Technical Report

REMPLI Discreet Event Simulation Results

Luís Marques

Filipe Pacheco

HURRAY-TR-070904

Version: 0

Date: 09-24-2007

REMPLI Discreet Event Simulation Results

Luís Marques, Filipe Pacheco

IPP-HURRAY!

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8340509

E-mail: {lmarques, ffp}@isep.ipp.pt

http://www.hurray.isep.ipp.pt

Abstract

This document makes a brief review on the results of the REMPLI Discreet Event Simulation system used to test the REMPLI Transport Layer.

1 Introduction

This document makes a brief review on the results of the REMPLI Discreet Event Simulation system used to test the REMPLI Transport Layer. An introduction on the REMPLI Discreet Event Simulation system is made on HURRAY-TR-070903.

2 Simulation Results and Analysis

Having a working simulator is not enough to accurately assess the applicability of the proposed protocol for real applications. Special care has to be put on the process of selecting and choosing simulation scenarios so that they closely resemble real life scenarios and at the same time allowing the detection of parameters or conditions that may significantly change the temporal behaviour of the proposed protocol. (e.g. scalability of the network – how uniformly changes the temporal behaviour when adding or removing stations). [A. M. Law and W. D. Kelton, "Simulation modelling and analysis", 3rd ed. New York: McGraw-Hill, 2000.]

Two simulation scenarios were chosen at this point resembling a real scenario with different network occupation rate.

2.1 Simulation Case Study 1

In this simulation scenario there are two Access Points interfacing the electrical service company computer network with the power line network. There are a hundred Nodes in the power line network, but only forty of them can actually communicate with the Access Points as the other sixty are too far away on the power line mesh for effective signal processing as signal-to-noise ratio is already very low. The next figure presents a logical topology of the power line communication network.

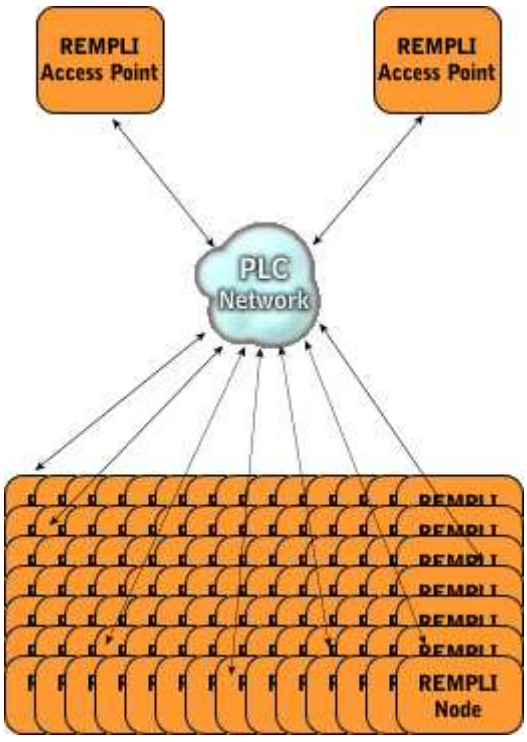

Figure 1 – 6.1 Simulation Case Study 1

At the Access Point 3 periodic messages were used to test the different services available for the Access Point applications. The messages had at least 4 bytes, used to store the message creation time, needed for simulation results and analysis. The messages are described below:

- Periodically each Access Point sends a Request with Response message to a random Node. The period is given by a uniform distribution function with minimum 5 seconds and maximum 15 seconds. The message length is 4 bytes plus a random number of bytes given by an exponential distribution function with average 26 bytes. This simulates an obvious required service of the power company – remote metering.

- Each Access Point also sends an Unconfirmed Unicast message periodically to a random Node. The period is given by a uniform distribution function with minimum 0.4 seconds and maximum 0.6 seconds. Again the message length is 4 bytes plus a random number of bytes given by an exponential distribution function with average 26 bytes. This simulates the updating of data for a specific Node that may be needed for its functioning.

- Finally each Access Point periodically sends an unconfirmed Multicast message to all the Nodes. The period is given by a uniform distribution function with minimum 1.8 seconds and maximum 2.2 seconds. This message length is also 4 bytes plus a random number of bytes given by an exponential distribution function with average 26 bytes. This simulates the updating of multiple Node data at the same time.

At the Node a periodic message was used to test the alarm services available for the Node applications. That message is described below:

- 6 pre-selected Nodes periodically send an Alarm message to all the Access Points. The period is given by a uniform distribution function with minimum 10 seconds and maximum 30 seconds. This message length is 4 bytes plus a random number of bytes given by an exponential distribution function with average 26 bytes. This simulates the Nodes informing the applications at the Access Point of anomalous situations (e.g. energy thieving, short-circuit, energy underrun, etc).

2.1.1 Alarm Service Results

During the simulation 1074 Alarm messages were sent from the Nodes to the APs. Both APs received each of the messages resulting in a total of 2148 Alarm Messages received. The temporal results are expressed in the following table and histogram barchart:

Alarms received by the APs	
Number of Alarms:	2148
Minimum Arrival Time:	41 ms
Maximum Arrival Time:	1946 ms
Average Arrival Time:	524.78 ms
Standard Deviation:	306.45
Variance:	93912.11

Table 1 – Case Study 1 - Alarms received by the APs

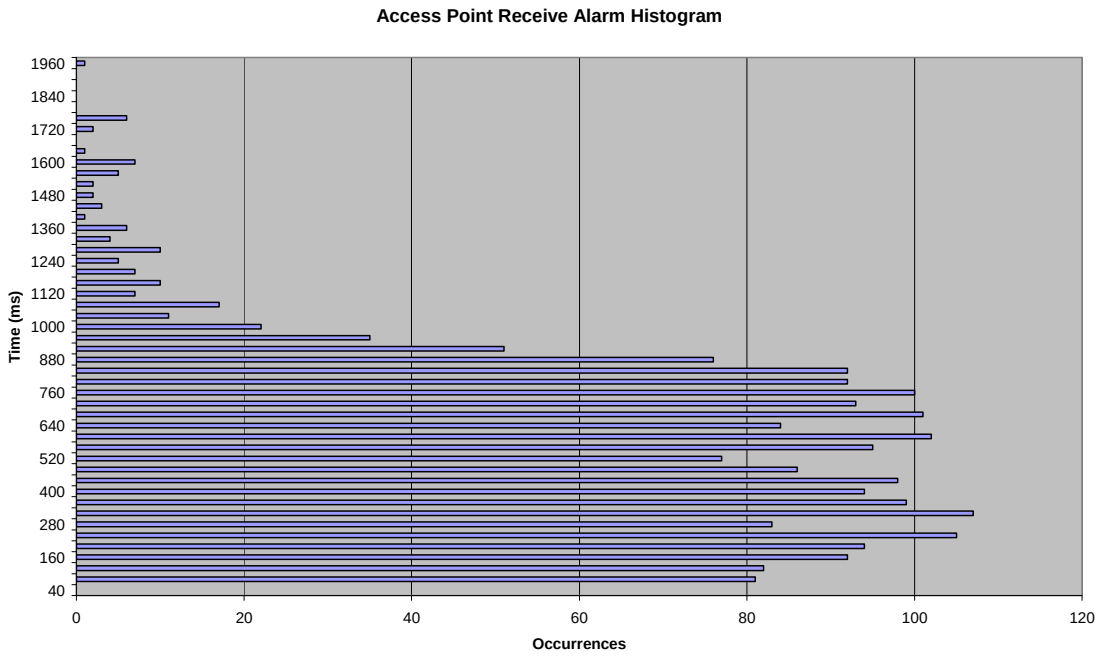

Figure 2 – Access Point Receive Alarm Histogram

Each of the Nodes received an AlarmOK message confirming the reception of the Alarm Message by each AP, resulting in a total of 2148 messages. Only the first AlarmOK for each distinct Alarm are considered for the analysis of the temporal behaviour of the Alarm Service. It is the first AlarmOK message that instructs the Node to stop sending the alarm message to other APs, so it is important to analyse the temporal behaviour of the reception of the first AlarmOKs. The following table and histogram bar chart show the temporal results for the first AlarmOKs received by the Nodes:

AlarmsOKs received by the Nodes	
Number of Alarm OKs:	1074
Minimum Arrival Time:	101 ms
Maximum Arrival Time:	1374 ms
Average Arrival Time:	432.39 ms
Standard Deviation:	227.03
Variance:	51543.52

Table 2 – Case Study 1 - AlarmsOKs received by the Nodes

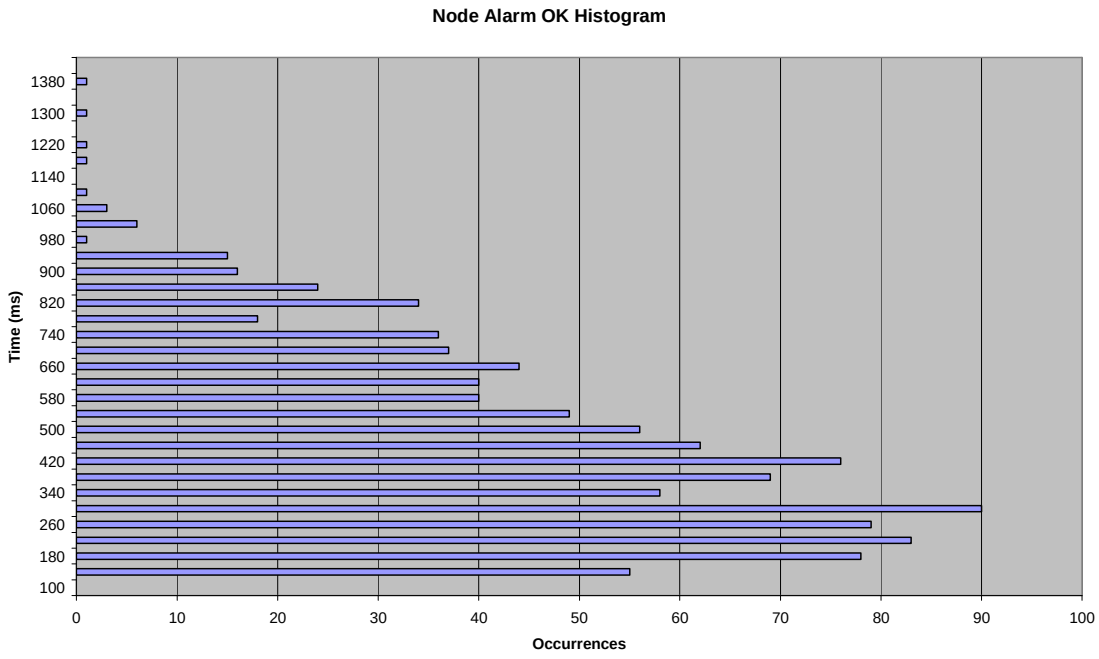

Figure 3 – Access Point AlarmOK Histogram

2.1.2 Request with Response Service Results

723 requests were sent from the Access Point Drivers during the simulation. The temporal results are expressed in the following table and histogram bar chart:

Requests to the Nodes	
Number of Requests:	723
Minimum Arrival Time:	41 ms
Maximum Arrival Time:	153 ms
Average Arrival Time:	59.91 ms
Standard Deviation:	17.51
Variance:	306.70

Table 3 – Case Study 1 Request's Temporal Results

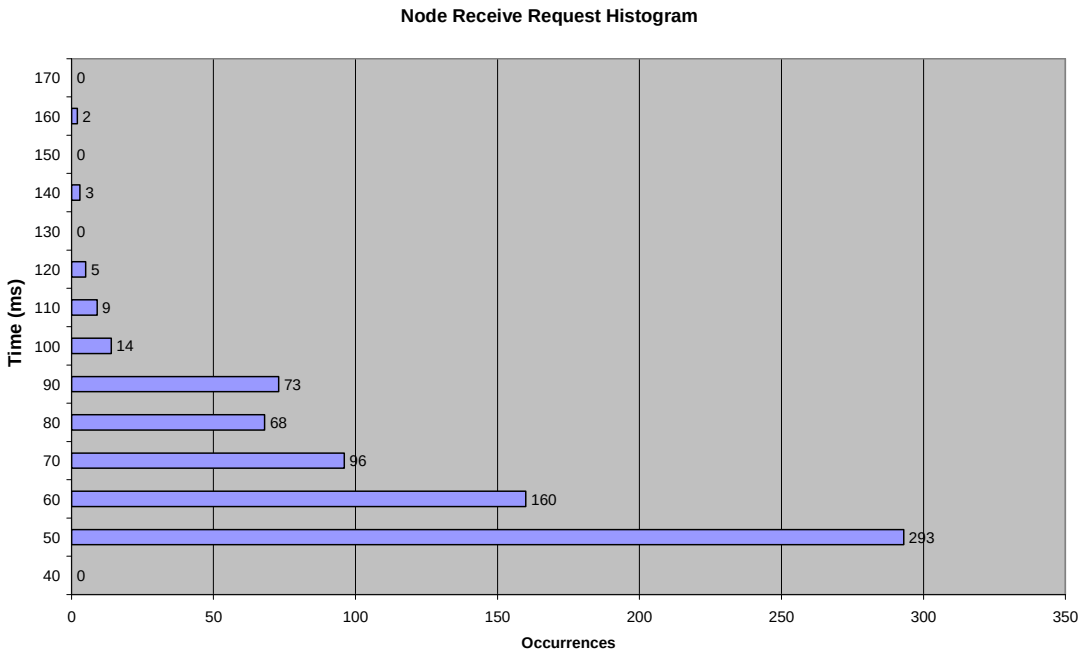

Figure 4 – Node Receive Request Histogram

All of the requests were successfully responded by the Nodes. The following table and histogram bar chart show the temporal results for the overall time since the creation of the request and the arrival of its response:

Responses to the Access Points	
Number of Responses:	723
Minimum Arrival Time:	181 ms
Maximum Arrival Time:	1909 ms
Average Arrival Time:	682.87 ms
Standard Deviation:	299.70
Variance:	89821.05

Table 4 – Case Study 1 Response's Temporal Results

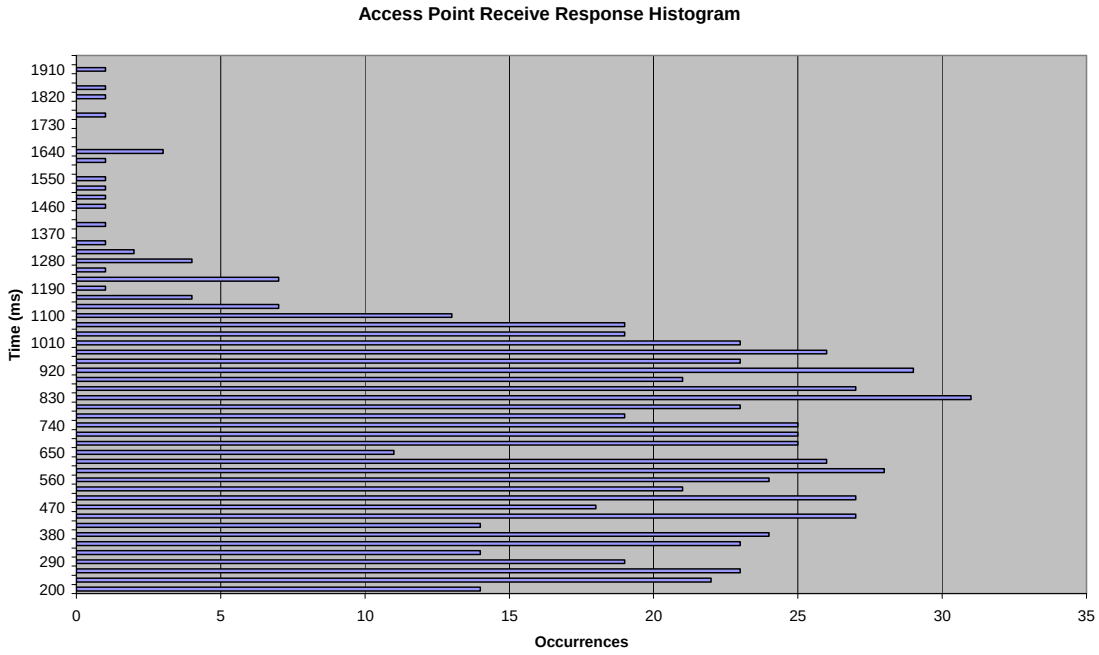

Figure 5 – Access Point Receive Response Histogram

2.1.3 Unicast Unconfirmed Requests Results

14423 unicast messages were sent from the Access Point Drivers during the simulation. The temporal results are expressed in the following table and histogram bar chart:

Unicast Messages Received	
Number of Requests:	14423
Minimum Arrival Time:	41 ms
Maximum Arrival Time:	203 ms
Average Arrival Time:	60.34 ms
Standard Deviation:	17.98
Variance:	323.11

Table 5 – Case Study 1 Unicast Temporal Results

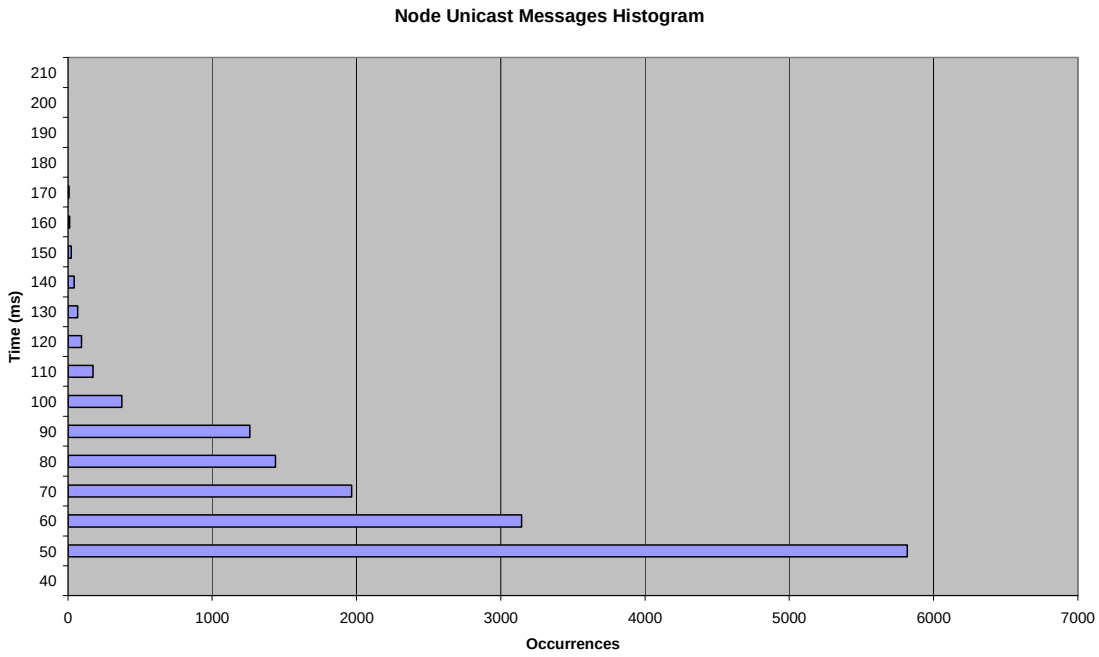

Figure 6 –Node Unicast Messages Histogram

2.1.4 Multicast/Broadcast Requests Results

3594 multicast messages were sent from the Access Point Drivers during the simulation. The temporal results are expressed in the following table and histogram bar chart:

Multicast Messages Received	
Number of Requests:	3594
Minimum Arrival Time:	41 ms
Maximum Arrival Time:	169 ms
Average Arrival Time:	59.75 ms
Standard Deviation:	17.56
Variance:	308.41

Table 6 – Case Study 1 Multicast Temporal Results

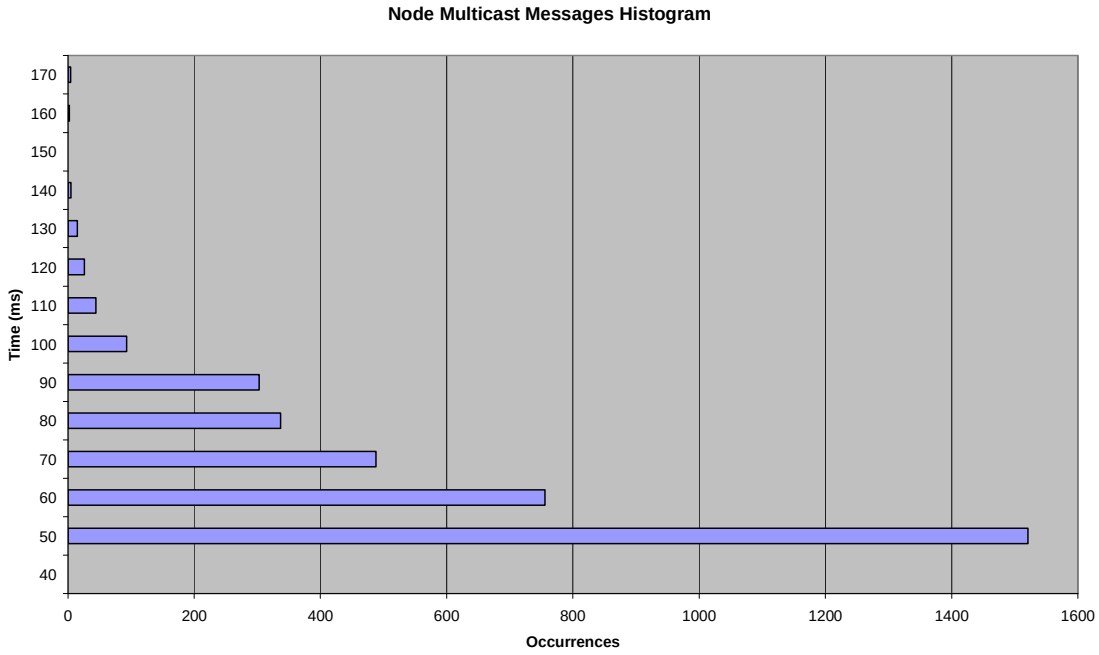

Figure 7 –Node Multicast Messages Histogram

2.1.5 Preliminary Case Study Results Analysis

It is notorious in all message types that the time between the instant a message is sent until the time it is received is quite variable in a way that can't be explained just by message lengths. The unpredictable behaviour of the PLC network is the obvious reason for such time differences. The proposed protocol has successfully delivered all of the messages so the error recovery mechanisms are adequate at least for small messages.

Alarm Service Results

Although the maximum time for an Alarm message being received by an Access Point since its sending was 1946ms the average time was just 524.78ms. In the histogram is easy to spot that just a few messages actually took more than 1000ms (107 messages to be more precise - 4.98% of all messages). Interestingly the messages arriving between 40ms and 880ms were almost uniformly distributed among the intervals, which demonstrate the unpredictability of the PLC network. This results in a high variance (and standard deviation).

For the AlarmOK messages the variance is actually lower. The reason is that the Node who has sent the Alarm Message needs the confirmation from just one the Access Points, probably the first receiving the Alarm, which results on smaller times. Because of this the time until reception of an Alarm message by the "slower" Access Points can be higher than the time between both the sending and the reception of the AlarmOK by the Node. This can be seen in the results as the worst time for receiving an Alarm message was greater that the worst time for receiving an AlarmOK message. This time the histogram shows that the messages are not so uniformly distributed in its intervals and that there is a greater amount of AlarmOKs that took less than the average 432.39ms.

Request with Response Service Results

As this is a master side service the result times were quite lower than for the Alarm Service. The Requests were received by the Nodes after an average of 59.91ms of being sent, with minimum 41ms and maximum 153ms. This resulted in the variance being just 306.70, which is relatively low.

For the Responses the results are considerably worst, which is not surprising as the Slave Network Layer can only send messages when there is an available time slot (reserved by the Master for Slave use as a response to a request). The Maximum Arrival Time was 1909ms which is very high when compared to the request Maximum Arrival Time of 153ms. The variance was high again (89821.05).

Unicast Unconfirmed Requests Results

This is also a master side service so the result times were quite lower than services that require slave side communication. The results were expected to be similar to the Request results of the Request with Response Service. Interestingly they were actually a little worst. While the Minimum Arrival Time was 41ms in both cases the Average Arrival Time for this service was 60.34ms, which is an increase of 0.43ms. This can be explained by the sending of slightly bigger messages in this service due to random number generation average being slightly different from estimated average. The variance was just 323.11.

Multicast/Broadcast Requests Results

For the Multicast/Broadcast Requests the results were similar to the Unicast Unconfirmed ones. This was also expected as the average message length is the same, and both are master-side services. The average Arrival Time was 59.75ms, which is slightly less, leading to a smaller variance (and standard deviation).

2.2 Simulation Case Study 2

For this second case study the simulation scenario remains the same as for the first. The characteristics of those messages are equal to the ones described on case study 1, with the exception of message's length which is far greater. Comparing the results from simulation case study 1 with the ones from this case allows the study of the message length impact on the network.

The message lengths for the periodic messages at the Access Points are described below:

- Each Request with Response message length is 4 bytes plus a random number of bytes given by an exponential distribution function with average 256 bytes.
- The Unconfirmed Unicast messages have a length of 4 bytes plus a random number of bytes given by an exponential distribution function with average 156 bytes.
- Finally each Unconfirmed Multicast message has a length of also 4 bytes plus a random number of bytes given by an exponential distribution function with average 156 bytes.

The message length for the Alarm message at the Node is described below:

- Each Alarm message has a length of 4 bytes plus a random number of bytes given by an exponential distribution function with average 56 bytes.

2.2.1 Alarm Service Results

During the simulation 1076 Alarm messages were sent from the Nodes to the Access Points. This time some Alarm messages were not received by both of them. This is due to the Alarm cancelling mechanism – some alarms were cancelled after one Access Point received the Alarm but not the other. A total of 2061 Alarm Messages were received by the Access Points, and the temporal results are expressed in the following table and histogram barchart:

Alarms Received by the APs	
Number of Requests:	2061
Minimum Arrival Time:	43 ms
Maximum Arrival Time:	4278 ms
Average Arrival Time:	675.00 ms
Standard Deviation:	409.44
Variance:	167644.08

Table 7 – Case Study 2 - Alarms received by the APs

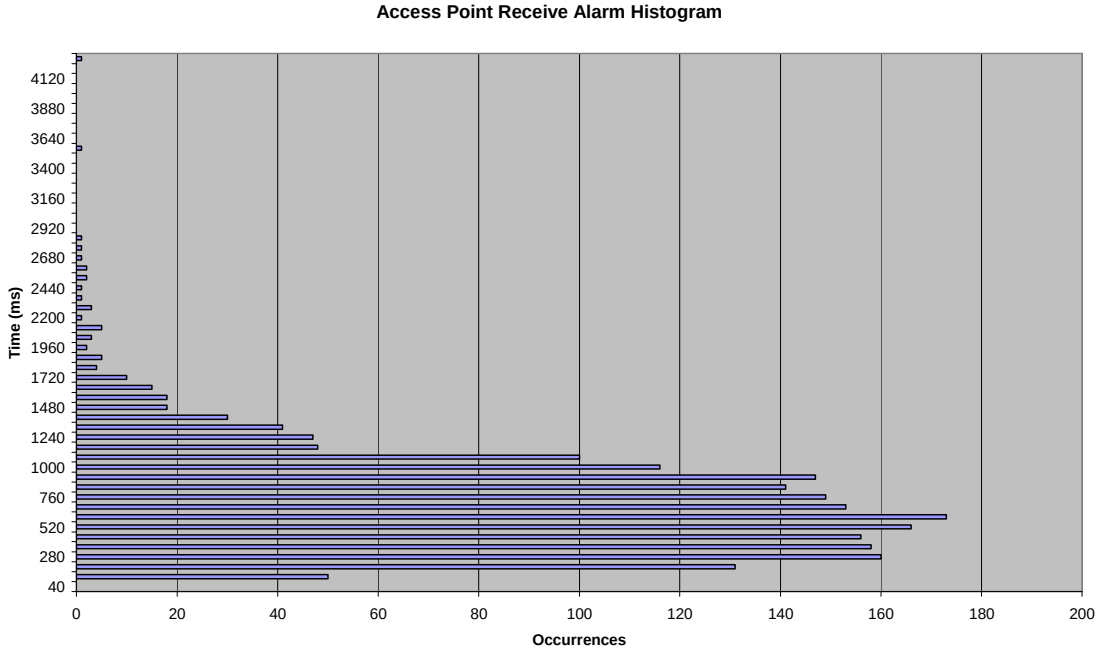

Figure 8 – Access Point Receive Alarm Histogram

Each of the Nodes received an AlarmOK message confirming the reception of the Alarm Message by at least one AP. Only the first AlarmOK for each distinct Alarm are considered for the analysis of the temporal behaviour of the Alarm Service. It is the first AlarmOK message that instructs the Node to stop sending the alarm message to other APs, so it is important to analyse the temporal behaviour of the reception of the first AlarmOKs. The following table and histogram bar chart show the temporal results for the first AlarmOKs received by the Nodes:

AlarmsOKs received by the Nodes	
Number of Alarm OKs:	1076
Minimum Arrival Time:	103 ms

Maximum Arrival Time:	3558 ms
Average Arrival Time:	550.70 ms
Standard Deviation:	304.64
Variance:	92808.38

Table 8 – Case Study 2 - AlarmsOKs received by the Nodes

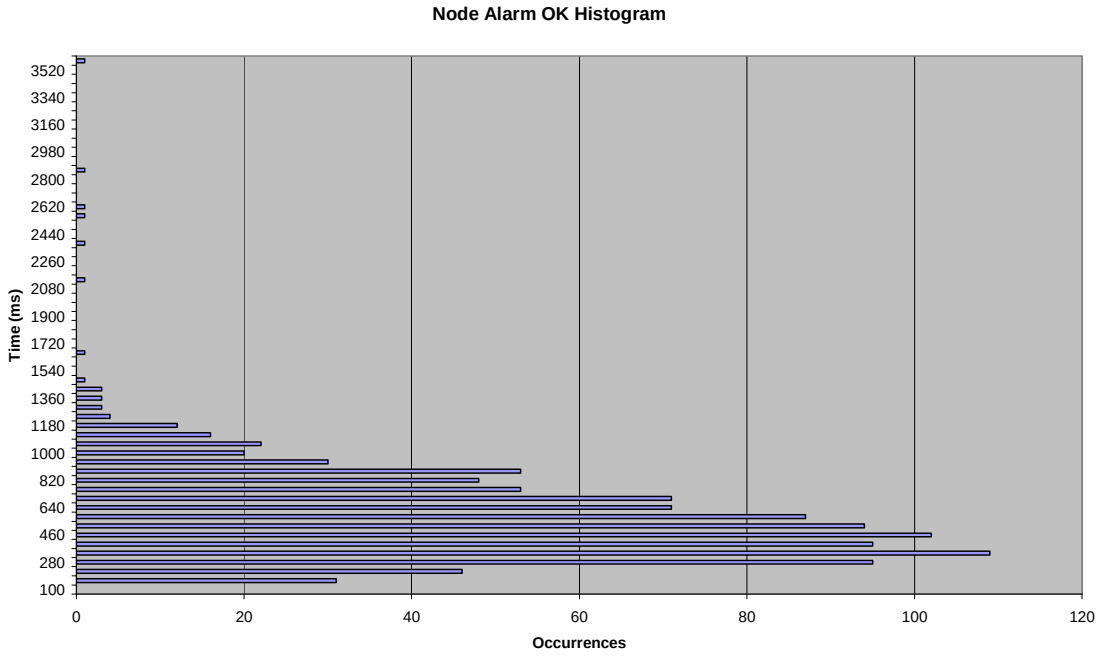

Figure 9 – Access Point AlarmOK Histogram

2.2.2 Request with Response Service Results

711 requests were sent from the Access Point Drivers during the simulation. The temporal results are expressed in the following table and histogram bar chart:

Requests to the Nodes	
Number of Requests:	711
Minimum Arrival Time:	41 ms
Maximum Arrival Time:	1034 ms
Average Arrival Time:	218.20 ms
Standard Deviation:	158.80
Variance:	25216.42

Table 9 – Case Study 2 Request's Temporal Results

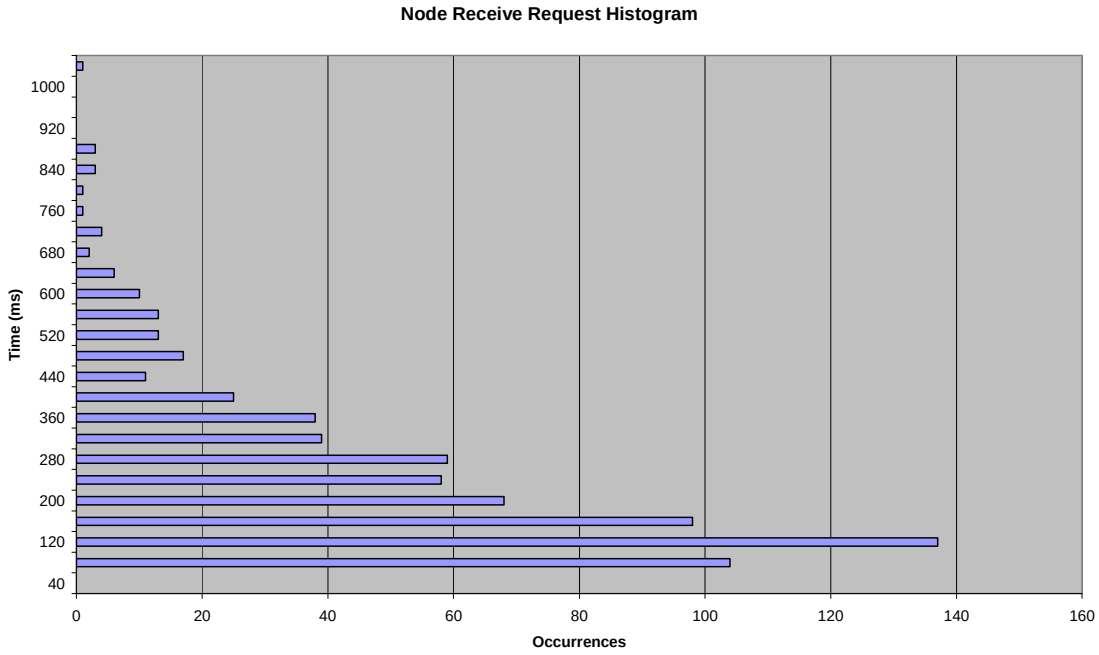

Figure 10 – Node Receive Request Histogram

All of the requests were successfully responded by the Nodes. The following table and histogram bar chart show the temporal results for the overall time since the creation of the request and the arrival of its response:

Responses to the Access Points	
Number of Responses:	711
Minimum Arrival Time:	194 ms
Maximum Arrival Time:	3438 ms
Average Arrival Time:	1119.32 ms
Standard Deviation:	521.92
Variance:	272399.38

Table 10 – Case Study 1 Response's Temporal Results

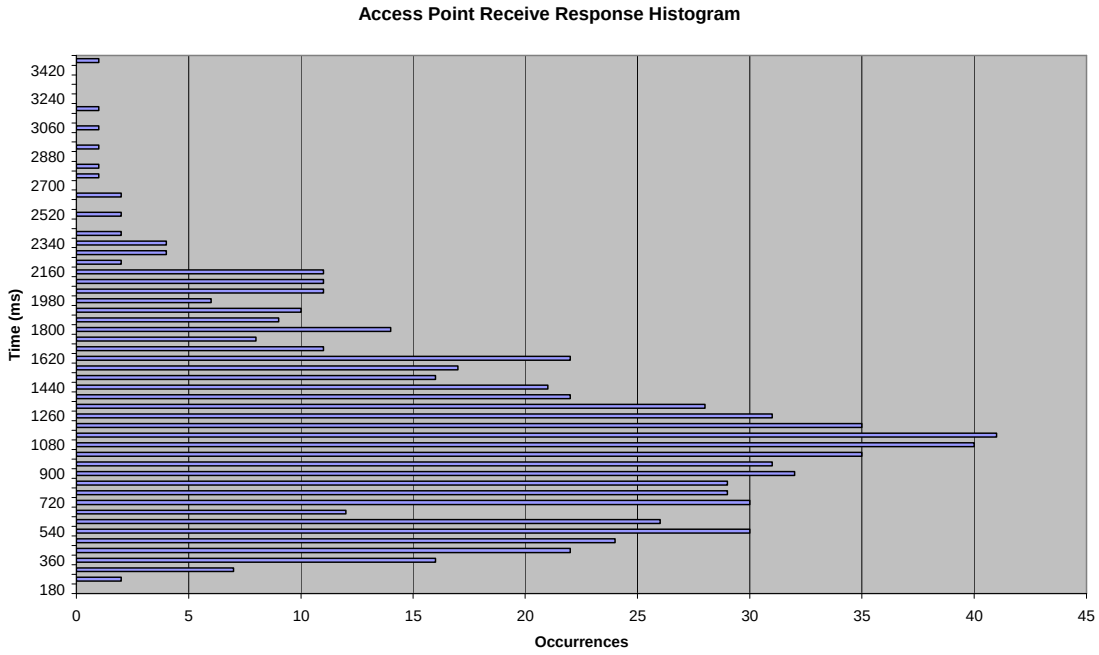

Figure 11 – Access Point Receive Response Histogram

2.2.3 Multicast/Broadcast Requests Results

3602 multicast messages were sent from the Access Point Drivers during the simulation. The temporal results are expressed in the following table and histogram bar chart:

Multicast Messages Received	
Number of Requests:	3602
Minimum Arrival Time:	41
Maximum Arrival Time:	845
Average Arrival Time:	135.84
Standard Deviation:	82.46
Variance:	6800.24

Table 11 – Case Study 2 Multicast Temporal Results

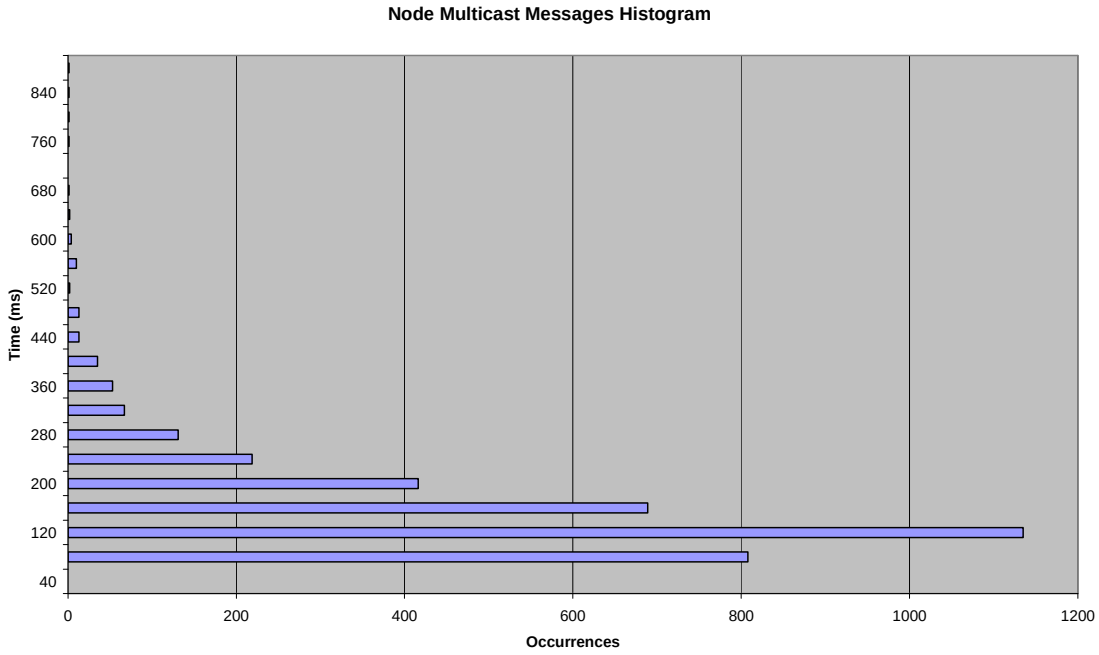

Figure 12 –Node Multicast Messages Histogram

2.2.4 Preliminary Case Study Results Analysis

The proposed protocol has successfully delivered all of the messages so the error recovery mechanisms are adequate at least for small messages.

Alarm Service Results

Again, the maximum time for the reception of an Alarm message (4278ms) is well higher than the average, but almost all messages took less than half that time (99.27% of the messages took less than 2120ms). Although the average size of Alarm messages is double of what was used on Case Study 1 (30 bytes on Case Study 1 vs 60 bytes on Case Study 2) its average time is just 675ms - an increased of just 27%. The reason for this is that as soon as the master knows that a Node has message fragments to send it polls that Node more often reducing the queuing time of the other fragments.

As expected, for the time results of the AlarmOK messages the variance is lower than those of the Alarm messages. In the histogram is visible that just a few AlarmOK messages arrived at the Node more than 1180ms since the sending of the respective Alarm message by the Node Driver (below 2% of AlarmOK messages).

Request with Response Service Results

As this is a master side service the result times were quite lower than for the Alarm Service. The Requests were received by the Nodes after an average of 59.91ms of being sent, with minimum 41ms and maximum 153ms. This resulted in the variance being just 306.70, which is relatively low.

For the Responses the results are considerably worst, which is not surprising as the Slave Network Layer can only send messages when there is an available time slot (reserved by the Master for Slave use as a response to a request). The Maximum Arrival Time was 1909ms which is very high when

compared to the request Maximum Arrival Time of 153ms. The variance was high again (89821.05).

Multicast/Broadcast Requests Results

For the Multicast/Broadcast Requests the results were similar to the Unicast Unconfirmed ones. This was also expected as the average message length is the same, and both are master-side services. The average Arrival Time was 59.75ms, which is slightly less, leading to a smaller variance (and standard deviation).