

Synchrophasor Communication Infrastructure

Impacts on Latency – Part II

2013 TECHNICAL UPDATE

Synchrophasor Communication Infrastructure

Impacts on Latency – Part II

EPRI Project Manager
P. Myrda



3420 Hillview Avenue
Palo Alto, CA 94304-1338
USA

PO Box 10412
Palo Alto, CA 94303-0813
USA

800.313.3774
650.855.2121

askepri@epri.com

www.epri.com

3002000604

Technical Update, August 2013

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THE FOLLOWING ORGANIZATIONS PREPARED THIS REPORT:

Electric Power Research Institute (EPRI)

CISCO Systems

This is an EPRI Technical Update report. A Technical Update report is intended as an informal report of continuing research, a meeting, or a topical study. It is not a final EPRI technical report.

NOTE

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

Copyright © 2013 Electric Power Research Institute, Inc. All rights reserved.



Acknowledgments

The following organizations prepared this report:

Electric Power Research Institute (EPRI)
3420 Hillview Ave
Palo Alto, CA 94304

Principal Investigators
P. Myrda
S. Sternfeld
L. Burnette

CISCO Systems
170 West Tasman Drive
San Jose, CA 95134-1706

Principal Investigator
P. Donner

This report describes research sponsored by EPRI.

This publication is a corporate document that should be cited in the literature in the following manner:

*Synchrophasor Communication
Infrastructure: Impacts on
Latency – Part II.*
EPRI, Palo Alto, CA: 2013.
3002000604.

Abstract

This report details ongoing work that began in 2011 to evaluate the benefits of various wide-area communication approaches for transporting synchrophasor measurements, with a primary focus on latency. The 2012 work focused on evaluating the North American SynchroPhasor Initiative network (NASPInet) architecture design elements and evaluating applications and network device settings that impact various closed-loop control use-cases. This collaborative effort between the Electric Power Research Institute (EPRI), Cisco, OSIsoft, SEL and SISCO led to the development of a test bed at EPRI's Smart Grid Substation Laboratory using their IP-based NASPInet architecture. The design permitted measurements with both phasor data concentrator (PDC) and non-PDC configurations.

The measurement results for Tests B and C corresponded to an acceptable response time for special protection schemes (SPS), or other scenarios where closed-loop control is desired, and matched the expected results for the controlled environment. However, the results from Test A, using mostly default configurations, resulted in measurements which would not have been acceptable for a control loop that required response in the millisecond range. The network architecture and design, including centralized vs. distributed architecture, data transmission distance, location of collection or automation engine relative to data source, and number of participating devices, is an important consideration in developing any control system. Significant improvements in the control loop latency were achieved through application of component level expertise and trial and error involving all components of the test bed, including configuration of device settings, network traffic control, and software interfaces. Significant improvements in the control loop latency were achieved through application of component level expertise and trial and error involving all components of the test bed.

Keywords

Synchrophasor
Phasor measurement unit
NASPInet
Latency
IP Multicast
IEC 61850-90-5

Table of Contents

Section 1: Introduction	1-1
Background.....	1-1
Project Objectives	1-1
Purpose of the Testing.....	1-2
Section 2: Synchrophasor Latency Timing Test Bed.....	2-1
Description of the Control Loop.....	2-1
System Timing.....	2-3
Intermediate Time References.....	2-4
Section 3: Testing Scenarios.....	3-1
Scenario 1 Description	3-2
Scenario 2 Description	3-3
Scenario 3 Description	3-3
Section 4: Results Discussion.....	4-1
Test A Results.....	4-1
Impact of Variation Observed in Test A	4-2
Test B Results	4-3
Test C Results.....	4-4
Test A, B, and C Summary Tables	4-6
General Test Observations	4-8
Section 5: Sub-System Opportunities	5-1
Historian Interface Buffer.....	5-1
Buffer Between PMU Interface and Historian	5-2
PDC Waiting Period.....	5-3
Event and Time Monitor.....	5-4
Section 6: Future Time Synchronization Opportunities	6-1
Section 7: Test Bed Network Connectivity Details	7-1
Section 8: Conclusions	8-1
Test Bed Limitations / Observations.....	8-2
Future Testing Areas	8-3

Appendix A: Test A Raw Data A-1
Appendix B: Test B Raw DataB-1
Appendix C: Test C Raw DataC-1

List of Figures

Figure 2-1 Geographic Locations of SGSL Facilities	2-1
Figure 2-2 Simplified Control Signal Diagram.....	2-2
Figure 2-3 Time synchronization across the Test Bed Control Loop	2-4
Figure 2-4 Intermediate Timing Step Identification	2-5
Figure 3-1 Scenario 1 configuration – IP Unicast with PDC.....	3-2
Figure 3-2 Scenario 2 configuration – IP Unicast without PDC.....	3-3
Figure 3-3 Scenario 3 configuration – IP Multicast.....	3-4
Figure 4-1 Test A Results – 60 samples per second.....	4-2
Figure 4-2 Test B Results – 60 samples per second	4-4
Figure 4-3 Test C Results – 30 samples per second.....	4-6
Figure 5-1 Historian Interface Buffer Configuration Screen.....	5-2
Figure 5-2 Historian Buffer Configuration Screen.....	5-3
Figure 5-3 PDC Configuration Screen with writing period settings.	5-3
Figure 5-4 Example of actual PDC Timing Results	5-4
Figure 6-1 System time adjustments over 15 minute interval periods	6-1
Figure 6-2 IEEE 1588v2 connections to all test bed devices.....	6-2
Figure 7-1 Test bed Network Diagram	7-1
Figure A-1 Test A Frequency of Occurrence.....	A-4
Figure B-1 Test B Frequency of Occurrence.....	B-4
Figure C-1 Test C Frequency of Occurrence.....	C-4

List of Tables

Table 2-1 Intermediate Timing Reference Points.....	2-6
Table 3-1 PMU Timing Test matrix	3-1
Table 4-1 Test A Summary Statistical Results	4-2
Table 4-2 Test B Summary Statistical Results.....	4-4
Table 4-3 Test C Summary Statistical Results	4-6
Table 4-4 Summary Table of Test Results – Mode	4-7
Table 4-5 Summary Table of Test Results – Minimum.....	4-7
Table 4-6 Summary Table of Test Results – Maximum.....	4-7
Table A-1 Test A Raw Data.....	A-1
Table A-2 Summary Statistics for Test A.....	A-3
Table B-1 Test B Raw Data.....	B-1
Table B-2 Summary Statistics for Test B.....	B-3
Table C-1 Test C Raw Data.....	C-1
Table C-2 Summary Statistics for Test C.....	C-3



Section 1: Introduction

Background

This report details ongoing work that began in 2011 evaluating the benefits of various wide-area communication approaches for transporting synchrophasor measurements, with a primary focus on latency. Recent discussions on synchrophasor use in the North American SynchroPhasor Initiative (NASPI) community have included consideration of automated closed-loop control over an Internet Protocol (IP) network infrastructure.

The focus in 2012 has been twofold. First the evaluation of the NASPI network (NASPInet) architecture design elements and second the evaluation of applications and network device setting that impact various closed-loop control use-cases. The project team's architecture for NASPInet uses open standards and off-the-shelf networking equipment, thereby allowing only those participants that can provide the required publish and subscribe functionality directly in the network.

This collaborative effort between EPRI, Cisco, OSIsoft, SEL and SISCO led to the development of a test bed at the Electric Power Research Institute's (EPRI's) Smart Grid Substation Laboratory using their IP-based NASPInet architecture. The design permitted measurements with both phasor data concentrator (PDC) and non-PDC configurations. This approach allows objective measurement of the impacts of these architectural choices. Results of the actual tests are reported.

Project Objectives

The primary objective was to improve the understanding of devices and settings that contribute to the overall latency in a synchrophasor based control system. These included the use of phasor measurement unit (PMU) capable intelligent electronic devices (IEDs) operated with the IEEE 37.118 protocol and also using IEC 61850 Generic Object Oriented System Event (GOOSE) messaging. The project also included the use of components from multiple vendors in an attempt to uncover vendor interoperability issues. Three scenarios were considered for IEEE C37.118 data transport. These included IP Unicast, IP Unicast with a Phasor Data Concentrator (PDC) and IP Multicast. These were selected to provide a sound technical basis for future power system control applications prior to any actual implementations. Preliminary testing of latency considerations seemed prudent as recent discussion at NASPI and elsewhere are considering the use of PMUs as potential elements of power system control scenarios.

Through these objectives one would be able to better understand latency contributions to the overall synchrophasor system by the network, transport protocols, and different vendor equipment and architecture choices. This should create industry awareness for approaches and methods that can be used to minimize latency.

Purpose of the Testing

The primary purpose for this testing was to make a determination to the viability of implementing a closed-loop control scheme utilizing existing substation power transmission and control equipment. Equally important was the desire to determine what configuration steps were necessary to minimize the overall latency of the PMU communication system and to determine areas for improvement.

A closed-loop control scheme was selected and modeled as an external automation system that monitored incoming PMU data and would send a control signal back to a substation IED, similar to what a SCADA system would do. The GOOSE signaling between substation IEDs is representative of a transfer-trip signal, or other substation-to-substation communication message.

The test system utilized PMUs with IEEE C37.118 protocol at rates of both 30 and 60 samples per second. The IEEE C37.118 signaling was propagated within a substation LAN and across a simulated WAN link to another substation. Three transport scenarios were evaluated which included:

- IEEE C37.118 signaling using IP Unicast originating at a PMU and sent directly to an event monitor
- IEEE C37.118 signaling using IP Unicast originating at a PMU and traversing a PDC en-route to an event monitor
- IEEE C37.118 signaling using IP Multicast originating at a PMU and sent directly to an event monitor

In addition, GOOSE messaging was also implemented in order to complete the control loop.

Section 2: Synchrophasor Latency Timing Test Bed

The testing was accomplished using the EPRI Smart Grid Substation Lab (SGSL) which is comprised of a multi site facility. The locations are Knoxville, TN, Charlotte, NC and Lenox, MA. These locations provide hardware simulations representative of typical utility environments and emulate typical utility communications methods. The Knoxville, TN site was the primary center for this project testing. Figure 2-1 depicts the locations and distances between lab facilities.

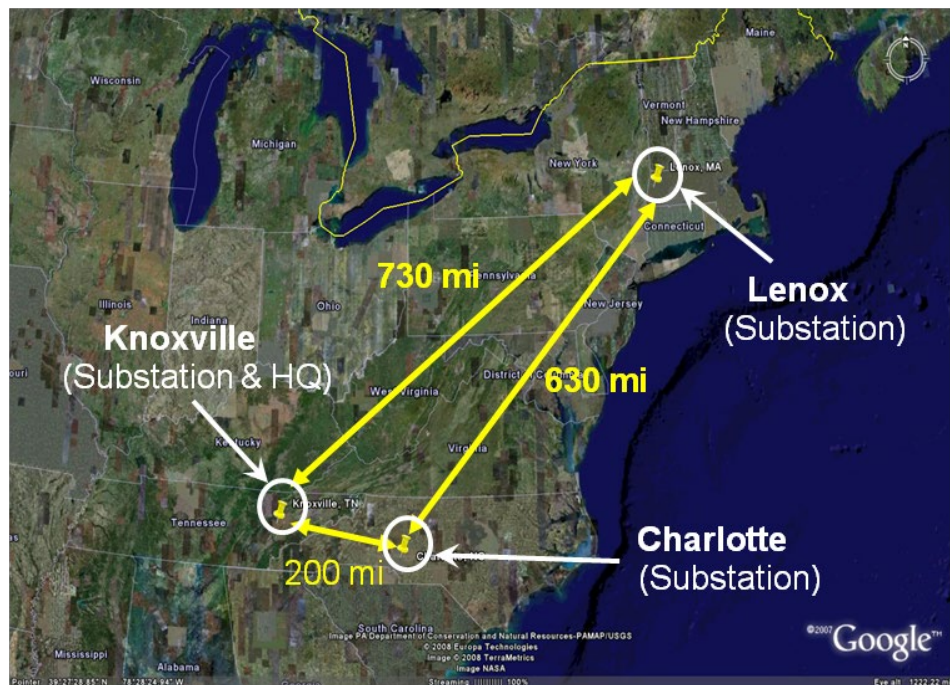


Figure 2-1
Geographic Locations of SGSL Facilities

Description of the Control Loop

The general use case considered for this project starts with phasor measurement unit (PMU) A in Knoxville (KNX) substation monitoring the grid. When the synchrophasor measurement reaches an established control limit, a triggering

event is initiated with a signal sent to PMU B in Lenox (LNX) substation. The specific event was not of concern in this project but could be any event such as a power transfer limit, oscillation, etc. that would be monitored for and when reached send a control trigger. This control event could be part of an automation or protection scheme. The overall goal of the control loop concept was to simulate an event trigger and then capture the overall latency from the time of event initiation to event control action. This was physically accomplished as follows.

The control loop was comprised of both IEEE C37.118 and IEC 61850 GOOSE signaling which spanned two substations. The substations were interconnected using a simulated WAN link over which control of various aspects of a typical WAN link could be modified (e.g. propagation delay, packet loss, jitter and others). Figure 2-2 illustrates a simplified view of the control signal flow.

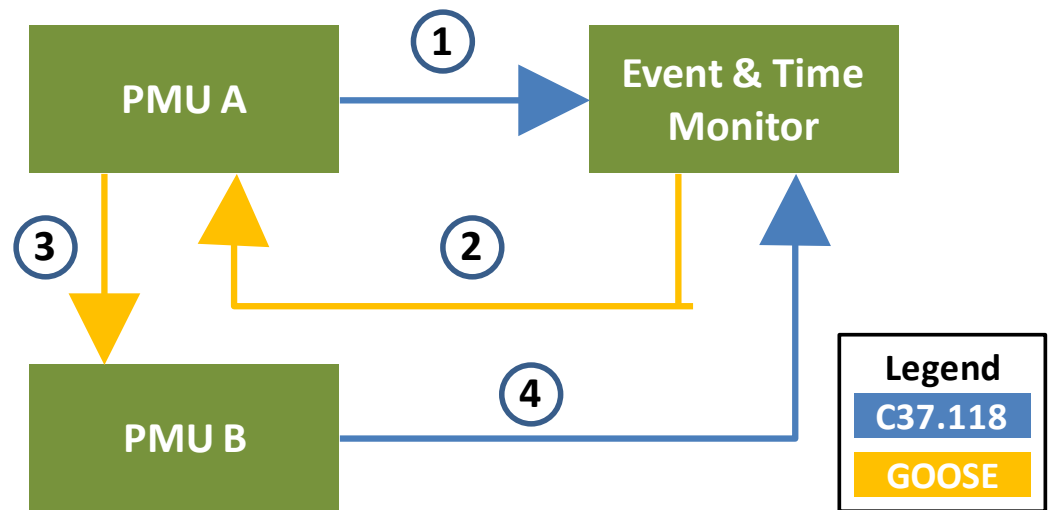


Figure 2-2
Simplified Control Signal Diagram

The four main flows of data in Figure 2-2 are:

1. The control loop simulation starts with a time stamped signal being sent from PMU A to the Event & Time Monitor (E&TM) and the E&TM records both the signal time stamp and the time of arrival.
2. Next the E&TM sends a GOOSE message back to PMU A.
3. Upon arrival of the GOOSE message at PMU A it sends a GOOSE message to PMU B.
4. PMU B sends a time stamped signal to the E&TM where the PMU B signal time stamp is recorded along with the arrival time.

System Timing

To determine the latency of the difference of the event times, a single time reference was needed. The variation of the time source must also be smaller than the PMU measurement window of 16.67ms for 60 samples/second or 33.33ms for 30 samples/second.

The synchrophasor timing test bed utilized a substation grade GPS-synchronized clock which distributed IRIG-B output signal to PMU A, PMU B and the Phasor Data Concentrator. This provides a “high accuracy IRIG” ($\sim 1\mu\text{s}$) to the PMUs.

The historian and event could be used to record the timestamp when the desired packet arrived (noted by a state change in a digital bit). Knowing the “arrival” timestamps would allow more granular understanding of the contributions of each component to the overall latency calculation. This would require the servers to be synchronized to the same time source as the PMUs.

Synchronizing the servers hosting the historian and the event monitor was more challenging. Unlike the PMUs, the servers do not have an IRIG input. The initial approach was to utilize Network Time Protocol (NTP) or a more accurate Precision Time Protocol (IEEE 1588/PTPv2) to propagate time signal information throughout the test bed to ensure a synchronized time measurement between all devices within 1ms or less.

To achieve this, a substation gateway product was utilized that received an IRIG-B signal from the same GPS synchronized clock as the PMUs. In turn, this gateway functioned as the NTP master for the entire substation network. All windows servers and NTP client-capable IEDs were set to synchronize with this NTP server. Figure 2-3 depicts the PMU and PDCs connected to IRIG-B signal and the servers connected to an NTP server which is receiving time synchronization from the substation GPS clock. In future diagrams, the detail of the IRIG and NTP connections will be omitted to improve clarity of the figures.

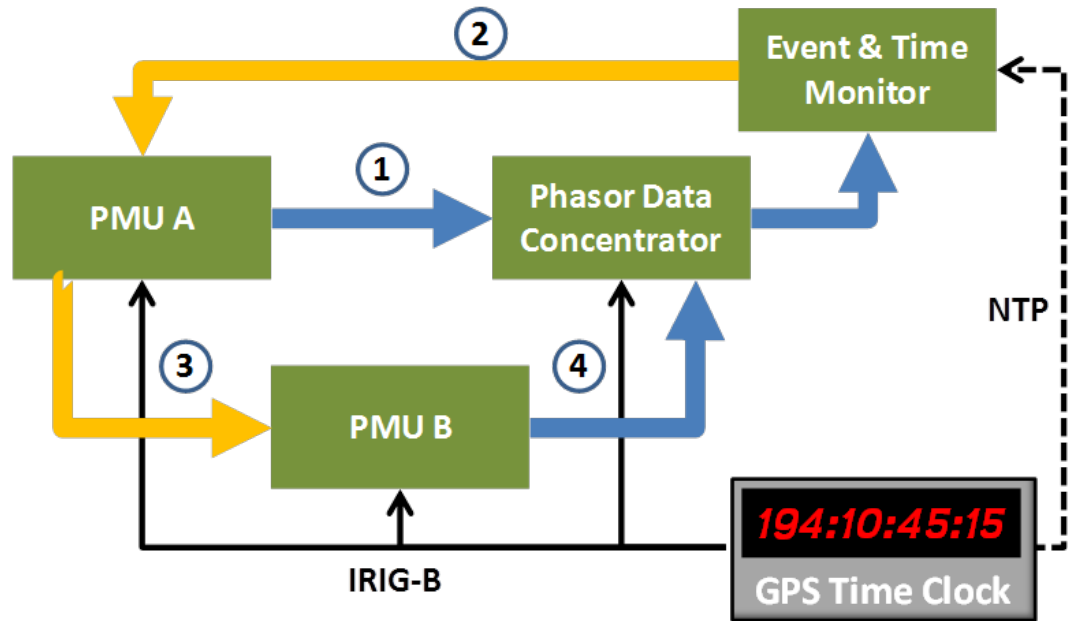


Figure 2-3
Time synchronization across the Test Bed Control Loop

The historian and subsequently the event monitor were used to capture and compare the timing of the various events that took place as an individual packet traversed the closed control loop. This was accomplished by extracting and recording the timestamp from the PMU IEEE C37.118 data stream, which remained unaltered by the historian or event monitor.

Intermediate Time References

Figure 2-4 identifies where the intermediate times that were captured during the testing sequence. The intermediate steps of the control loop are broken down and as follows:

- 1) Push-button on PMU A toggles a status change in Remote Bit 01 (RB01).
 - 1.1) PMU A's "RB01" status change is mirrored to PMU A digital bit PSV36.
- 2) The status change of PMU A's PSV36 is identified by the event monitor through the C37.118 protocol.
 - 2.1) Upon receipt, the event monitor sends a GOOSE message to PMU A.
- 3) PMU A receives this GOOSE message as Virtual Bit 01 (VB01).
 - 3.1) PMU A's "VB01" status change is mirrored to PMU A digital bit PSV24.
 - 3.2) PMU A sends a GOOSE message to PMU B.

- 4) PMU B receives this GOOSE message as Virtual Bit 01 (VB01).
- 4.1) PMU B's "VB01" status change is mirrored to PMU B digital bit PSV25.
- 5) The status change of PMU B's PSV25 is identified by the event monitor through the C37.118 protocol.

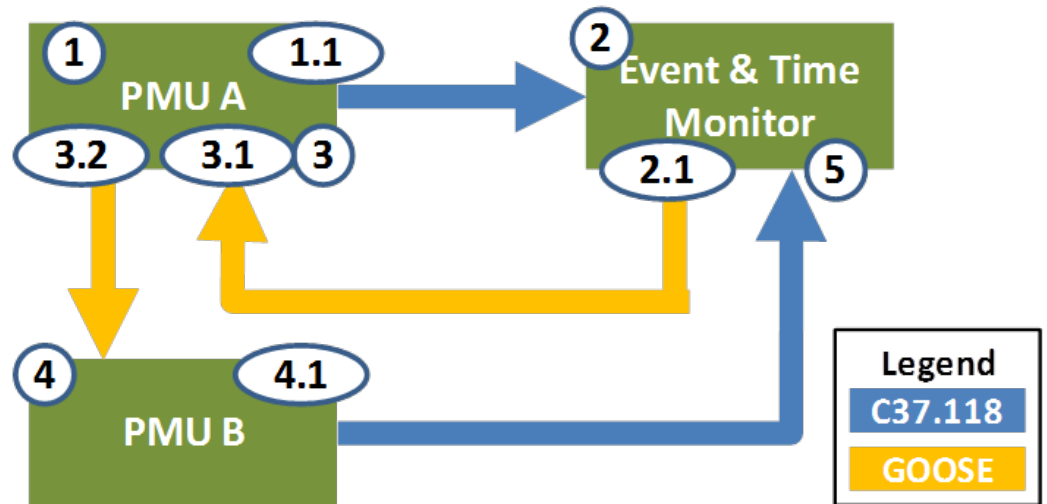


Figure 2-4
Intermediate Timing Step Identification

These are further illustrated in Table 2-1 which graphically shows the different sources of time data potentially available during this testing. The blue boxes represent an event time that appears in the sequence of events recorder (SER) from PMU A or PMU B or manual calculations available using the SER timestamp. The orange boxes represent an event time from the E&TM based on either data arrival time or GOOSE transmittal time. The yellow boxes represent the GPS timestamps from the PMUs and transmitted via C37.118 protocol.

Section 3: Testing Scenarios

Three transport scenarios were evaluated during this testing. These were:

1. IP Unicast with PDC
2. IP Unicast without PDC
3. IP Multicast

PMU sample rates of 60 samples/second and 30 samples/second were also tested as variables.

The test plan consisted of three tests with three configurations each, for a total of nine tests performed.

Test A was conducted with the default configuration of various sub-system components. The results were not as expected, and through support from the component vendors, a modified configuration was achieved. Tests B and C were then conducted using this configuration at 60 and 30 samples/second. Table 3-1 provides a summary of the tests conducted.

*Table 3-1
PMU Timing Test matrix*

Test	Sample Rate	Scenario	Sub-System Configuration
A	60 samples/sec	Unicast w/o PDC Unicast w/ PDC Multicast	Default configuration
B	60 samples/sec	Unicast w/o PDC Unicast w/ PDC Multicast	Modified configuration 1
C	30 samples/sec	Unicast w/o PDC Unicast w/ PDC Multicast	Modified configuration 1

When configuring PMU and PDC output connections, there are several transport options. For IP Unicast, TCP, UDP_U and UDP_T are available transport modes. For IP Multicast, UDP_S is the only setting compatible for

this mode¹. In our test environment, the PMUs and PDC supported all of these modes, as well as the historian that was receiving the PMU data. Presently, not all PMU devices or vendors presently support IP Multicast or UDP_S mode. For our tests, UDP_T was used for all IP unicast tests, and UDP_S was used for all IP Multicast tests.

Scenario 1 Description

In the first scenario, PMU A & B transmit IEEE C37.118 formatted data to the PDC using IP unicast with UDP_T mode, then to the historian before being received by the event monitor.

To enable this test, the PMU is reconfigured to transmit to the PDC in UDP_T mode and the PDC is configured to transmit to the historian in UDP_T mode. The historian is also configured to receive data from the PDC. This scenario is shown in Figure 3-1.

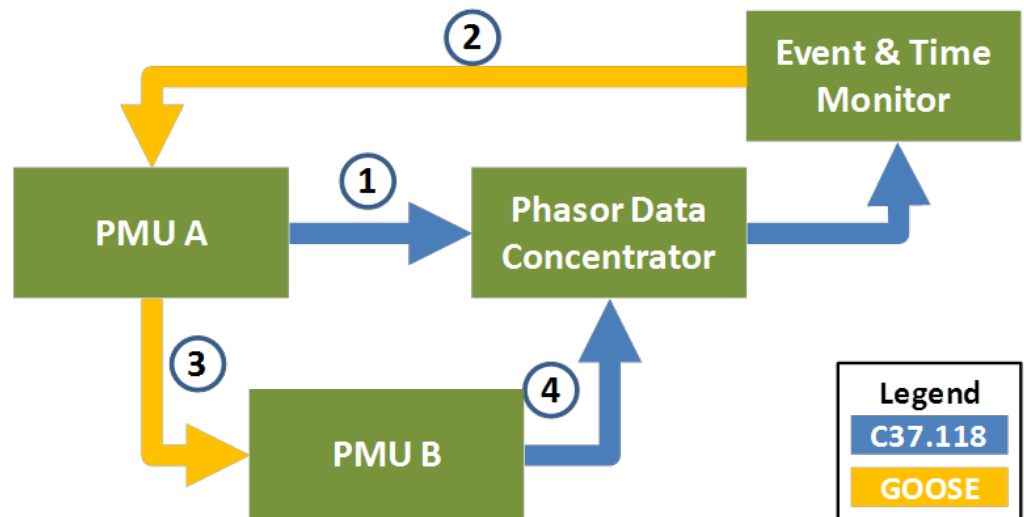


Figure 3-1
Scenario 1 configuration – IP Unicast with PDC

¹“...general guidance on using the various protocols. Use TCP to ensure that all data are received (e.g., for archiving). TCP uses handshaking that allows data to be retransmitted if lost in transit. TCP should not be used in high latency (>5 s) situations because disconnects and drops may occur. Use UDP, UDP_T, or UDP_S for applications where it is OK if some data are lost (e.g., visualization). UDP_S is used in applications where users want to ignore all incoming communications (e.g. used on the output of a substation PDC). UDP_S also enables the use of IP multicast. Use UDP_T or UDP_U in applications where you want TCP handshaking for the commands, and want minimal latency added to the data stream.” Source: SEL-3373 Station PDC Instruction Manual, Date Code 20121029.

Scenario 2 Description

In the second scenario, PMU A & B transmit IEEE C37.118 formatted data directly to the historian using IP unicast with UDP_T mode, before being received by the event monitor.

To enable this test, the PMU destination is reconfigured to transmit to the historian in UDP_T mode. The historian is then configured to receive data from the PMU in UDP_T mode. This scenario is shown in Figure 3-2.

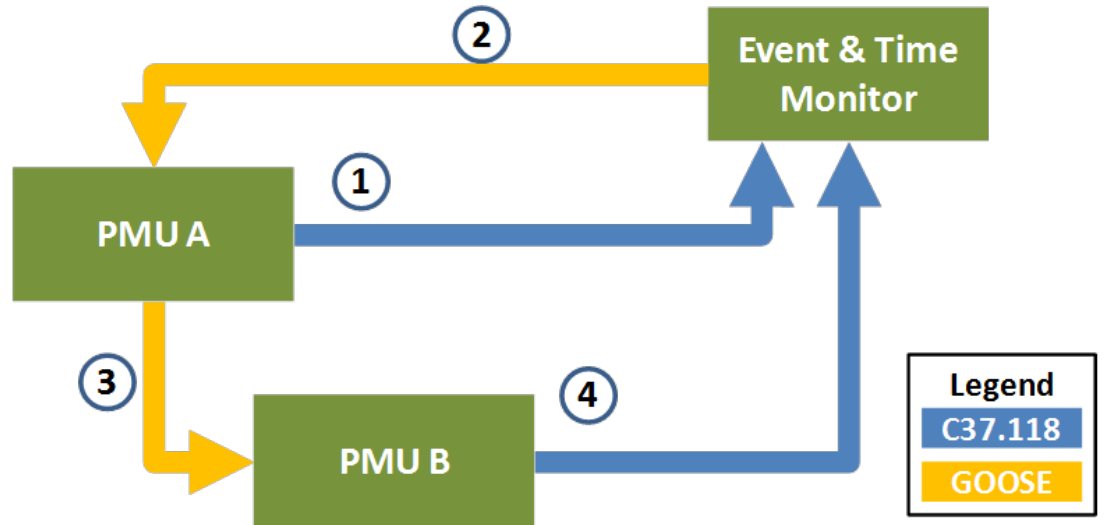


Figure 3-2
Scenario 2 configuration – IP Unicast without PDC

Scenario 3 Description

In the third scenario, PMU A & B transmit IEEE C37.118 formatted data directly to the historian using IP Multicast with UDP_S mode, before being received by the event monitor.

To enable this test, the PMU was reconfigured to transmit in IP Multicast (UDP_S mode) and the historian was reconfigured to receive using IP Multicast. This scenario is shown in Figure 3-3.

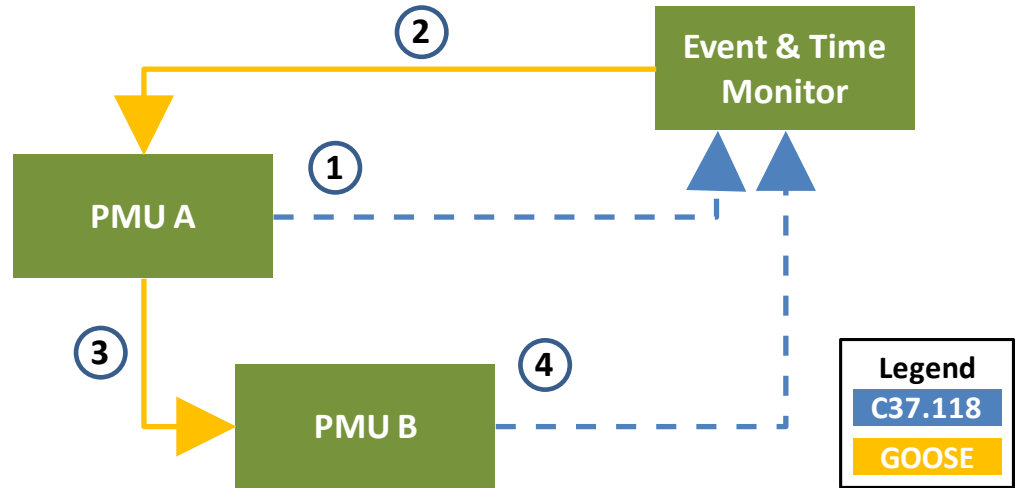


Figure 3-3
Scenario 3 configuration – IP Multicast



Section 4: Results Discussion

For Test A, timing tests were performed in each of the three configurations (IP Unicast with PDC, IP Unicast without PDC and IP Multicast). Due to the significant variance recorded during initial pre-testing, test sequences were repeated 60 times to increase the sample size and provide an improved confidence. Once completed, the Average, Maximum, Minimum, Standard Deviation and Mode were calculated for each configuration. Later when the configurations of the sub-system components were modified and the variation was reduced for Tests B and C, 60 samples were also obtained for consistency with Test A.

Latency measurements presented below are calculated as the timestamp of the control signal received at PMU B subtracted from the timestamp of the initiation of the control loop at PMU A. The Event and Time Monitor provided the monitoring and started the control scheme after receiving the data from the historian.

Test A Results

The chart in Figure 4-1 below shows the latency timing results from Test A that was repeated 60 times for each scenario. These tests yielded results that varied from less than 100ms to greater than 400ms.




Figure 4-1
Test A Results – 60 samples per second

Table 4-1
Test A Summary Statistical Results

	Multicast	Unicast w/ PDC	Unicast w/o PDC
Average	0.223	0.239	0.215
Maximum	0.450	0.416	0.333
Minimum	0.066	0.083	0.049
Standard Deviation	0.085	0.086	0.071
Mode	0.283	0.316	0.283

Impact of Variation Observed in Test A

Prior to conducting the tests, the expectation was for timing results to be in the range of 50-100ms. Since our test bed was relatively small and in a mostly controlled and isolated network environment, obtaining variations from 33 to 416ms was unexpected. For example, IP Unicast through a PDC had the same average as IP Unicast without a PDC, but the mode (value that occurred most often) was lower with a PDC than without. It was determined that there were

other network or other unknown delays/issues that were impeding the ability to get repeatable, accurate measurements with low variability.

To state another way: the variation in “unknown/network” was more significant than the measurement window expected to be received which was around 50-100ms at maximum.

If not resolved, the non-deterministic nature of the latency results observed in Test A could have an impact for use as a control mechanism.

This finding led to deeper investigation about system timing and potential sources impacting data latency. If not resolved or understood, the non-deterministic nature of the latency results could have a significant impact on use as a control mechanism. Across the three tests, the variations spanned from 3-5 cycles at the minimum time to 20-25 cycles at the maximum time. Under tests B & C, the maximum time was reduced 14-16 cycles. This upper range would still not be adequate in most transmission systems where the critical clearing time may be as short as 3-4 cycles. Even for other control actions within the transmission domain the latency could still not be fast enough.

The project team began to identify areas to reduce or eliminate the variations observed in Test A. As a result, multiple software buffering points were identified at various points along the testing data path in the test bed.

The configuration steps outlined greatly reduced the overall latency measurements and significantly reduced the variability. The results after these configuration changes were primarily in the 33-66 ms range (2-4 electrical cycles which will support some transmission system applications).

These configuration changes are described in more detail in Section 5.

Test B Results

Test B was performed and the results are presented in Figure 4-2. Test B was repeated 60 times for each scenario. These tests yielded results that were mostly less than 66 msec to some reaching 264 msec. The variations were significantly reduced, and the frequency of excursion was less.

- IP Multicast had the least amount of latency at 33ms
- IP Unicast (without PDC) measured either 33 or 66ms
- IP Unicast with PDC measured primarily 66ms latency.

These results were consistent with our expectations.

Test B results were more in line with expectations.




Figure 4-2
Test B Results – 60 samples per second

Table 4-2
Test B Summary Statistical Results

	Multicast	Unicast w/ PDC	Unicast w/o PDC
Average	0.057	0.072	0.059
Maximum	0.250	0.250	0.233
Minimum	0.016	0.033	0.033
Standard Deviation	0.055	0.051	0.043
Mode	0.033	0.050	0.050

Test C Results

Test C was performed with the same system configuration as Test B, however the PMU output rate was lowered from 60 to 30 samples per second.

This test was performed since 30 samples per second is a minimum sample rate often required by Independent System Operators (ISO) or other consumers of PMU data for the purpose of system response and analysis. This sample rate has become the de facto standard for PMU sample rates within North America by

the North American Synchrophasor Initiative and seems to provide the level of granularity needed to observe oscillations on the power grid and other grid behavior.

Reducing the sample rate from 60 to 30 samples per second reduces the volume of data and also reduces the amount of bandwidth required. Note that bandwidth utilization or utility bandwidth requirements were not examined during these tests.

- 60 samples/second = $1/60$ seconds = .0166 sec = 16.66ms sampling window for the PMU measurement.
- 30 samples/second = $1/30$ seconds = .0333 sec = 33.33ms sampling window for the PMU measurement.

So it would be anticipated that the results would be about 1-2 cycles longer (16-33ms) due to the PMU transmittal of the data at a slower rate.

Test C results in Figure 4-3 also displayed the reduced variability from Test A, and confirmed the expected latency impact with a slight shift to the right representing a one cycle increase in latency. This was due to the reduced sampling rate which increases the size of the sampling window.

- IP Multicast still had the least amount of latency at 33ms
- IP Unicast (without PDC) measured either 33 or 66ms
- IP Unicast with PDC measured primarily 66ms latency, with a slightly higher mode.

Overall, some occasional measurements had increased latency above 100ms level. Future efforts can identify the factors causing this and investigate ways to reduce these as much as possible. These test artifacts may not necessarily be trivial, and a delay as significant as 200ms could have a large impact on future control or other high frequency application that may develop that would be relying on this data stream.




Figure 4-3
Test C Results – 30 samples per second

Table 4-3
Test C Summary Statistical Results

	Multicast	Unicast w/ PDC	Unicast w/o PDC
Average	0.067	0.094	0.061
Maximum	0.233	0.333	0.233
Minimum	0.033	0.033	0.033
Standard Deviation	0.067	0.060	0.047
Mode	0.033	0.066	0.033

Test A, B, and C Summary Tables

A summary of Tests A, B and C results are presented below as Mode, Minimum and Maximum latencies. The actual test data is located the appendices.

Table 4-4
Summary Table of Test Results – Mode

Mode	Samples /Cycle	IP Unicast	IP Unicast w/PDC	IP Multicast	Configuration
Test A	60	.283	.100	.233	Initial Configuration
Test B	60	.033	.050	.050	Modified Config 1
Test C	30	.033	.066	.033	Modified Config 1

Brief analysis of Mode table:

- Test A – The mode of IP Unicast with PDC was less than the case of IP Unicast without PDC. This is not logical and indicates a problem.
- The latency for IP Multicast in test B was greater than test C by 16 ms (1 cycle). The anticipated results did not match the pattern shown by the other two scenarios. In those cases, test C had equal or slightly longer measurements than test B.

Table 4-5
Summary Table of Test Results – Minimum

Minimum	Samples /Cycle	IP Unicast	IP Unicast w/PDC	IP Multicast	Configuration
Test A	60	.049	.083	.066	Initial Configuration
Test B	60	.033	.050	.033	Modified Config 1
Test C	30	.033	.066	.033	Modified Config 1

Brief analysis of Minimum table

- Results were as expected. Unicast with PDC was 16ms (1 cycle) longer when using 30 samples/second.

Table 4-6
Summary Table of Test Results – Maximum

Maximum	Samples /Cycle	IP Unicast	IP Unicast w/PDC	IP Multicast	Configuration
Test A	60	.316	.416	.333	Initial Configuration
Test B	60	.233	.166	.233	Modified Config 1
Test C	30	.233	.266	.233	Modified Config 1

Note that the maximum values may be impacted by currently unknown or uncharacterized impacts such as bursts of network activity, which may affect the transmission of data.

General Test Observations

- Multicast and Unicast measurements were approximately the same.
 - Multicast was often faster by just a slim margin.
 - The test network was not complex enough to fully test or demonstrate the differences between the two approaches.
- It was anticipated that the use of a PDC, when compared to results without a PDC, would increase latency, since it requires another in-line devices to “hold and forward” PMU data. There is also a (slightly) longer network path and associated processing time used for this configuration. Test results from B and C varied only slightly with and without the PDC, and it was determined that the PDC added only a minimal delay, but timing with a PDC could be further impacted by other factors such as additional PMUs feeding into the PDC. The results were smaller than the sampling window for the PMU rate and fell between measurement reporting windows (33ms, 49ms, 66ms, etc).
- Uncharacterized delays may be significant and these could be considered in future areas of research discussed in more detail in Section 8.



Section 5: Sub-System Opportunities

There were a number of subsystems that were part of the test platform. These provided additional opportunities for timing improvement. These configuration opportunities identified fell in the following main areas:

1. Buffer between the historian PMU interface to the PDC or PMU
2. Buffers between the historian interface and the historian
3. PDC waiting period
4. Event and Time monitor configurations

Historian Interface Buffer

Within the historian PMU IEEE C37.118 interface, Figure 5-1 shows a “Flush Events” option. This was changed from “No” to “Yes” to force the interface to write the values to the historian as soon as they are received, and eliminate a delay with a variable buffer until the values are written. The exact impact was not characterized, but this modification appeared to have a significant impact on reducing variability. It was assumed that this buffer was variable, and not able to be fixed or modified, only disabled.

The description in the “hover-over” text describes this function as: “Specify whether or not to flush events after each processing loop.”

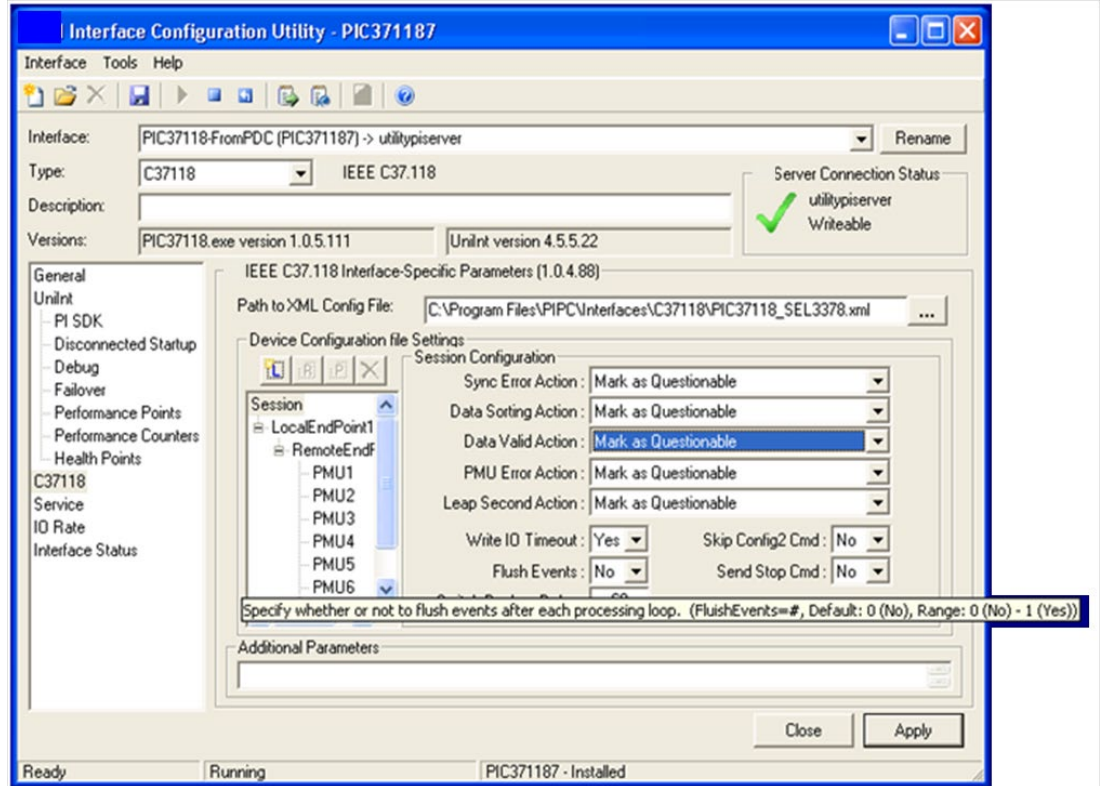


Figure 5-1
Historian Interface Buffer Configuration Screen

Buffer Between PMU Interface and Historian

There was also an additional buffer identified between the PMU interface and the historian. As shown in Figure 5-2, the “send rate” buffer had a default setting of 100. This was reduced to 1ms. Setting this buffer to 0ms would disable this buffer. Disabling the buffer did not appear to be a realistic implementation for a utility as disabling the buffer would have the effect of discarding data if a link to the historian was lost. Leaving the buffer at 1ms would still allow for a period of data to be buffered. However, this does imply that the system will ‘hold’ the data for 1ms, effectively increasing the transmission path delay by 1ms. The change from a default value of 100ms to 1ms had the impact of approximately 100ms reduction in delay. The relationship was not tested extensively to verify if it was linear.

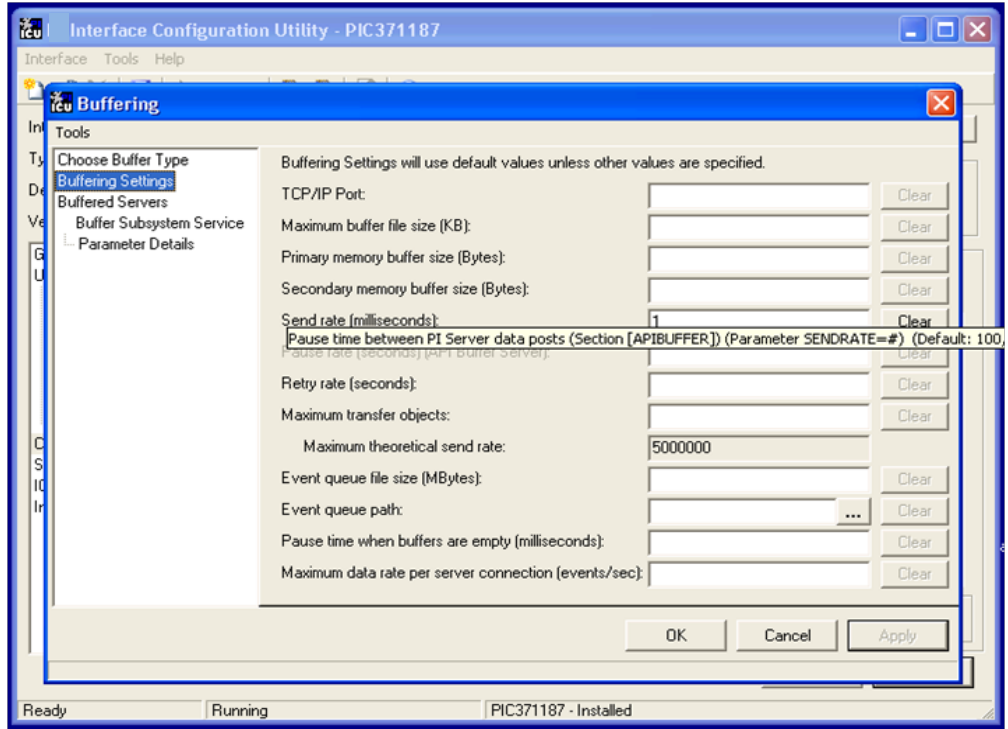


Figure 5-2
Historian Buffer Configuration Screen

PDC Waiting Period

The PDC waiting period was also modified from 1500ms, 200ms, 50ms, and 1ms. The tests were conducted at 1ms. Figure 5-3 shows a configuration screen of the PDC with the waiting period settings option.

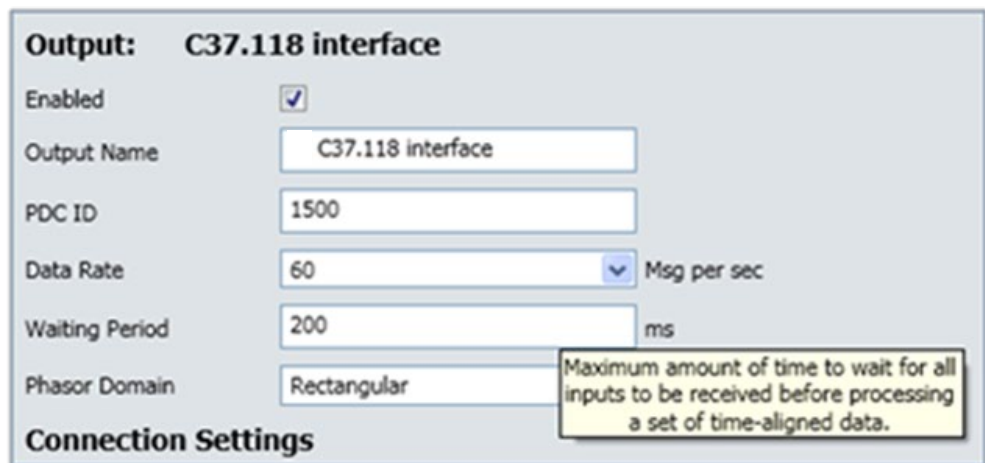


Figure 5-3
PDC Configuration Screen with waiting period settings

The description in the “hover-over” text describes this function as: “Maximum amount of time to wait for all inputs to be received before processing a set of time-aligned data.”

The number of incoming PMUs was also reduced from an initial configuration of seven PMUs to just include the two incoming streams from KNX (PMU A) and LNX (PMU B) substations to allow the PDC to forward the incoming packets as soon as both inputs were received. If other PMUs in the test bed were to send their data to the PDC and be included in the output stream, it is possible that the PDC would insert additional delays for the testing. This would be due to the fact that the PDC will wait up to this maximum time if another PMU has delays in its transmission to the PDC for any reason. Any PMU data received by a PDC after this waiting period is either discarded or transmitted but marked ‘invalid’.

In Figure 5-4 below, the data from the two PMUs were received within 4ms of each other. The key is that the ‘differences’ measurement is very small, so this is what the ‘time delay’ window will look for. Since these are coming in at the same time, they will be sent out by the PDC with very little delay.

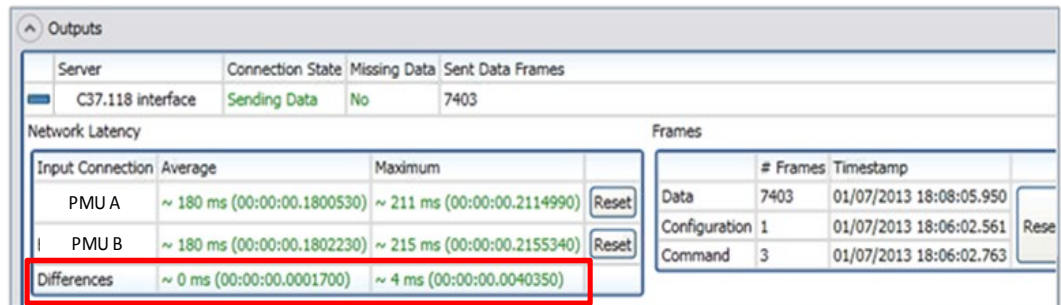


Figure 5-4
Example of actual PDC Timing Results

Event and Time Monitor

The Event and Time Monitor was also optimized to respond as soon as an incoming event changed. This was achieved by setting the “Max Sleep” time to 1ms. This allows the Event and Time Monitor to poll the historian every 1 ms to identify the change in state.

The number of incoming PMUs was also reduced from seven initially to just include the two incoming streams from KNX and LNX substations. This allowed the PDC to forward the incoming packets as soon as both inputs were received. In Figure 5-4 below, the data from the two PMUs were received within 4ms of each other. The key is that the ‘differences’ measurement is very small, so this is what the ‘time delay’ window will look for. Since these are coming in at the same time, they will be sent out by the PDC with very little delay.

Section 6: Future Time Synchronization Opportunities

When the historian and E&TM were configured to poll the NTP server every 15 minutes for time accuracy checks, the adjustment required to resynchronize the system time with NTP time proved to be unacceptable for testing, as the variation between 15 minute periods was often larger than the expected target measurement window. Further confirming that the use of the server time synchronized over NTP was unacceptable, in some tests comparing the PMU packet arrival time at the server, the packet would “arrive” before the “send” time at the PMU. Figure 6-1 depicts the output of an NTP time synchronization software with the adjustments required on a windows server.

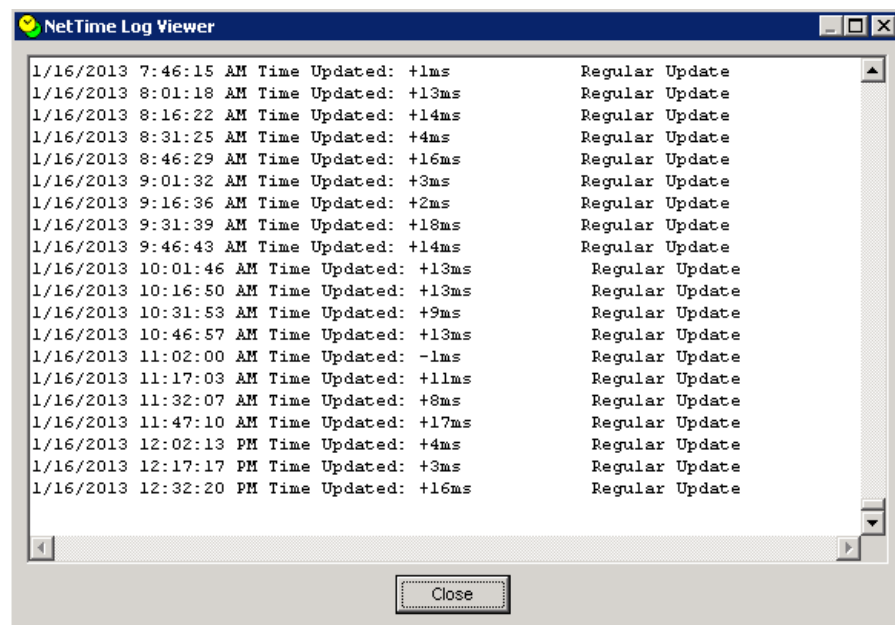


Figure 6-1
System time adjustments over 15 minute interval periods

Some of the variables identified as being responsible for the significant variance that appeared during the test configuration stage include the variability of using a non-real-time Operating System (OS). The variations were further compounded by the installation of the historian and the event monitor in a Virtual Machine

(VM) environment. The event monitor was later moved to run directly on the host OS because of the unbounded variability in the VM-based software preventing high-rate timing updates as well as the variability in the host OS when handling VM tasks. Variations observed were sometimes as high as 18 or 23ms. Variations were improved by relocation outside of a VM environment, but still not reduced enough for these tests examining timing measurements expected to be measured in increments of 16.66ms.

Opportunities to provide the desired synchronization and resolution across the PMUs and servers include: installing an IRIG-B card into the servers, examining IEEE 1588 hardware or software clients the servers. Migrating to a non-Windows client may reduce variability, and more updated versions of Windows OS (Server 2008 or Server 2013) are reported to be better with reducing variation in system time. The historian and event monitor currently only support Windows clients.

Further investigation is planned in the area of IEEE 1588, which can provide a higher level of timing accuracy and can adjust to delays in transmission of the time signal over the network. These adjustments are possible due to the negotiation occurring between grandmaster, transparent, boundary and ordinary clocks participating in a peer-to-peer network. Multiple masters are also allowed in the same network, with constant negotiations occurring to ensure that the most accurate master is always being used. As of the time of this report, none of the PMUs available in the EPRI lab provide support for IEEE 1588v2 protocol; however several vendors report to be working on this capability. Figure 6-2 depicts a future goal or IEEE 1588 connections to all test bed devices.

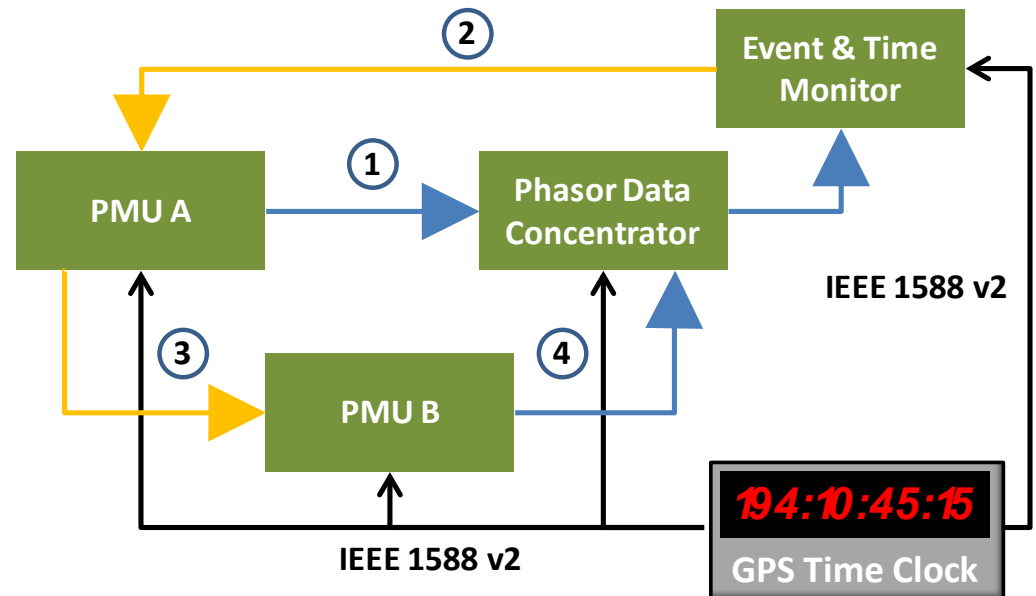


Figure 6-2
IEEE 1588v2 connections to all test bed devices

Section 7: Test Bed Network Connectivity Details

The test bed comprises two substations implementing a LAN at each station and a simulated interconnecting WAN link between the station routers. The WAN link can be manipulated through a packet modification tool (PMOD) to control latency, loss, out-of-order packets, jitter, propagation delay and other attributes. For analysis and troubleshooting, the TCPdump tool was used and logically connected to switches at both substations. Figure 7-1 below shows the network connectivity of the various components utilized in the test bed. The individual items in the test bed include two routers, three switches, the PMOD and TCPdump tools, two PMUs, a PDC, a historian and historian interface, and the event monitor.

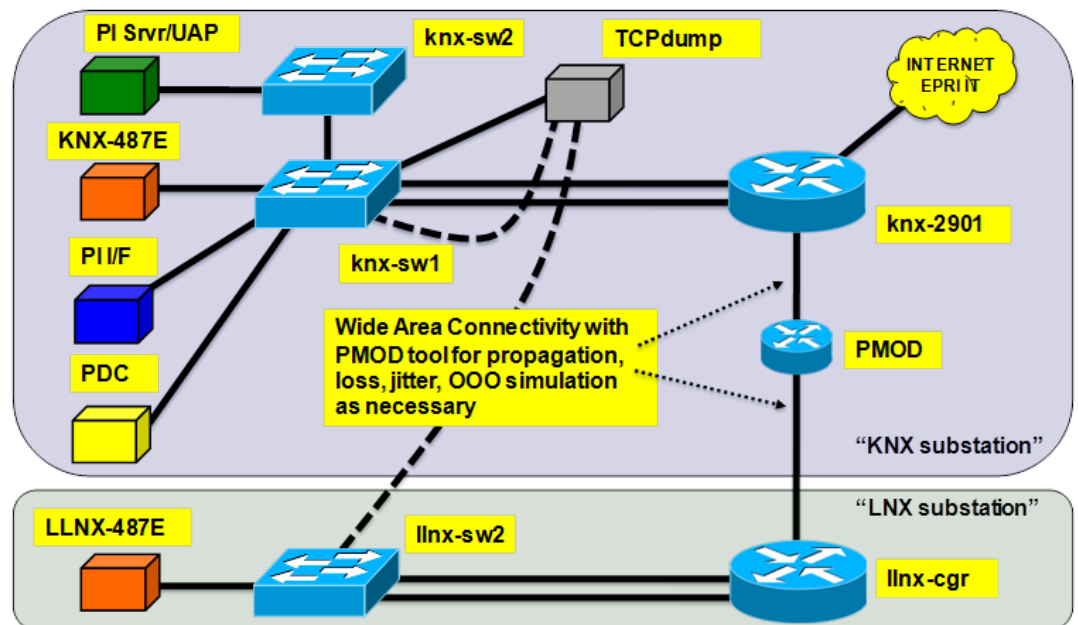


Figure 7-1
Test bed Network Diagram

Since it is a layer 2 protocol, GOOSE is typically LAN based and can only be transported within a single substation. However, in this test bed, we were able to transport GOOSE messages over a WAN (between substations). This was

accomplished by implementing Layer 2 Tunneling Protocol Version (L2TPv3) on the routers, in which Layer 2 Ethernet-encapsulated GOOSE messages were transported over the WAN within IP Packets. IEEE C37.118 IP Unicast and IP Multicast traffic were transported within the LANs and across the WAN link as IP packets.

GOOSE signaling which originated at PMU A and was destined for PMU B, was specifically marked with a VLAN ID of 4 to allow it to be differentiated from all other network traffic (marked as VLAN ID 1). This GOOSE traffic was then permitted to be transported over L2TPv3 between the substations, using the WAN link. All other traffic was filtered and not permitted to traverse the L2TPv3 tunnel between the substations.

L2TPv3 is an IETF standard related to L2TP that can be used as an alternative protocol to Multiprotocol Label Switching (MPLS). In computer networking, Layer 2 Tunneling Protocol (L2TP) is a tunneling protocol used to support virtual private networks (VPNs) or as part of the delivery of services by ISPs. It does not provide any encryption or confidentiality by itself; it relies on an encryption protocol that it passes within the tunnel to provide privacy.



Section 8: Conclusions

The measurement time-window of PMU data transmitted via IEEE C37.118 is 16.66ms at 60 samples/second and 33.33ms at 30 samples/second. In Tests B and C, the control loop measured latency was primarily in the 33-66 millisecond range. This corresponds to 2-4 electrical cycles, which is an acceptable response time for a for special protection scheme (SPS), or other scenarios where closed-loop control is desired. This matched the expected results for the controlled environment.


However, the results from Test A, using mostly default configuration settings from the vendor products, resulted in measurements averaging 225ms, or 14 electrical cycles. This would not have been acceptable to use for a control loop that required response in the under 50 millisecond range. It still would be acceptable for control applications that required second or minute accuracy.

The network architecture and design is an important consideration in developing any control system. Important design considerations include centralized vs. distributed architecture, data transmission distance, location of collection or automation engine relative to data source, and number of participating devices (PMUs, PDCs). Our testing was all based on commercially available equipment.

Our test bed architecture included routing data through a historian before reaching the event and time manager. For a protection scheme based on PMU data, the architecture would likely not include passing data through a historian first. An alternate architecture could include a dedicated, real-time calculation/automation platform that is directly receiving and acting on the data. This platform would receive a data feed directly from the PDC.

Significant improvements in the control loop latency were achieved through application of component level expertise and trial and error involving all components of the test bed. This included configuration of device settings, network traffic control, and software interfaces among the most important. Some specific issues include the following:

- Network Traffic configuration utilizing VLANs to differentiate traffic within and between the substations.
- Understanding of the effects of network propagation delays
- Proper set up of buffers and subsystems in various software components



Understanding the many intricacies of each piece of equipment required domain experts.

- Understanding the affects of PDCs on the system latency, with respect to their “wait windows”. Multiple layers of PDCs could exacerbate the impact.
- There were some unexplained latency behavior components which were not clearly understood (variation or trends over longer periods of time than the testing period).
- Impact of non real-time operating systems, including the use of Virtual Machines.

Understanding many intricacies of each equipment component required experts to be available during testing.

The configuration of each of the pieces of equipment minimize latency was a somewhat laborious process which could possibly be prone to error if done in a wide-spread fashion. It might be prudent to develop an automated coordination system which, given the parameters and specific goals of a defined control loop, would configure and optimize the equipment participating in the control loop. Distributed components, such as substation PDC configurations, event and time monitors, or historian interfaces are a few examples of components that could be very numerous throughout the utility environment. A short term approach would be for the default setting to be changed to the utility standard defined.

It was very difficult to measure the individual latency components of the test bed. This was due to the inability to collect and record the time at all critical test bed components with a single synchronized time reference. Lack of the universal time source with the same clock accuracy also hindered more detailed timing collection. However, the goal of measuring the timing of the overall control loop activity was able to be captured allowing us to be able to evaluate the performance of the system as a whole.

While the results indicate a similarity in latency performance for IP Multicast and IP Unicast (without an intermediate PDC), there are some benefits to using IP multicast that need to be considered. First of all, IP multicast utilizes a subscription model permitting any client to request and subsequently receive PMU data from any IP multicast feed it chooses. Theoretically, this could permit a very large number of clients to receive this data for a very small overhead tradeoff to make this happen. IP Unicast, on the other hand requires significantly more overhead to support each client. But more importantly, since IP Unicast it does not implement and integrated subscription, it requires establishment and management of any connection which takes place between a client and a PMU.

Test Bed Limitations / Observations

The section describes some of the limitations of the test bed and various observations made during the actual testing.

- Not all of the network impacts were identified and resolved. However, the latency variation was significantly reduced to levels that Tests B and C were able to be performed and results evaluated.

The root cause of the “bursts of network activity” seen during the tests was not determined. One possible solution is to create network isolation through different VLANs for PMU, GOOSE, and other network traffic. See Section 7: Test Bed Network Connectivity Details on network details.

- Timing Synchronization Issues

The IRIG-B timing at the PMU was used as the single reference to provide results. During the testing we were not able to break down times within the sequence without synchronization. See Section 6: Future Time Synchronization Opportunities on timing issues and future plans.

- A single vendor PMU and PDC were used

In the current testing only a single vendor’s hardware was used and while many utilities do concentrate on a single vendor, that may not always be representative of a utility environment. Future tests could use other vendors if multicast is supported. It would be useful though to use the same internal components with good time sync across these devices.

- Bandwidth usage was not examined

The testing did not measure actual bandwidth usage to determine loading of the network. This could have an impact on the amount that utilities pay for their bandwidth.

- The test network is not complex enough

The network used was not complex enough to determine how PDC layering for example may actually contribute to the overall latency. Also the number of PMUs producing data was kept to a minimum. Future testing should scale up the complexity (expand to use more physical diversity) and the data volume. This would help to highlight differences between Unicast and Multicast.

Future Testing Areas

The following areas may be investigated in future work:

- IEC 61850-90-5

Analysis of latency timing differences using IEC 61850-90-5, when compared to unicast and multicast can be performed. This would be helpful to explain differences from multicast and unicast and identify the benefits / drawbacks of the approach.

- Security overlay on WAN link

Lack of security in C37.118 protocol implies that security needs to be considered to protect against manipulation or modification of the data in transit. It could help to identify the effects of encryption/decryption on latency.

- Additional network traffic/loading and additional PMUs can be added

This would be helpful to determine the impact to delays at PDC or network. Concentration of 2 PMUs vs. 10 PMUs may be impacted by PDC 'delay time'. Tests with more than 2 PMUs can be performed under this area.

- Network Configuration

Determining the impact of various network configuration and their impact on WAN propagation delays, network impairment / availability, traffic prioritization / isolation and the use of Vlans to minimize / isolate traffic.

- IEEE 1588 / PTP synchronization to windows servers

The ability to take timing measurements using routers, switches or windows servers to break down overall 'trip time' and analyze which components hold up the data the longest/shortest. This would allow the greatest component of total timing duration through the loop.

- Expanding the variety of vendor hardware and software included

The addition of alternate vendor PMUs and intermediate hardware PDCs, event and time monitors, historians, can increase the knowledge of other settings etc will all add to the variability and settings relevant to the overall system optimization with various solutions.

Appendix A: Test A Raw Data

The following table contains the raw measurement data from the configuration used for test A described earlier in this report. In total 60 test runs were performed. It is followed by a summary table of statistical calculations based on the data.

Table A-1
Test A Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
1	0.300	0.183	0.333
2	0.300	0.166	0.266
3	0.133	0.083	0.100
4	0.283	0.316	0.233
5	0.250	0.166	0.200
6	0.266	0.316	0.299
7	0.166	0.150	0.199
8	0.149	0.316	0.199
9	0.066	0.333	0.100
10	0.183	0.300	0.216
11	0.183	0.166	0.266
12	0.066	0.383	0.233
13	0.233	0.283	0.216
14	0.283	0.233	0.250
15	0.300	0.233	0.316
16	0.283	0.316	0.199
17	0.150	0.250	0.149
18	0.183	0.183	0.283
19	0.250	0.300	0.149
20	0.199	0.283	0.216

Table A-1 (continued)
 Test A Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
21	0.283	0.283	0.250
22	0.233	0.316	0.083
23	0.149	0.150	0.183
24	0.283	0.266	0.283
25	0.133	0.183	0.149
26	0.266	0.133	0.116
27	0.199	0.250	0.183
28	0.283	0.166	0.200
29	0.133	0.166	0.166
30	0.166	0.266	0.133
31	0.250	0.316	0.200
32	0.183	0.199	0.283
33	0.416	0.116	0.133
34	0.366	0.133	0.250
35	0.250	0.233	0.150
36	0.350	0.383	0.250
37	0.200	0.150	0.100
38	0.383	0.283	0.300
39	0.133	0.316	0.233
40	0.450	0.366	0.216
41	0.066	0.283	0.200
42	0.266	0.100	0.299
43	0.250	0.316	0.266
44	0.183	0.250	0.133
45	0.266	0.316	0.300
46	0.233	0.250	0.200
47	0.216	0.100	0.283
48	0.083	0.350	0.266
49	0.199	0.166	0.083
50	0.250	0.233	0.250

Table A-1 (continued)
 Test A Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
51	0.233	0.416	0.283
52	0.133	0.100	0.316
53	0.233	0.083	0.049
54	0.066	0.133	0.300
55	0.283	0.333	0.100
56	0.116	0.283	0.300
57	0.216	0.299	0.283
58	0.266	0.333	0.183
59	0.149	0.183	0.283
60	0.333	0.166	0.216

Table A-2
 Summary Statistics for Test A

	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
Average	0.223	0.239	0.215
Maximum	0.450	0.416	0.333
Minimum	0.066	0.083	0.049
Standard Deviation	0.085	0.086	0.071
Mode	0.283	0.316	0.283




Figure A-1
 Test A Frequency of Occurrence

Appendix B: Test B Raw Data

The following table contains the raw measurement data from the configuration used for test B described earlier in this report. In total 60 test runs were performed. It is followed by a summary table of statistical calculations based on the data.

Table B-1
Test B Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
1	0.033	0.050	0.033
2	0.033	0.050	0.033
3	0.033	0.050	0.033
4	0.033	0.100	0.050
5	0.033	0.083	0.033
6	0.050	0.066	0.049
7	0.033	0.049	0.033
8	0.033	0.049	0.033
9	0.050	0.049	0.033
10	0.033	0.233	0.033
11	0.016	0.066	0.066
12	0.033	0.066	0.033
13	0.033	0.049	0.050
14	0.033	0.050	0.050
15	0.033	0.050	0.050
16	0.033	0.066	0.050
17	0.116	0.049	0.050
18	0.183	0.066	0.066
19	0.033	0.066	0.049
20	0.233	0.050	0.066

Table B-1 (continued)
 Test B Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
21	0.199	0.066	0.050
22	0.250	0.050	0.049
23	0.133	0.066	0.050
24	0.083	0.050	0.033
25	0.033	0.050	0.033
26	0.133	0.049	0.050
27	0.083	0.050	0.050
28	0.033	0.050	0.166
29	0.033	0.066	0.050
30	0.033	0.050	0.166
31	0.050	0.149	0.050
32	0.050	0.250	0.033
33	0.033	0.250	0.050
34	0.033	0.050	0.050
35	0.033	0.066	0.100
36	0.066	0.033	0.049
37	0.033	0.033	0.066
38	0.033	0.066	0.050
39	0.033	0.050	0.050
40	0.033	0.049	0.049
41	0.033	0.050	0.050
42	0.033	0.233	0.033
43	0.033	0.066	0.033
44	0.033	0.050	0.033
45	0.050	0.050	0.050
46	0.033	0.050	0.033
47	0.033	0.050	0.033
48	0.233	0.166	0.033
49	0.033	0.050	0.199
50	0.033	0.050	0.050

Table B-1 (continued)
 Test B Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
51	0.033	0.066	0.050
52	0.033	0.066	0.049
53	0.033	0.050	0.166
54	0.033	0.066	0.033
55	0.033	0.066	0.050
56	0.033	0.050	0.133
57	0.083	0.050	0.033
58	0.033	0.050	0.050
59	0.033	0.066	0.233
60	0.033	0.050	0.050

Table B-2
 Summary Statistics for Test B

	System Latency Relay Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
Average	0.057	0.072	0.059
Maximum	0.250	0.250	0.233
Minimum	0.016	0.033	0.033
Standard Deviation	0.055	0.051	0.043
Mode	0.033	0.050	0.050




Figure B-1
 Test B Frequency of Occurrence

Appendix C: Test C Raw Data

The following table contains the raw measurement data from the configuration used for test C described earlier in this report. In total 60 test runs were performed. It is followed by a summary table of statistical calculations based on the data.

Table C-1
Test C Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
1	0.033	0.066	0.066
2	0.033	0.066	0.066
3	0.033	0.066	0.033
4	0.033	0.200	0.066
5	0.033	0.066	0.066
6	0.133	0.133	0.033
7	0.233	0.066	0.033
8	0.133	0.066	0.199
9	0.166	0.066	0.033
10	0.033	0.066	0.033
11	0.033	0.200	0.233
12	0.033	0.066	0.033
13	0.033	0.066	0.066
14	0.033	0.066	0.066
15	0.033	0.199	0.066
16	0.233	0.066	0.066
17	0.033	0.066	0.033
18	0.033	0.066	0.066
19	0.133	0.066	0.066
20	0.033	0.066	0.033

Table C-1 (continued)
 Test C Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
21	0.233	0.066	0.033
22	0.033	0.099	0.066
23	0.066	0.066	0.033
24	0.033	0.133	0.033
25	0.033	0.066	0.066
26	0.033	0.066	0.033
27	0.033	0.066	0.033
28	0.033	0.066	0.033
29	0.033	0.066	0.033
30	0.200	0.066	0.100
31	0.066	0.033	0.066
32	0.033	0.166	0.033
33	0.033	0.066	0.133
34	0.033	0.066	0.066
35	0.033	0.066	0.033
36	0.033	0.066	0.033
37	0.033	0.066	0.066
38	0.033	0.066	0.066
39	0.066	0.333	0.033
40	0.033	0.066	0.033
41	0.166	0.066	0.166
42	0.233	0.099	0.066
43	0.033	0.100	0.033
44	0.033	0.199	0.033
45	0.033	0.066	0.033
46	0.033	0.066	0.033
47	0.033	0.066	0.033
48	0.033	0.133	0.033
49	0.033	0.066	0.033
50	0.033	0.266	0.033

Table C-1 (continued)
 Test C Raw Data

Run	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
51	0.033	0.066	0.033
52	0.033	0.066	0.066
53	0.233	0.066	0.033
54	0.033	0.266	0.033
55	0.033	0.066	0.033
56	0.233	0.066	0.100
57	0.033	0.100	0.133
58	0.033	0.166	0.066
59	0.033	0.066	0.233
60	0.033	0.066	0.066

Table C-2
 Summary Statistics for Test C

	System Latency Time		
	Multicast	Unicast w/ PDC	Unicast w/o PDC
Average	0.067	0.094	0.061
Maximum	0.233	0.333	0.233
Minimum	0.033	0.033	0.033
Standard Deviation	0.067	0.060	0.047
Mode	0.033	0.066	0.033




Figure C-1
Test C Frequency of Occurrence

Export Control Restrictions

Access to and use of EPRI Intellectual Property is granted with the specific understanding and requirement that responsibility for ensuring full compliance with all applicable U.S. and foreign export laws and regulations is being undertaken by you and your company. This includes an obligation to ensure that any individual receiving access hereunder who is not a U.S. citizen or permanent U.S. resident is permitted access under applicable U.S. and foreign export laws and regulations. In the event you are uncertain whether you or your company may lawfully obtain access to this EPRI Intellectual Property, you acknowledge that it is your obligation to consult with your company's legal counsel to determine whether this access is lawful. Although EPRI may make available on a case-by-case basis an informal assessment of the applicable U.S. export classification for specific EPRI Intellectual Property, you and your company acknowledge that this assessment is solely for informational purposes and not for reliance purposes. You and your company acknowledge that it is still the obligation of you and your company to make your own assessment of the applicable U.S. export classification and ensure compliance accordingly. You and your company understand and acknowledge your obligations to make a prompt report to EPRI and the appropriate authorities regarding any access to or use of EPRI Intellectual Property hereunder that may be in violation of applicable U.S. or foreign export laws or regulations.

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI's members represent approximately 90 percent of the electricity generated and delivered in the United States, and international participation extends to more than 30 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; and Lenox, Mass.

Together...Shaping the Future of Electricity

© 2013 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

3002000604