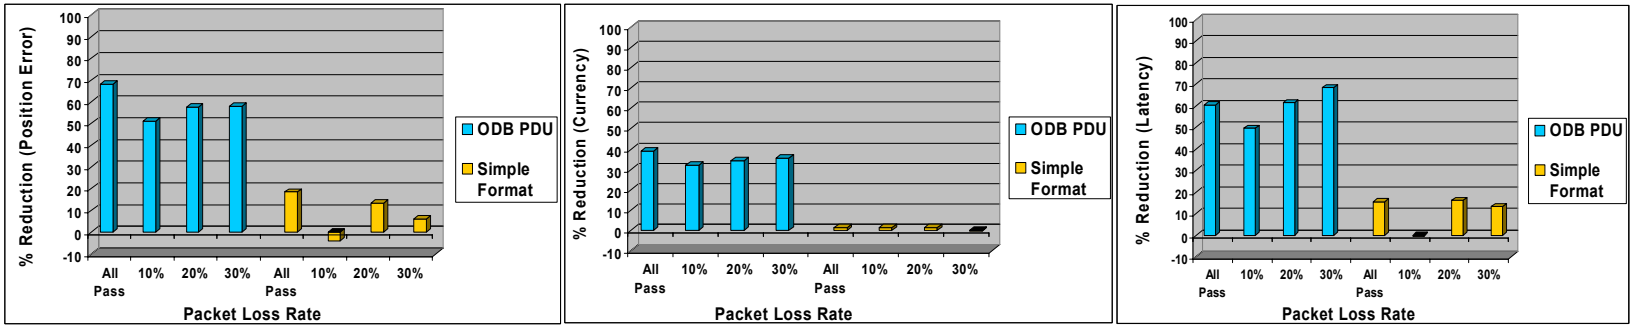

Figure 7-7: Percent Reduction in Position Error, Currency and Latency Due to Use of Data Compression.

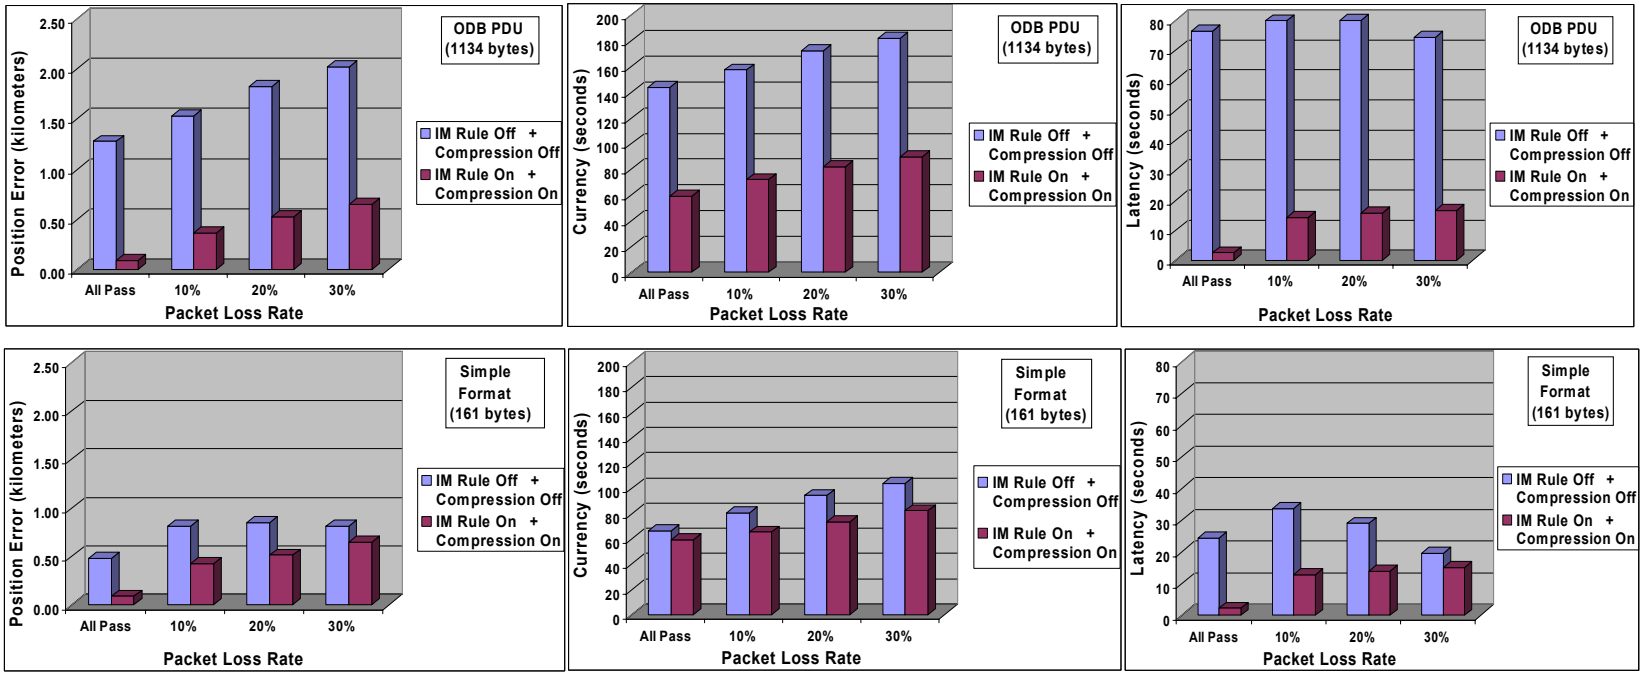

Figure 7-8: Combined Effect of IM Rule and Data Compression on Network-Averaged Position Error, Currency and Latency.

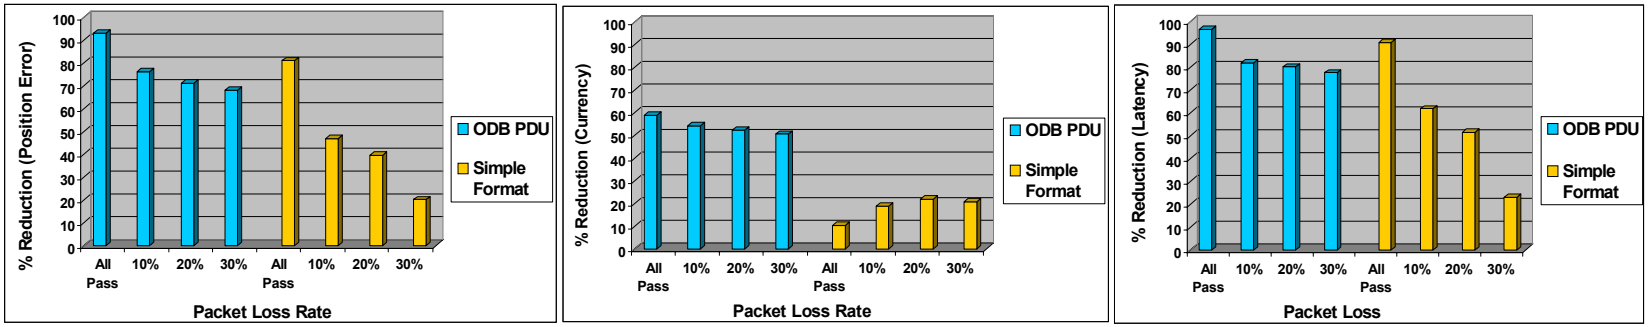

Figure 7-9: Percent Reduction in Position Error, Currency and Latency Due to Combined Use of IM Rule and Data Compression.

7.3.2.5 Summary of Results

For the ODB PDU format, both the IM rule and the data compression, used alone, result in significant reductions in network-averaged position error, latency and currency. The reductions achieved when employed together are greater than when they are employed alone. Data compression reduces payload size, while the IM rule reduces the offered load to the network. For the Simple Format, significant reductions in network-averaged position error, latency and currency are achieved when the IM rule alone is used. However, little reduction is achieved when data compression alone is used since most of the gain achievable through reducing payload size has already been achieved through the more compact payload format. For the Simple Format, virtually all of the reductions achieved when the techniques are employed together are attributable to use of the IM rule.

7.4 SUMMARY AND CONCLUSIONS

This chapter has presented results of simulations conducted to examine the impact of application-layer information management techniques on the throughput of operationally important information over a congested tactical radio subnet. In the simulations, information exchange is accomplished via data replication (Section 4.2) based on an all-informed distribution model (Section 4.2.2.6). An experimental Replication Transport Layer satisfying many of the characteristics identified in Section 4.2.2.4 was employed. The results are specific to the particular network and load conditions used in the simulations, but serve to illustrate the positive impact that application-layer information management techniques that reduce payload size or offered load can have on information flow over disadvantaged tactical communication grids. In particular, the value of using a context-sensitive information management rule (Section 4.2.2 and 4.2.2.4.1) to adjust, without user intervention, the offered load to the network in response to network conditions is demonstrated.