

HSR/PRP Solutions on Sitara Processors for Grid Substation Communication

Catalog Processors

ABSTRACT

HSR/PRP is Ethernet based communication technology commonly deployed in smart grid substation for low cost, easy to maintain and interoperable common network infrastructure with built-in redundancy. This application report presents the overview of HSR/PRP technology followed by HSR/PRP solutions on Sitara devices. TI's PRU-ICSS technology with features like real-time and determinism allows Sitara devices to offer HSR/PRP solution. HSR/PRP solution is available as part of Processor SDK for Linux and RTOS.

This document points to various collateral and resources available for customers to quick ramp on HSR/PRP solution on Sitara devices including PRU-ICSS technology.

Contents

1	HSR/PRP for Smart Grid Substation Communication	2
2	HSR/PRP Solution on Sitara Processors	5
3	References	11

List of Figures

1	Typical Electric Smart Grid	2
2	Process Bus vs Station Bus	3
3	Typical HSR Network Topology.....	4
4	Typical PRP Network Topology	5
5	PRU-ICSS Block Diagram	6
6	Sitara Devices HSR/PRP Software Block Diagram	7
7	PRU-ICSS HSR/PRP Firmware Block Diagram	8
8	Sitara Devices Linux HSR/PRP Software Block Diagram	9
9	Sitara Devices RTOS HSR/PRP Software Block Diagram	10

List of Tables

1	Linux HSR/PRP	11
2	TI-RTOS HSR/PRP	11

Trademarks

All trademarks are the property of their respective owners.

1 HSR/PRP for Smart Grid Substation Communication

A smart grid substation requires fast, fault tolerant communication that helps operators achieve cost efficient and reliable operations. A substation is a key component of the electricity-grid infrastructure, located everywhere from power generation facilities throughout the distribution network to the low-voltage feeders serving residences and businesses. Substations are a primary factor in transforming voltage levels for transmission and performing important functions such as switching, monitoring, and protecting sub-systems in order to maintain grid efficiency and reliability. Substation networks require communications that have built-in redundancy and time synchronization with no down time due to a single fault in the network. If a failure does occur, the network must recover within a given short time. Redundancies create more than one path between the source and destination to reroute traffic at the time of failure.



Figure 1. Typical Electric Smart Grid

Operators need to continually monitor the health of networks and take action to maintain the operation with efficiency. This need leads to the requirement for reliable and low-latency communications between the control center of the operator and high-value nodes such as substations. Due to the different requirements of station and process bus networks, different protocols need to be implemented according to their specific performance characteristics.

A ‘station bus’ and ‘process bus’ are Ethernet networks that connect various intelligent devices and smart control/SCADA systems. The ‘station bus’ network is mainly used to carry event-driven Ethernet messages for supervising the system, it interconnects the whole substation and provides connectivity between central management and the individual bays. The ‘station bus’ typically carries Generic Object Oriented Substation Event (GOOSE) traffics and TCP/UDP traffics. Traffics in the station bus can tolerate frame losses.

The ‘process bus’ is mainly used to carry measurement traffic, in the form of sampled values (SV) traffic. These are small Ethernet frames carrying measurement values, sent by merging units for example. SV traffic does not tolerate frame loss, therefore, media redundancy must work virtually without interruption.

The process bus is redundantly connected to the station bus through HSR or PRP. The process bus traffic is sent to the station bus for analysis and for additional monitoring,

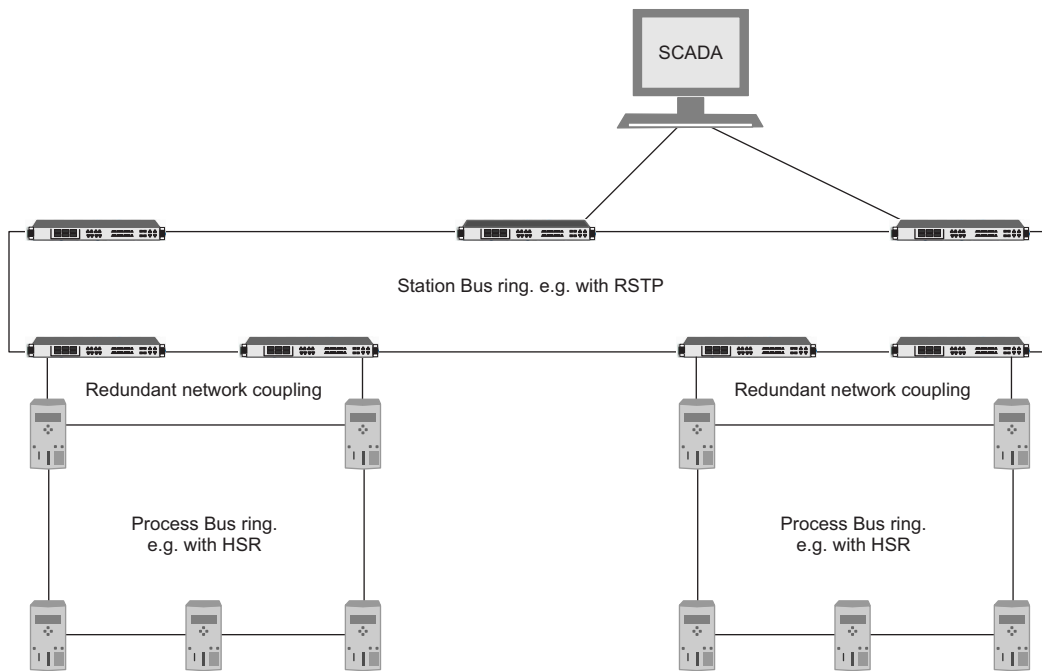


Figure 2. Process Bus vs Station Bus

The need for high performance, high reliability, and predictable Ethernet networks drove the motivation for the IEC 62439 standard with the goal to create low cost, easy to maintain and interoperable common network infrastructure with built-in redundancy.

1.1 High Availability Seamless Redundancy (HSR)

IEC 62439-3 clause 5 defines the High Availability Seamless Redundancy (HSR) protocol as a redundancy protocol for substation automation. The HSR is application-protocol independent redundancy protocol based on double transmission of message frames over ring-topology networks in both directions. In the case of an error, the message frame will be transmitted without any delay. No reconfiguration time or relearning of the communication paths is necessary for the network.

Since communication over an HSR network still continues without interruption when there is a network failure, HSR can be used in process bus applications.

HSR frames are identified uniquely by the HSR tag. Field devices are connected to a Double Attached Node HSR (DANH in a Ring topology and only DANH compliant nodes can be connected to the HSR network). Other standard Ethernet devices (Singly Attached Nodes, SANs) need to be connected through a Redundancy Box (RedBox) to work with HSR networks. The HSR tag is used to manage redundancy in the HSR network.

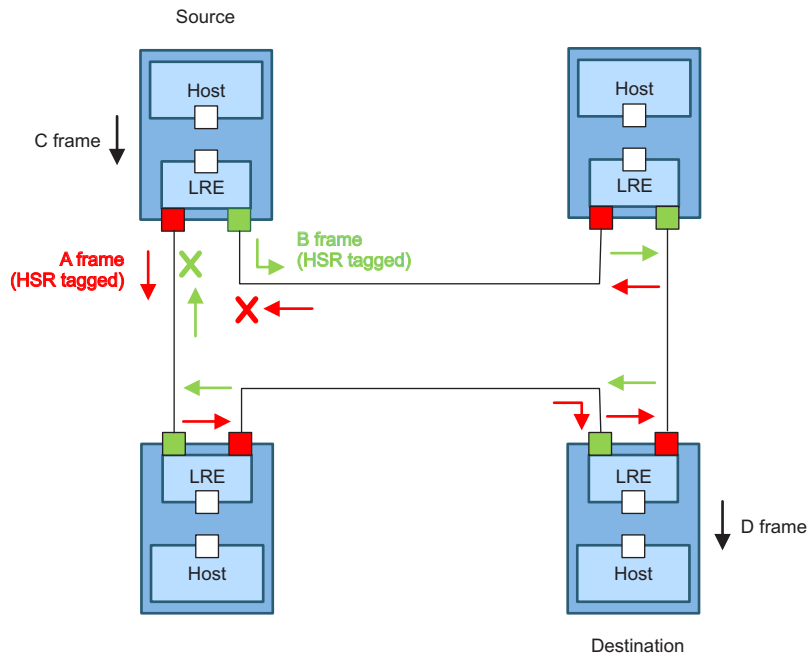


Figure 3. Typical HSR Network Topology

A DANH node has two ports operated in parallel. A source DANH prefixes a frame passed from its upper layers with an HSR tag to identify frame duplicates and sends the frame over each port.

A destination DANH receives, in the fault-free state, two identical frames (one from each port) within a certain interval. It removes the HSR tag of the first frame before passing it to its upper layers and discards any duplicate frames. The nodes support the IEEE 802.1D bridge functionality and forward frames from one port to the other, except if they have already sent the same frame in that same direction. In particular, the node will not forward a frame that it injected into the ring. A destination node of a unicast frame does not forward a frame for which it is the only destination, except for testing. Frames circulating in the ring carry the HSR tag inserted by the source, which contains a sequence number. The doublet (source MAC address, sequence number) uniquely identifies copies of the same frame. The frame discarding prevents the network from being flooded with unnecessary traffic.

1.2 Parallel Redundancy Protocol (PRP)

PRP simply doubles the necessary network infrastructure and is defined in IEC 62439-3 Clause 4. All devices that need to take advantage of this redundant network infrastructure must be connected with a dual attachment network interface that connects both interfaces, each to one of the two redundant LANs A and B. Both networks are used simultaneously and both networks carry the same data that is sent redundantly by the Dual Attached Nodes PRP; the DANP. Each DANP duplicates its entire network traffic that travels over both LAN A and LAN B.

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

When one of the two LANs experiences a fault, the network traffic still runs over the other LAN without interruption. The PRP can, therefore, be used with the process bus, because in case of a failure in one network, the SV traffic can still safely travel through the other network to its destination.

A DANP node has two ports that operate in parallel and that are attached to the same upper layers of the communication stack through the Link Redundancy Entity (LRE). For the basic communication, the LRE presents toward its upper layers the same interface as a non-redundant network adapter, so the upper layers are unaware of redundancy. The LRE has two tasks: handling of duplicates and management of redundancy. When receiving a frame from the node's upper layers, the LRE appends a Redundancy Check Trailer (RCT) containing a sequence number to the frame and sends the frame through both ports at nearly the same time. The two frames are nearly identical except for the LAN identifier (and the checksum).




Figure 4. Typical PRP Network Topology

The two frames transit through the two LANs with different delays, ideally they arrive at nearly the same time at the destination node. When receiving frames from the network, the LRE forwards the first received frame of a pair to its node's upper layers and discards the duplicate frame (if it arrives). It removes the RCT, if required.

2 HSR/PRP Solution on Sitara Processors

2.1 PRU-ICSS as Key Technology to Enable HSR/PRP on Sitara Device

Texas Instruments Sitara processors embed PRU-ICSS technology that enables customers to add HSR/PRP Dual Attached Node support to the system.

The Programmable Real-Time Unit Subsystem and Industrial Communication Subsystem (PRU-ICSS) consists of dual 32-bit RISC cores (Programmable Real-Time Units, or PRUs), shared, data, and instruction memories, internal peripheral modules, and an interrupt controller (INTC). 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 in offloading tasks from the other processor cores of the system-on-chip (SoC).

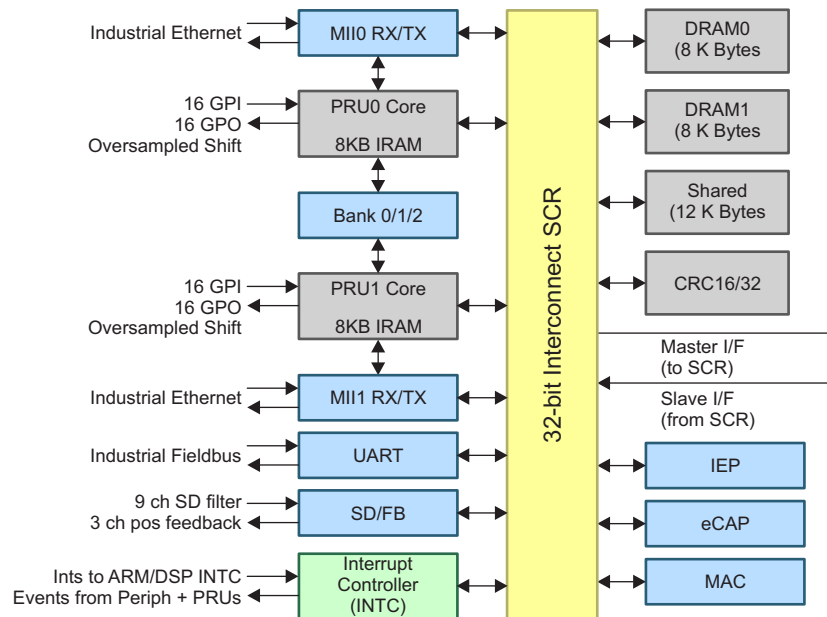


Figure 5. PRU-ICSS Block Diagram

The PRU cores within the subsystems have access to all resources on the SoC through the Interface/OCP Master port, and the external host processors can access the PRU-ICSS resources through the Interface/OCP Slave port. The 32-bit interconnect bus connects the various internal and external masters to the resources inside the PRU-ICSS. The INTC handles system input events and posts events back to the device-level host CPU. The PRU cores are programmed with a small, deterministic instruction set. Each PRU can operate independently or in coordination with each other and can also work in coordination with the device-level host CPU. This interaction between processors is determined by the nature of the firmware loaded into the PRU's instruction memories.

PRU-ICSS training series (video) can be found at: [Introduction to the Programmable Real-Time Unit \(PRU\) Training Series](#).

For more information about PRU-ICSS, see the [PRU-ICSS](#) wiki.

2.2 Sitara HSR/PRP Solution Features and Architecture

The implementation consists of ICSS-PRU firmware, Drivers (Linux/RTOS), application and SNMP support. LRE is mostly handled in PRU-ICSS firmware and Driver, whereas, Host (Cortex Ax) runs drivers and application code. TI provides PRU-ICSS HSR firmware and drivers (Linux and TI-RTOS) are free of charge, customers develop protocol and applications.

Figure 6 shows HSR/PRP switch architecture on Sitara devices that are made of firmware, drivers, protocols and applications.

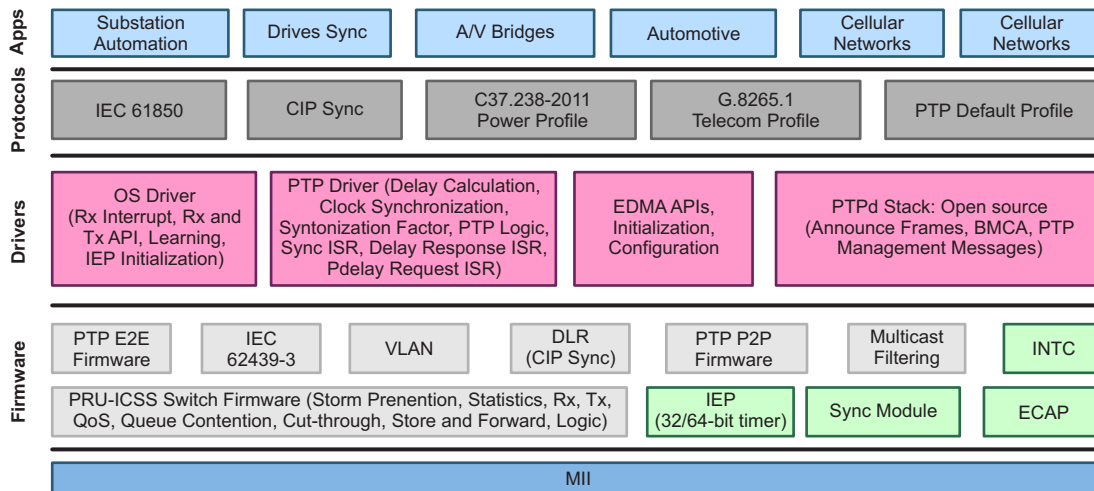


Figure 6. Sitara Devices HSR/PRP Software Block Diagram

The firmware layer is run on ICSS-PRU that implements features like cut-through, maintenance of node table, duplicate discards, and so forth. Driver layer connects firmware to protocol layer and implement features like frame duplication in transmit direction. Protocol layer implements features like IEC 61850, CIP sync; and interface with grid automation application.

2.3 HSR/PRP Features and Capability

- Enhance ICSS cut-through switch with HSR capability
- HSR End Node (DANH) complying to IEC62439-3 Edition 2 Clause 5
- PRP End Node (DANP) complying to IEC6239-3 Edition 2 Clause 4
- 10/100 Mbits/s Full Duplex Ethernet Interface
- Cut-through with switching time < 3 μs
- Store and forward with data integrity check
- Support for HSR Modes H;N,T,U,M
- Continuous update of node table
- Configurable node table size
- Up to 256 entries
- Firmware can be extended to support 512 entries on AM57x
- QoS based on 8-bit VLAN PCP field. The eight levels are mapped to four different queues
- Support all statistics MIB variables defined in the Standard
- Node table enhanced with traffic capture (statistics) for all frames received by the host.
- Storm prevention (AM57x only)

2.4 PRU-ICSS Firmware

A common firmware is used across TI RTOS and Linux implementations of HSR/PRP. PRU-ICSS firmware handles features like cut-through, node table management, Port to host/port duplicate detection and elimination.

The firmware runs the Link Redundancy Entity (LRE) on the two PRU cores letting the host manage the high level stack and user application as well as the necessary platform initialization

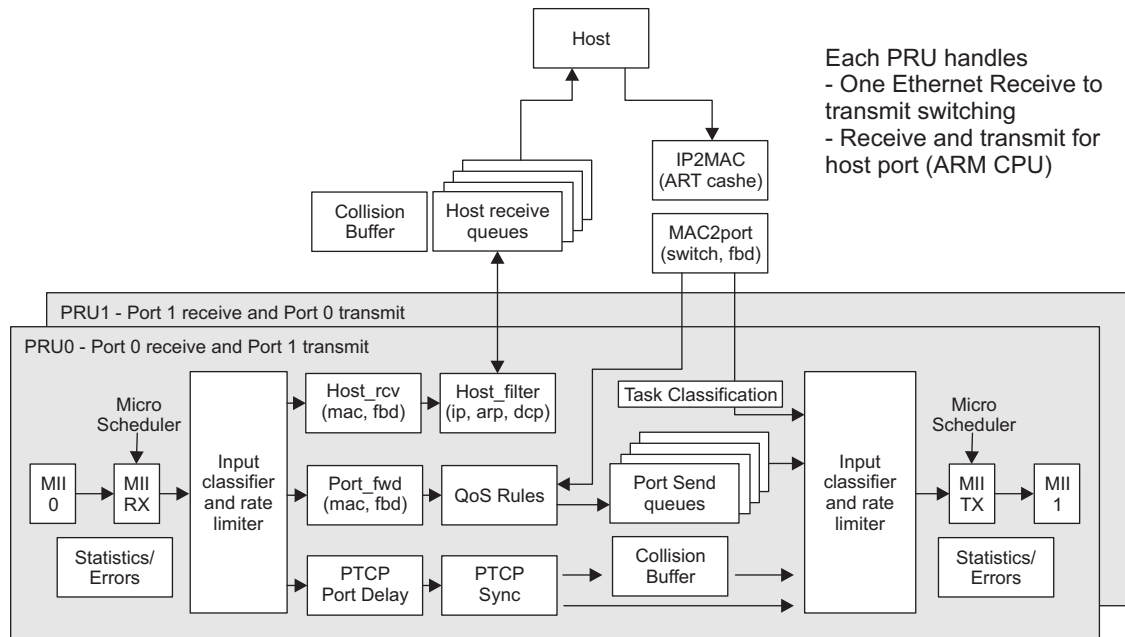


Figure 7. PRU-ICSS HSR/PRP Firmware Block Diagram

The following functions are executed on the PRU0 Core:

- Initialization done by the MS
- Receive frame on port A (RX task)
- Store and forward frame to port B and/or host (HSR ONLY)
- Cut-through frame received on Port A to Port B (HSR ONLY)
- Transmit frame from host on port B (duplication managed at the host side)
- Discard duplicate for frames received on port A/B and transmitted to the host
- Discard duplicate for frames received on port A/B and forwarded to the port peer (HSR ONLY)
- HSR tag detection and removal for ingress traffic (HSR ONLY)
- Redundancy control trailer (RCT) removal (PRP ONLY)
- Statistics for received frames on port A and transmit frames on port B
- MIB counters update
- Queue contention handling for host and port B
- Receive frame in shadow buffer in case of queue contention on host and Port B
- Update node table
- Update/check duplicate table
- Update/check node table

The functions are exactly symmetrical to PRU1 with the following exceptions:

- Initialization, host duplicate check task and port duplicate check task are only performed by PRU0
- The Node table check task is only performed by PRU

For more details on the ICSS-PRU HSR/PRP firmware, see the [ICSS Firmware](#).

2.5 Linux HSR/PRP Solution on Sitara Devices

Figure 8 shows Linux components that has been modified or developed for HSR/PRP. Linux HSR/PRP driver provides a standard Ethernet interface to application layer and hide the details of the redundancy protocol under the driver layer. Linux HSR/PRP driver implements the Link Redundancy Entity (LRE) that is central to implementing redundancy using respective protocol.

The Driver performs the task of duplicating a packet and sending it out on both ports. It also performs the task of node table, port-port and port-host duplicate table management.

The LRE replaces EMAC LLD Rx and Tx API's with identical API's providing the redundancy.

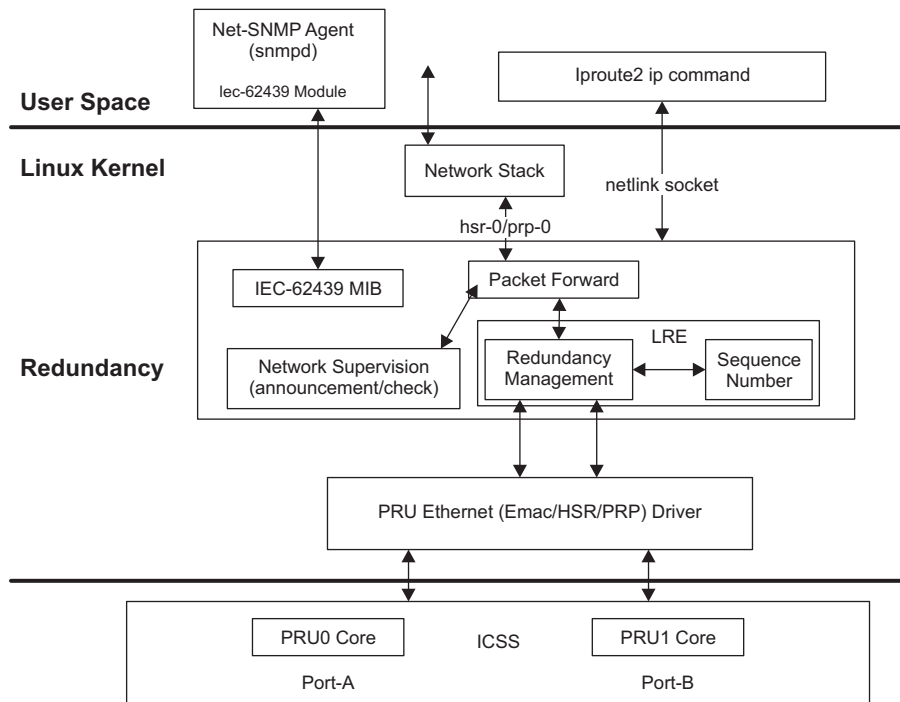


Figure 8. Sitara Devices Linux HSR/PRP Software Block Diagram

Starting from the Processor SDK 04.00.00 release, the PRU Ethernet driver supports HSR/PRP firmware offload. When a HSR/PRP firmware is loaded, the PRU Ethernet driver will do proper firmware configurations and indicate offload capabilities in the netdev feature flag. The Linux HSR/PRP driver stack use these feature flag to determine if offload is supported or not at the lower level Ethernet driver and disable certain processing at the its layer.

For more information on the Linux-based HSR/PRP solution including features, setup, test procedures and performance, see [Linux Software](#).

2.5.1 TI-RTOS HSR/PRP Solution on Sitara Devices

The Processor-SDK for RTOS provides fundamental platform software and tools for development, deployment and execution of RTOS-based applications

Figure 9 shows the architecture of HSR/PRP solution on Sitara AM335x device.

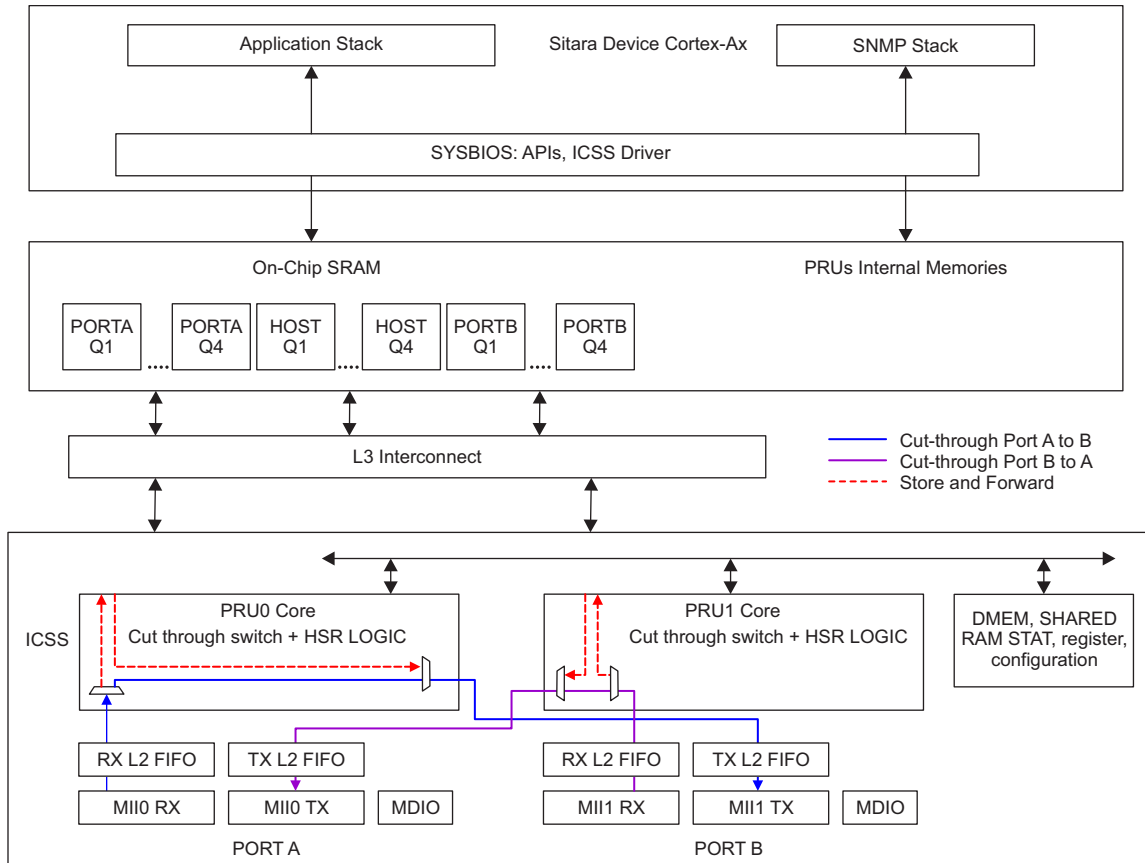


Figure 9. Sitara Devices RTOS HSR/PRP Software Block Diagram

For more Information regarding TI-RTOS, see the [Processor SDK RTOS](#).

For more information on TI-RTOS HSR/PRP, including features and setup, see the [PRU ICSS HSR PRP](#) wiki.

2.5.2 HSR/PRP Software Availability From TI

The Texas Instruments PRU ICSS Industrial Packages are a single scalable software platform for Industrial Protocols/Drives solution including HSR/PRP that offers streamlined development across different TI Sitara Processors. The Industrial Protocol packages provide fundamental software for development of various HSR/PRP solutions. The packages also include sample demo applications that showcase the ability of the software/hardware to implement those.

This package runs on top of the Processor SDK and consists of (depending on the package):

- PRU-ICSS firmware to implement HSR/PRP LRE
- Associated drivers
- Stack adaptation layer
- Evaluation versions of protocol stack library for selected protocols
- Sample applications

The PRU-ICSS Industrial package for HSR/PRP is available for download at: [PRU-ICSS-HSR-PRP-DAN 01_00_03_02](#).

For more information about the package, see the [Industrial Protocol Package Getting Started Guide](#) wiki.

Table 1. Linux HSR/PRP

	AM335x	AM437x	AM57x	K2Gx
HSR	√	√	√	√
PRP	√	√	√	√

√ = Released or In-development.

Table 2. TI-RTOS HSR/PRP

	AM335x	AM437x	AM57x	K2Gx
HSR	√	√	√	√
PRP	√	√	√	√

√ = Released or In-development.

3 References

- [Introduction to the Programmable Real-Time Unit \(PRU\) Training Series](#)
- [PRU-ICSS](#) wiki
- [ICSS Firmware](#)
- [Linux Software](#)
- [Processor SDK RTOS](#)
- [PRU ICSS HSR PRP](#) wiki
- [PRU-ICSS-HSR-PRP-DAN 01_00_03_02](#)
- [Industrial Protocol Package Getting Started Guide](#) wiki

IMPORTANT NOTICE FOR TI DESIGN INFORMATION AND RESOURCES

Texas Instruments Incorporated ("TI") technical, application or other design advice, services or information, including, but not limited to, reference designs and materials relating to evaluation modules, (collectively, "TI Resources") are intended to assist designers who are developing applications that incorporate TI products; by downloading, accessing or using any particular TI Resource in any way, you (individually or, if you are acting on behalf of a company, your company) agree to use it solely for this purpose and subject to the terms of this Notice.

TI's provision of TI Resources does not expand or otherwise alter TI's applicable published warranties or warranty disclaimers for TI products, and no additional obligations or liabilities arise from TI providing such TI Resources. TI reserves the right to make corrections, enhancements, improvements and other changes to its TI Resources.

You understand and agree that you remain responsible for using your independent analysis, evaluation and judgment in designing your applications and that you have full and exclusive responsibility to assure the safety of your applications and compliance of your applications (and of all TI products used in or for your applications) with all applicable regulations, laws and other applicable requirements. You represent that, with respect to your applications, you have all the necessary expertise to create and implement safeguards that (1) anticipate dangerous consequences of failures, (2) monitor failures and their consequences, and (3) lessen the likelihood of failures that might cause harm and take appropriate actions. You agree that prior to using or distributing any applications that include TI products, you will thoroughly test such applications and the functionality of such TI products as used in such applications. TI has not conducted any testing other than that specifically described in the published documentation for a particular TI Resource.

You are authorized to use, copy and modify any individual TI Resource only in connection with the development of applications that include the TI product(s) identified in such TI Resource. NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE TO ANY OTHER TI INTELLECTUAL PROPERTY RIGHT, AND NO LICENSE TO ANY TECHNOLOGY OR INTELLECTUAL PROPERTY RIGHT OF TI OR ANY THIRD PARTY IS GRANTED HEREIN, including but not limited to any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information regarding or referencing third-party products or services does not constitute a license to use such products or services, or a warranty or endorsement thereof. Use of TI Resources may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

TI RESOURCES ARE PROVIDED "AS IS" AND WITH ALL FAULTS. TI DISCLAIMS ALL OTHER WARRANTIES OR REPRESENTATIONS, EXPRESS OR IMPLIED, REGARDING TI RESOURCES OR USE THEREOF, INCLUDING BUT NOT LIMITED TO ACCURACY OR COMPLETENESS, TITLE, ANY EPIDEMIC FAILURE WARRANTY AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF ANY THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

TI SHALL NOT BE LIABLE FOR AND SHALL NOT DEFEND OR INDEMNIFY YOU AGAINST ANY CLAIM, INCLUDING BUT NOT LIMITED TO ANY INFRINGEMENT CLAIM THAT RELATES TO OR IS BASED ON ANY COMBINATION OF PRODUCTS EVEN IF DESCRIBED IN TI RESOURCES OR OTHERWISE. IN NO EVENT SHALL TI BE LIABLE FOR ANY ACTUAL, DIRECT, SPECIAL, COLLATERAL, INDIRECT, PUNITIVE, INCIDENTAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES IN CONNECTION WITH OR ARISING OUT OF TI RESOURCES OR USE THEREOF, AND REGARDLESS OF WHETHER TI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You agree to fully indemnify TI and its representatives against any damages, costs, losses, and/or liabilities arising out of your non-compliance with the terms and provisions of this Notice.

This Notice applies to TI Resources. Additional terms apply to the use and purchase of certain types of materials, TI products and services. These include; without limitation, TI's standard terms for semiconductor products (<http://www.ti.com/sc/docs/stdterms.htm>), [evaluation modules](#), and [samples](http://www.ti.com/sc/docs/sampterm.htm) (<http://www.ti.com/sc/docs/sampterm.htm>).

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2018, Texas Instruments Incorporated