

TI Designs: TIDEP-0103

Linux® でのサブステーション・オートメーション用の並列冗長性プロトコル・イーサネットのリファレンス・デザイン



概要

この Linux® ベースのリファレンス・デザインでは、スマート・グリッド送配電網ネットワークにおけるサブステーション・オートメーション機器用の、高信頼性でレイテンシの短いネットワーク通信を紹介します。このリファレンス・デザインは、IEC 62439 規格の PRP (並列冗長性プロトコル) 仕様と、IEEE 1588 の PTP (高精度タイム・プロトコル) 仕様をサポートしています。この独自手法は、FPGA を使用する手法に比べて低コストで柔軟性と性能が高く、追加コンポーネントなしで IEC 61850 サポートなどの機能を追加できます。

リソース

TIDEP-0103	デザイン・フォルダ
AM572x	プロダクト・フォルダ
AM437x	プロダクト・フォルダ
AM335x	プロダクト・フォルダ
AM572x IDK	ツール・フォルダ
AM437x IDK	ツール・フォルダ
AM335x ICE	ツール・フォルダ
プロセッサ SDK	ツール・フォルダ




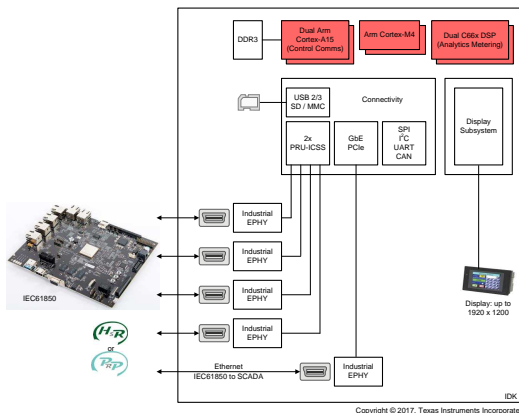
[E2E エキスパートに質問](#)

特長

- PRP イーサネット通信の IEC 62439-3 Clause 6 仕様に準拠
- ポート A とポート B の 2 ポートで DANP として動作
- 仮想ローカル・エリア・ネットワーク (VLAN) ID、マルチキャストおよびブロードキャストのサポート、内蔵のストーム防止および監督機構によるトラフィックのフィルタリング
- VLAN PCP による QoS (サービス品質) で 8 レベルをサポート
- IEEE 1588 の通常クロックおよびピアツーピア透過型クロック・プロファイルによるネットワーク同期
- IEEE 1588 の境界クロックによるネットワーク同期 (AM57x のみ)
- ネットワーク障害時に復旧時間が不要
- デュアル・ポート、全二重の 100Mbps イーサネット
- AM335x、AM437x、AM57x デバイス・ファミリをサポート
- 完全にプログラム可能なソリューションで、Linux ベースの追加アプリケーションの統合用プラットフォームとして機能

アプリケーション

- グリッド・オートメーション
- 保護リレー
- スマート・グリッド通信



使用許可、知的財産、その他免責事項は、最終ページにあるIMPORTANT NOTICE (重要な注意事項)をご参照くださいますようお願いいたします。

1 System Description

A substation is a key component of the electricity grid infrastructure, which is located everywhere from power generation facilities throughout the distribution network to the low-voltage feeders serving residences and businesses. Substations play a key role in transforming voltage levels for transmission and performing important functions, such as switching, monitoring, and protecting subsystems, to maintain grid efficiency and reliability. Traditional substation systems focused on fault monitoring, which can be manually fixed by switching to backup subsystems.

Consumers, regulators, and grid operators demand ever-greater reliability of electricity delivery. Automatic switching and protecting subsystems can require substations to start automating operations and communications such as monitoring grid conditions and communicating that information to grid operators reliably and rapidly. Operators must continually monitor the health of their network and take action to maintain high-speed operation. This leads to the requirement for reliable and low-latency communications between the operator’s control center and high-value nodes such as substations.

The International Electro-technical Commission (IEC) released specifications for industrial Ethernet communications under the IEC 62439 standard. The PRP specification is a static redundancy, Ethernet-based protocol, which supports critical real-time systems that require continuous monitoring.

The IEEE 1588 PTP is designed to provide high-accuracy network time synchronization between subsystems.

This design provides a reliable, high-speed, PRP communication solution for substation automation. This design implements PRP-compliant with IEC 62439-3 Clause 6.

This reference design is an affordable alternative to ASIC or FPGA-based Ethernet solutions while delivering equivalent performance. The programmable nature of the solution allows operating different redundancy Ethernet protocols without modifying hardware and adding applications, such as IEC 61850, without requiring extra system cost.

図 1 shows the overall system architecture. The PRP supports dual-port, full-duplex, Ethernet communication between network devices. The system includes Ethernet PHY as layer one and the Network Stack with PRP as the upper layers. PRP capability can be demonstrated simply with standard Linux applications like ping. RT Linux priorities can be used to meet the requirements of hard, real-time applications to minimize latencies. Standard TCP/IP traffic can use the well-recognized Linux network stack.

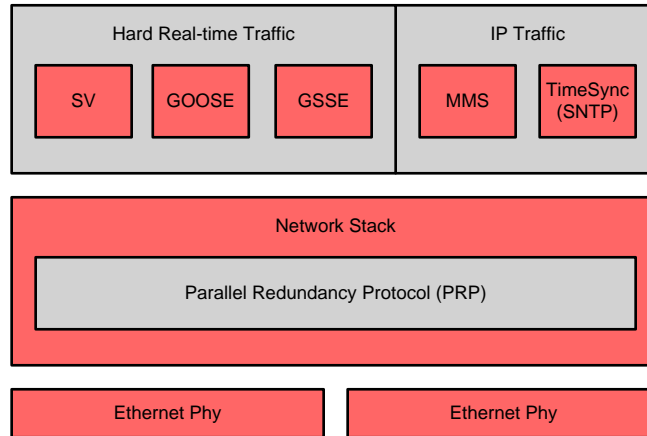


図 1. System Architecture

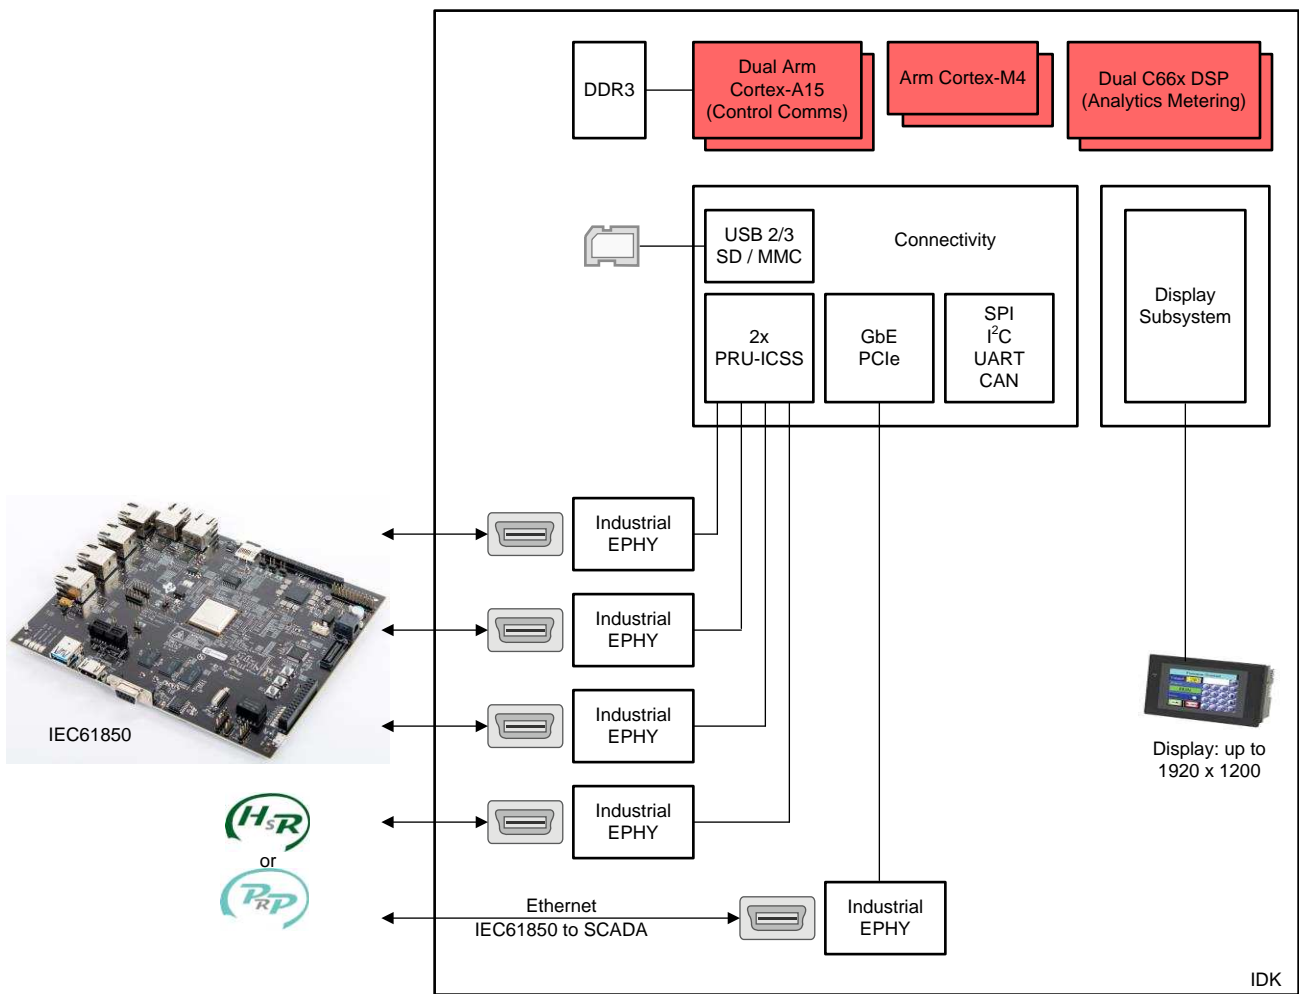
2 System Overview

2.1 Block Diagram

Figure 2 shows the block diagram. The primary device for this design is the dual-core AM5728 ARM® Cortex®-A15 microprocessor (MPU) as the host processor to support PRP and user applications running under a real-time Linux environment. This support also scales across the AM335x and AM437x device families as well.

This design uses this highly-integrated device for these benefits:

- The programmable real-time unit and industrial communication subsystem (PRU-ICSS) allows independent operation for real-time communication stacks.
- The high-performance ARM cores support the real-time applications for substation automation.
- The programmable, flexible software design allows upgrades to different Ethernet-based redundancy protocols without hardware modification.
- The C66x DSP on certain devices allows the device to extend to the Network Edge solution that requires local data analytics. This can reduce the backbone traffic greatly by not sending raw samples.



Copyright © 2017, Texas Instruments Incorporated

Figure 2. TIDEP-0103 Block Diagram

2.2 Highlighted Products

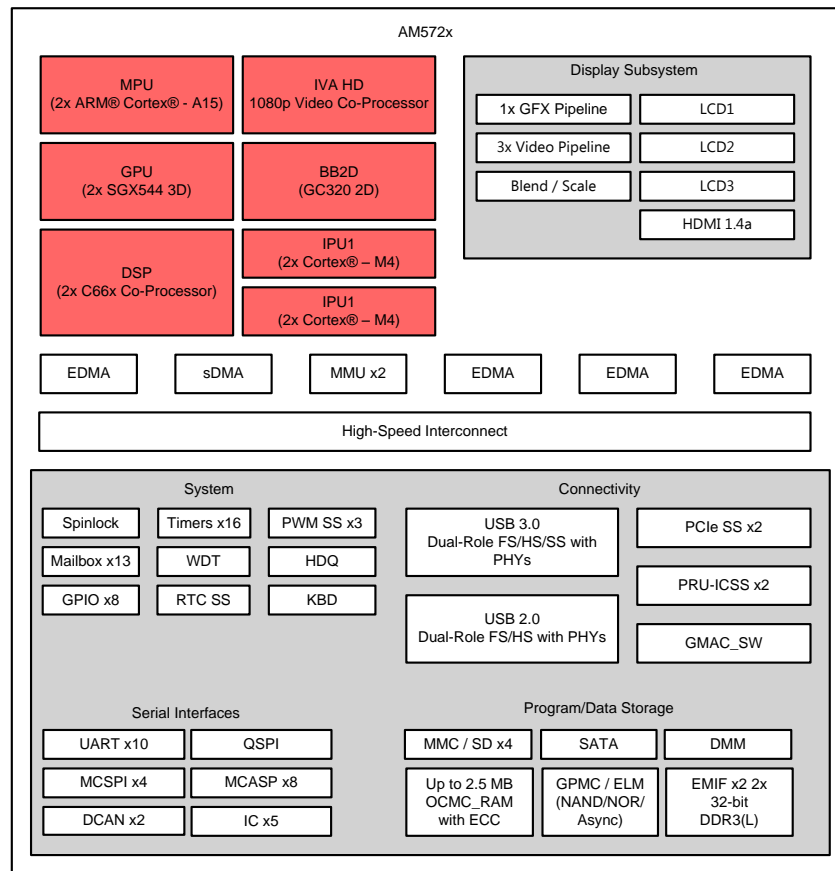
2.2.1 TI AM572x Sitara™ Processor Overview

The AM572x Sitara processor brings high-processing performance through the maximum flexibility of a fully-integrated, mixed processor solution. The devices also combine programmable video processing with a broad and highly-integrated peripheral set, which is well suited for industrial applications.

Programmability is provided by dual-core, ARM Cortex-A15 RISC CPUs with ARM NEON™ technology and two TI C66x VLIW, floating-point, digital signal processor (DSP) cores. The ARM Cortex lets developers keep separate control functions from other algorithms that are programmed on the DSPs and coprocessors. The separated control functions reduce the complexity of the system software. The ARM Cortex-A15 CPU supports multiple operating frequencies at a range of up to 1.5 GHz.

The AM572x processor is configured with two, dual-core PRU-ICSSs. The PRU-ICSS can be used for communication protocols, such as EtherCAT® master and slave, PROFINET®, Ethernet/IPTM, SERCOS®, HSR, PRP and so forth. The PRU-ICSS is separate from the ARM core, which allows for independent operation and clocking for greater efficiency and flexibility. The PRU-ICSS unit contains two PRUs, each of which includes 32-bit RISC processor capable of running at 200 MHz, that support real-time protocol for HSR and additional interfaces of media independent interface (MII) and reduced media independent interface (RMII) to connect to the Ethernet PHY devices directly.

Additionally, the programmable nature of the PRU-ICSS, along with its access to pins, events, and all system-on-chip (SoC) resources, provides flexibility in implementing fast, real-time responses, and specialized data handling.



Copyright © 2017, Texas Instruments Incorporated

図 3. AM5728 Functional Block Diagram

2.2.2 AM437x

The TI AM437x high-performance processors are based on the ARM Cortex-A9 core. The processors each provide a rich graphical user interface (GUI). The AM437x has PRU-ICSS coprocessors for deterministic, real-time processing including industrial communication protocols, such as EtherCAT®, PROFIBUS®, HSR, PRP and others as well as industrial drive protocols such as EnDat, Tamagawa, Sigma Delta, and so forth. Other Operating Systems are available from TI's Design Network and ecosystem partners.


These devices offer an upgrade to systems based on lower performance ARM cores and provide updated peripherals, including memory options such as QSPI-NOR.

High-performance interconnects provide high-bandwidth data transfers for multiple initiators to the internal and external memory controllers and to on-chip peripherals. The device also offers a comprehensive clock-management scheme.

One on-chip analog-to-digital converter (ADC0) can couple with the display subsystem to provide an integrated touch-screen solution. The other ADC (ADC1) can combine with the pulse width module to create a closed-loop motor control solution.

The RTC provides a clock reference on a separate power domain. The clock reference enables a battery-backed clock reference. The camera interface offers configuration for a single- or dual-camera parallel port. Cryptographic acceleration is available in every AM437x device. Secure boot is available only on AM437xHS devices for anticloning and illegal software update protection.

Figure 4 shows the AM437x block diagram.



Copyright © 2017, Texas Instruments Incorporated

Figure 4. AM437x Block Diagram

2.2.3 AM3359

The AM3359 microprocessors (based on the ARM Cortex-A8 processor) are enhanced with image, graphics processing, peripherals, and an industrial-interface option for PRP and high-availability seamless redundancy (HSR). The PRU-ICSS is separate from the ARM core, allowing for independent operation and clocking for increased efficiency and flexibility. The PRU-ICSS unit contains two PRUs; each PRU includes a 32-bit RISC processor capable of running at 200 MHz that supports real-time protocol for PRP and HSR. The PRUs support additional media-independent interfaces (MIIs) and reduced media-independent interfaces (RMII) to connect to the Ethernet PHY devices directly.

The programmable nature of the PRU-ICSS, the access to pins, events, and all system-on-chip (SoC) resources provide flexibility when implementing fast, real-time, specialized-data handling (see [Figure 5](#)).

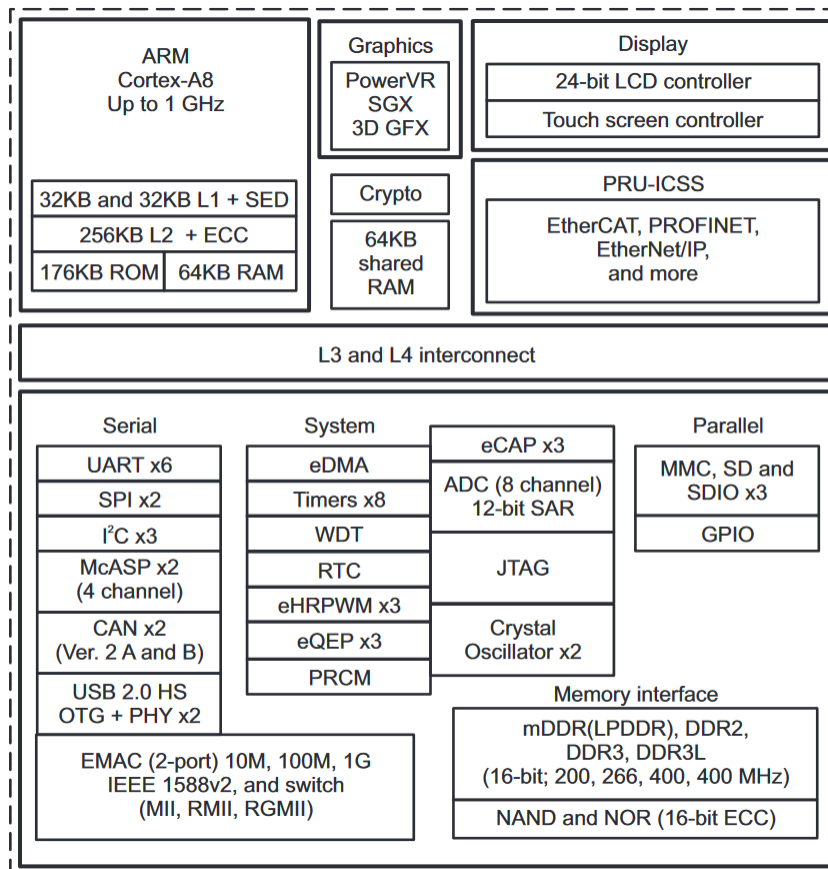


図 5. AM3359 Functional Block Diagram

2.2.4 PRU-ICSS

The Programmable Real-Time Unit and Industrial Communication Subsystem (PRU-ICSS) is separate from the ARM core and allows independent operation and clocking for greater efficiency and flexibility. The PRU-ICSS enables additional peripheral interfaces and real-time protocols such as EtherCAT, PROFINET IRT®, EtherNet/IP™, PROFIBUS, Ethernet POWERLINK™, Sercos III™, HSR, PRP and others. The second PRU-ICSS subsystem of the AM437x enables EnDat 2.2, Tamagawa, Sigma Delta and another industrial communication protocol in parallel. Additionally, the programmable nature of the PRU-ICSS, along with their access to pins, events, and all SoC resources, provides flexibility in implementing fast real-time responses, specialized data-handling operations, custom peripheral interfaces, and off-loading tasks from the other processor cores of the SoC.

Figure 6 shows an example of the PRU-ICSS block diagram.

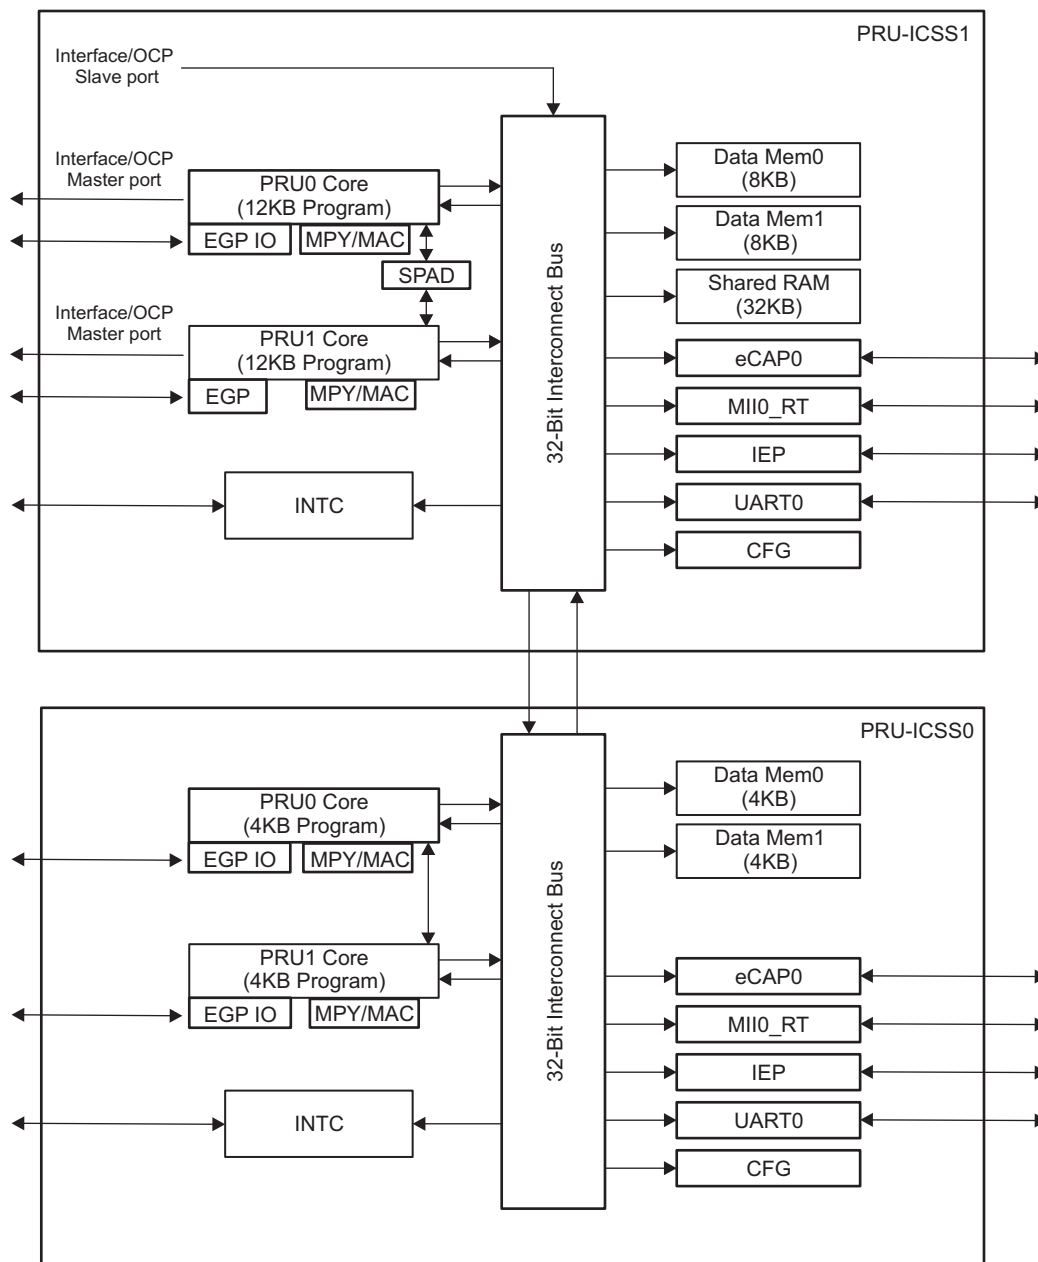


Figure 6. PRU-ICSS Block Diagram

2.3 System Design Theory

2.3.1 Parallel Redundancy Protocol (PRP)

PRP is a redundancy protocol for Ethernet networks (standardized using IEC 62439-3 Clause 4) that is selected for substation automation in the IEC 61850 standard. PRP is application-protocol independent and can be used by most industrial-Ethernet applications that require reliable high-speed communications.

The PRP supports star topology, which ensures a fixed-hop delay with cost of infrastructure.

Figure 7 shows PRP operation over star topology. Once a packet is generated at Node 1, the same packet is sent through LAN A and B, and the destination (Node 3) receives duplicated packets if both deliveries succeed. The redundancy provides zero-recovery time if the packet fails to be delivered in one direction.

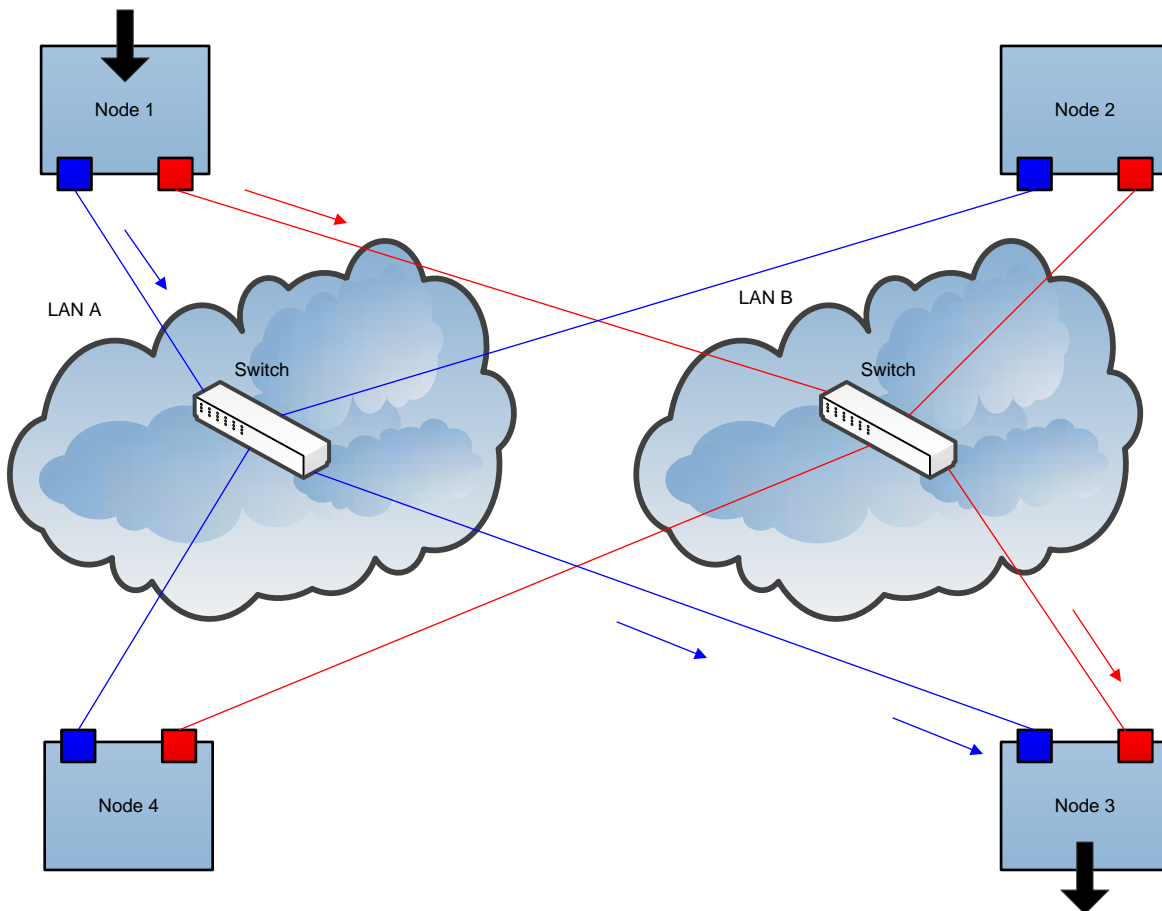


Figure 7. PRP Star Topology

2.3.2 IEEE 1588 (PTP)

The IEEE 1588 is a protocol used to provide high-accuracy time synchronization over networks. Originally defined in the IEEE 1588 standard, this is designed to fill a niche not well served by NTP and GPS. In this design with the HSR protocol, IEEE 1588 v2 peer-to-peer transparent clock profile is supported at layer 2 to synchronize network time by measuring mean path delay using peer delay request and response mechanism. The PTP supports transmissions over IEEE 802.3, and only multicast PTP messages shall be used.

3 Hardware, Software, Testing Requirements, and Test Results

3.1 Required Hardware and Software

3.1.1 Hardware

Figure 8 shows the AM572x industrial development kit (IDK) revision 1.3B. In addition to the AM572x processor, the EVM includes Ethernet PHYs, various storage devices, DDR memory, and power management support. The EVM is designed to support multiple communication standards by providing various interfaces such as Ethernet, CAN, and RS-485.




Figure 8. AM572x IDK Rev 1.3B

3.1.2 Software

The required software for this setup is included in the Processor SDK for real-time Linux. The SDK can be downloaded [here](#). The download page includes useful links to documentation to help users get started quickly.

Figure 9 provides a high-level diagram of the Linux PRP implementation. PRP is added to the Linux Kernel to provide redundant Ethernet interfaces. These interfaces are abstracted to the application layer, which requires no changes to take advantage of PRP. For example, a simple ping command can send data over both interfaces. If a network cable were to be compromised on one of the interfaces, the ping still succeeds on the other interface.

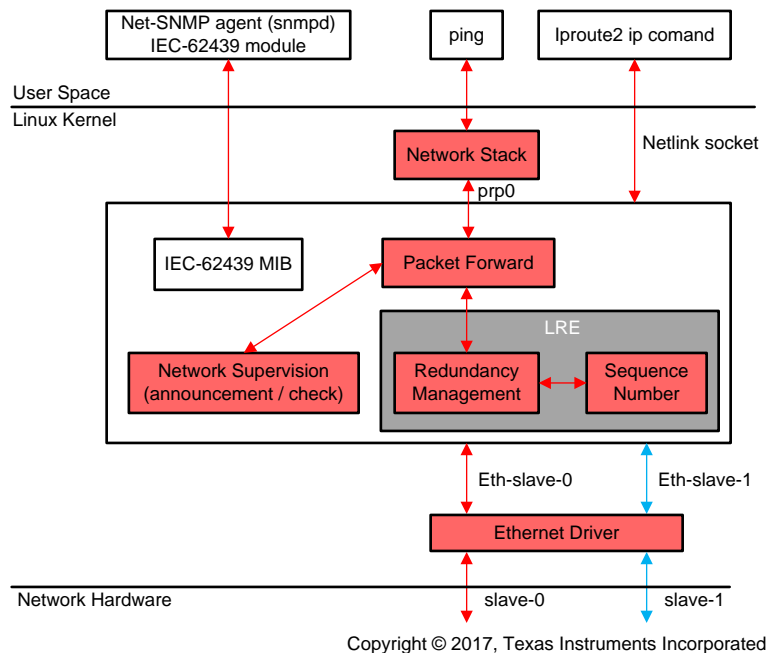


Figure 9. Linux® PRP Architecture

Figure 10 shows a diagram of how the PRU-ICSS on the AM5728 device can be used to aid in packet processing. The PRU-ICSS is a parallel core to the ARM Cortex-A15 that runs Linux. The PRU is designed to minimize latency with deterministic processing.

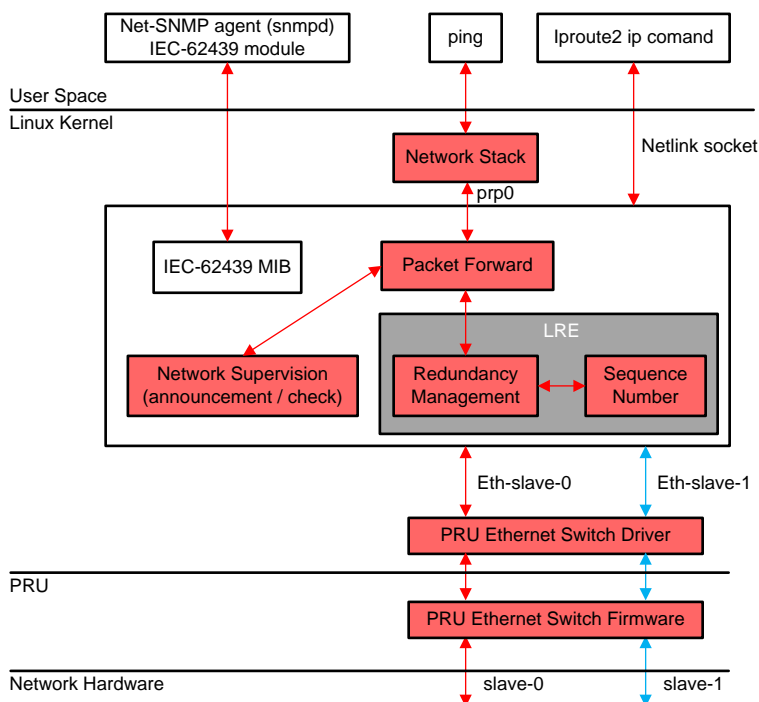


Figure 10. PRU-Enabled PRP

The PRU uses queue priority ranges from 0 to 3 to prioritize packet processing. Lower queue numbers are higher priority, and transmission occurs in the order of priority. The packet length is the MAC PDU frame size. The TX packet is in forms of Ethernet MAC frame that consist of the Ethernet MAC header and payload. The Ethernet MAC header includes a 6-byte destination address, a 6-byte source address, and a 2-byte Ethernet Type. The 4-byte cyclic-redundancy check (CRC) is not required at the application level. The underlying PRP appends the CRC after adding the PRP trailer to the MAC frame.


The Ethernet type can be set to a VLAN tagged frame with an optional 802.1-Q header. Using a VLAN tag allows prioritized packet reception based on the PCP (priority-code point) subfield in the 802.1Q header.

When a packet is received in the PRP firmware, the 3-bit PCP subfield of a VLAN tag is read, and the packet is copied to the appropriate RX queue. Similarly, on exit, four queues are used to store packets. Separate queues on ingress and egress allow for prioritization of packets as required by applications for latency or other quality of service requirements.

3.2 Testing and Results

3.2.1 Test Setup

☒ 11 shows the test setup with two nodes. Each node has two Ethernet connections, and connection from one EVM to the other forms one LAN of the two required for PRP. For these experiments, common Linux applications and commands were used to measure the performance of delivery ratio and latency. For the target TX and RX, a PC is attached to each node to configure test modes and parameters with a serial terminal program. In addition, the underlying firmware generates background traffic necessary for the protocol and statistics.



☒ 11. Test Setup

Looking closely at ☒ 11, notice that the EVM on the right is inverted, so the connections can be shown clearly. Note the same ports from each board are connected to form the two LANs (LAN_A and LAN_B) of the PRP network. This testing could also be done using switches in a star topology as shown in ☒ 7.

Each EVM is running the Processor SDK for real-time Linux from an SD card. Instructions and scripts for creating the SD card are included with the Processor SDK. With Linux running, the following procedure can be used to create the PRP connection between the boards.

1. Configure the appropriate interfaces with the same Ethernet MAC addresses.
2. Bring both interfaces up.
3. Leverage the extended IP command to add a PRP interface using the two ports.
4. Assign the new PRP interface an IP address, and bring the interface up.
5. Use whatever network applications that require the PRP interface. A standard ping can easily test connectivity and latency.

The details of this procedure are fully documented in the [PRP documentation](#) provided with the SDK.

3.2.2 Test Results


The goal of these experiments is to evaluate that TI PRP solution meets the performance requirement for substation automation. 表 1 summarizes the performance requirement for substation automation. The required communication recovery time means the time duration in which a network recovers failure and the application recovery tolerated delay (or grace time) is the time duration in that the substation tolerates an outage of the automation system. The sampled values (SV) are sampled at a nominal value of 4 kHz. Therefore, the target application recovery tolerated delay for SV is 500 μ s (= 2 \times ¼ kHz).

表 1. Recovery Delay Demands as Shown in IEC 61850-5 [6]

COMMUNICATING PARTNERS	SERVICE	APPLICATION RECOVERY TOLERATED DELAY	REQUIRED COMMUNICATION RECOVERY TIME
SCADA to IED, client-server	IEC 61850-8-1	800 ms	400 ms
IED to IED interlocking	IEC 61850-8-1	12 ms (with T _{min} set to 4 ms)	4 ms
IED to IED, reverse blocking	IEC 61850-8-1	12 ms (with T _{min} set to 4 ms)	4 ms
Protection trip excluding busbar protection	IEC 61850-8-1	8 ms	4 ms
Busbar protection	IEC 61850-9-2 on station bus	< 1 ms	Bumpless
Sampled values	IEC 61850-9-2 on process bus	Less than two consecutive samples	Bumpless

3.2.2.1 Latency

The goal of this experiment is to validate if the latency performance meets the requirement for substation automation applications. To measure the latency, one node sends a ping with packet sizes from 64 to 1500 bytes in increments of 200 bytes. The target node replies if it receives the request. The round-trip delay is measured at the originator of the packet by calculating the time gap between TX and RX. Then, the latency is calculated by a half of the round-trip delay. The latency measurement was performed 100 times, and the latency was averaged. The latency is considered a one-way delay based on the definition in the IEC/TR 61850-90-6.



☒ 12. Latency Performance

3.2.2.2 Delivery Ratio

The goal of this experiment is to verify zero network recovery time that is one of requirements for substation automation. For this purpose, the delivery ratio performance was measured while emulating the link failures by disconnecting a link intentionally in the middle of data transmissions.

For this experiment, Node 1 is configured as TX mode with 10,000 packet transmissions, 1528-byte frame size, and 1-ms packet interval. The other nodes are configured as RX mode. During the experiment, Link 1 is disconnected to emulate link failure. To validate the impact of hops and packet types on delivery ratio performance, various experiments were performed with different hops and unicast or broadcast traffic. Each experiment emulated link failure by disconnecting Link 1.

Each experiment captured the number of TX packets at Node 1 and the number of RX packets at the other nodes. The delivery ratio is calculated by the number of TX packets divided by the number of RX packets.

表 2 shows delivery ratio performance over various scenarios. For all scenarios, the result shows a 100% delivery ratio even with link failure, which implies that link failure is recovered immediately. This is expected because redundant communication recovers the link failure immediately.

表 2. Delivery Ratio Performance

TEST SCENARIO	DELIVERY RATIO (%)
Unicast	100
Broadcast	100

4 Design Files

4.1 Schematics

To download the schematics, see the design files at [TIDEP-0103](#).

4.2 Bill of Materials

To download the bill of materials (BOM), see the design files at [TIDEP-0103](#).

4.3 Gerber Files

To download the Gerber files, see the design files at [TIDEP-0103](#).

4.4 Assembly Drawings

To download the assembly drawings, see the design files at [TIDEP-0103](#).

5 Software Files

To download the software files, see the design files at [TIDEP-0103](#).

6 Related Documentation

1. Wikipedia, [Parallel Redundancy Protocol](#)
2. Texas Instruments, [AM572x Sitara™ Processors Silicon Revision 2.0 Data Sheet](#)
3. Texas Instruments, [AM437x Sitara Processors](#)
4. Texas Instruments, [AM335x Sitara Processors](#)
5. Texas Instruments, [Processor SDK for Linux Getting Started Guide](#)
6. Texas Instruments, [Processor SDK Linux HSR, PRP](#)
7. University of Brescia, [IEC61850 One World, One Technology, One Standard](#)

6.1 商標

Sitara is a trademark of Texas Instruments.

NEON is a trademark of Arm Limited.

ARM, Cortex are registered trademarks of Arm Limited.

EtherCAT is a registered trademark of Beckhoff Automation GmbH.

EtherCAT is a registered trademark of Beckhoff Automation GmbH, Germany.

Ethernet POWERLINK is a trademark of Bernecker + Rainer Industrie-ElektronikGes.m.b.H..

EtherNet/IP is a trademark of ODVA, Inc.

PROFINET, PROFIBUS, PROFINET IRT are registered trademarks of PROFIBUS and PROFINET International (PI).

Sercos III is a trademark of Sercos International.

SERCOS is a registered trademark of Sercos International.

Linux is a registered trademark of The Linux Foundation.

すべての商標および登録商標はそれぞれの所有者に帰属します。

7 About the Author

RON BIRKETT is an Applications Engineering Manager at Texas Instruments, where he is responsible for providing technical support and training on Linux running on Embedded Systems. He received his Master of Science in Computer Science from the University of Texas at Dallas, Dallas, TX.

改訂履歴

資料番号末尾の英字は改訂を表しています。その改訂履歴は英語版に準じています。

2017年12月1日発行分から2018年12月31日発行分への変更 (from * Revision (December 2017) to A Revision)	Page
• Added AM437x section	6
• Added AM3359 section	8
• Added PRU-ICSS section	9

重要なお知らせと免責事項

TI は、技術データと信頼性データ(データシートを含みます)、設計リソース(リファレンス・デザインを含みます)、アプリケーションや設計に関する各種アドバイス、Web ツール、安全性情報、その他のリソースを、欠陥が存在する可能性のある「現状のまま」提供しており、商品性および特定目的に対する適合性の黙示保証、第三者の知的財産権の非侵害保証を含むいかなる保証も、明示的または黙示的にかかわらず拒否します。

これらのリソースは、TI 製品を使用する設計の経験を積んだ開発者への提供を意図したものです。(1) お客様のアプリケーションに適した TI 製品の選定、(2) お客様のアプリケーションの設計、検証、試験、(3) お客様のアプリケーションが適用される各種規格や、その他のあらゆる安全性、セキュリティ、またはその他の要件を満たしていることを確実にする責任を、お客様のみが単独で負うものとします。上記の各種リソースは、予告なく変更される可能性があります。これらのリソースは、リソースで説明されている TI 製品を使用するアプリケーションの開発の目的でのみ、TI はその使用をお客様に許諾します。これらのリソースに関して、他の目的で複製することや掲載することは禁止されています。TI や第三者の知的財産権のライセンスが付与されている訳ではありません。お客様は、これらのリソースを自身で使用した結果発生するあらゆる申し立て、損害、費用、損失、責任について、TI およびその代理人を完全に補償するものとし、TI は一切の責任を拒否します。

TI の製品は、TI の販売条件 (www.tij.co.jp/ja-jp/legal/termsofsale.html)、または ti.com やかかる TI 製品の関連資料などのいずれかを通じて提供する適用可能な条項の下で提供されています。TI がこれらのリソースを提供することは、適用される TI の保証または他の保証の放棄の拡大や変更を意味するものではありません。

Copyright © 2019, Texas Instruments Incorporated
日本語版 日本テキサス・インスツルメンツ株式会社

重要なお知らせと免責事項

TI は、技術データと信頼性データ(データシートを含みます)、設計リソース(リファレンス・デザインを含みます)、アプリケーションや設計に関する各種アドバイス、Web ツール、安全性情報、その他のリソースを、欠陥が存在する可能性のある「現状のまま」提供しており、商品性および特定目的に対する適合性の黙示保証、第三者の知的財産権の非侵害保証を含むいかなる保証も、明示的または黙示的にかかわらず拒否します。

これらのリソースは、TI 製品を使用する設計の経験を積んだ開発者への提供を意図したものです。(1) お客様のアプリケーションに適した TI 製品の選定、(2) お客様のアプリケーションの設計、検証、試験、(3) お客様のアプリケーションが適用される各種規格や、その他のあらゆる安全性、セキュリティ、またはその他の要件を満たしていることを確実にする責任を、お客様のみが単独で負うものとします。上記の各種リソースは、予告なく変更される可能性があります。これらのリソースは、リソースで説明されている TI 製品を使用するアプリケーションの開発の目的でのみ、TI はその使用をお客様に許諾します。これらのリソースに関して、他の目的で複製することや掲載することは禁止されています。TI や第三者の知的財産権のライセンスが付与されている訳ではありません。お客様は、これらのリソースを自身で使用した結果発生するあらゆる申し立て、損害、費用、損失、責任について、TI およびその代理人を完全に補償するものとし、TI は一切の責任を拒否します。

TI の製品は、TI の販売条件 (www.tij.co.jp/ja-jp/legal/termssofsale.html)、または ti.com やかかる TI 製品の関連資料などのいずれかを通じて提供する適用可能な条項の下で提供されています。TI がこれらのリソースを提供することは、適用される TI の保証または他の保証の放棄の拡大や変更を意味するものではありません。

Copyright © 2019, Texas Instruments Incorporated
日本語版 日本テキサス・インスツルメンツ株式会社