

Melissa Watkins
application engineer

Carlos Betancourt
product marketing manager
Texas Instruments

Ensuring real-time predictability

Leveraging TI's Sitara™ Processors Programmable Real-Time Unit

Executive summary

System developers have discovered that many, if not most, contemporary applications require a finely tuned combination of high-speed software processing and real-time hardware performance. In other words, high-performance is not the equivalent of real-time performance. In fact, they both have their own unique and often mutually exclusive set of requirements which prevents either one from performing the role of the other very well.

For example, high-performance processors like ARM® Cortex®-A cores have an entirely different set of resources and processing capabilities than those of real-time processing cores, like the Programmable Real-Time Unit (PRU) coprocessor in TI's Sitara processors. The capabilities that make an ARM core so powerful at processing software could also impede its real-time determinism and predictability. And, in many of today's most sophisticated applications, real-time capabilities are just as critical as high-performance, if not more so.

Real-time requirements

Outside of the data center, many systems need the low-latency predictability of real-time processing in one way or another. In fact, even many general-purpose systems that require a high-level operating system (HLOS) often have a real-time component or subsystem, such as communication protocol processing, audio processing, lighting control, sensor monitoring, factory or home automation, motor control and others.

In these types of systems, a general-purpose processor (GPP), no matter how high its performance, cannot deliver the guaranteed response time within strict time constraints that typify real-time applications and subsystems. Many of the features of a GPP like instruction pipelines and memory and interconnect architectures which make it so effective running a HLOS (Figure 1) will often become counter-productive in real-time applications. The typical architecture surrounding a GPP core on a system-on-a-chip (SoC) will include several layers of memory, which may be internal to the processing core or external, as well as shared memory or dedicated to one core. Moreover, the usual SoC architecture will also include several layers

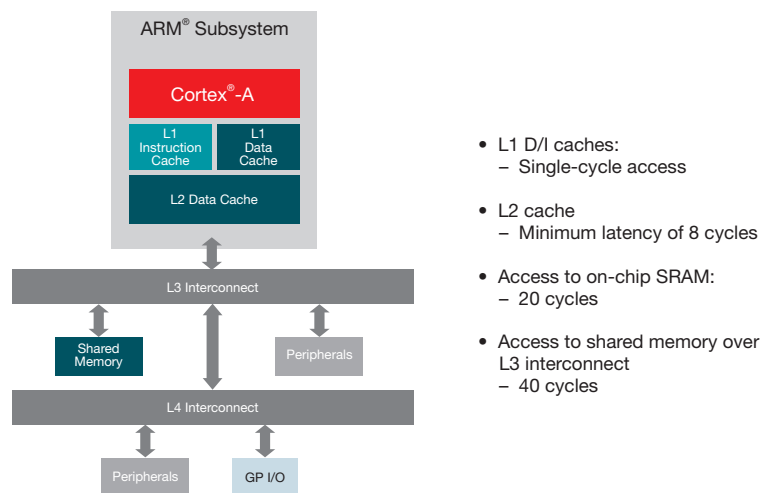


Figure 1. A typical Cortex-A-based SoC general-purpose processing architecture

of interconnects which link the various on-chip modules, peripherals and eventually lead off the chip via specialized or general-purpose input/output (GPIO) pins. All of these facilities and structure can get in the way of real-time processing.

The architecture surrounding a GPP cannot provide the predictable low-latency response times that must be guaranteed in a real-time application. Accessing data on any of the various levels of memory and communicating over any or all of the several layers of interconnects will add to the core's response time and ensure the unpredictability of the response. The processor's response time to an incoming time-sensitive interrupt from a sensor, for example, will vary according to a number of factors, such as where the needed data is stored, how many layers of interconnects must be traversed to access or store data, the processing load currently executing on the GPP core and others. One experiment on a certain GPP architecture showed that a simple toggle on a GPIO pin could take as much as 200 nanoseconds (ns) whereas the equivalent toggle on a real-time processor with direct GPIO pin access would have a response time of five ns, or forty times faster. In addition to the often slower response typical of GPP processors, the actual response time is usually unpredictable. In other words, the response time for a GPP to handle a real-time event will likely vary each time the event occurs.

To overcome the real-time limitations of GPP cores, certain capabilities such as real-time coprocessors are often integrated into an SoC architecture. In the example below (Figure 2), the Texas Instruments (TI) Programmable Real-Time Units (PRU) form the basis for a real-time subsystem on processors. Such an architecture gives the SoC direct and fast access to the outside world since each PRU has its own single-cycle I/O. Additionally, local memory and peripherals dedicated to each real-time engine means that each unit is able to guarantee low-latency responsiveness. Plus an incoming interrupt has direct access to a real-time processing engine without encountering the delays caused by crossing several layers of interconnects and memory.

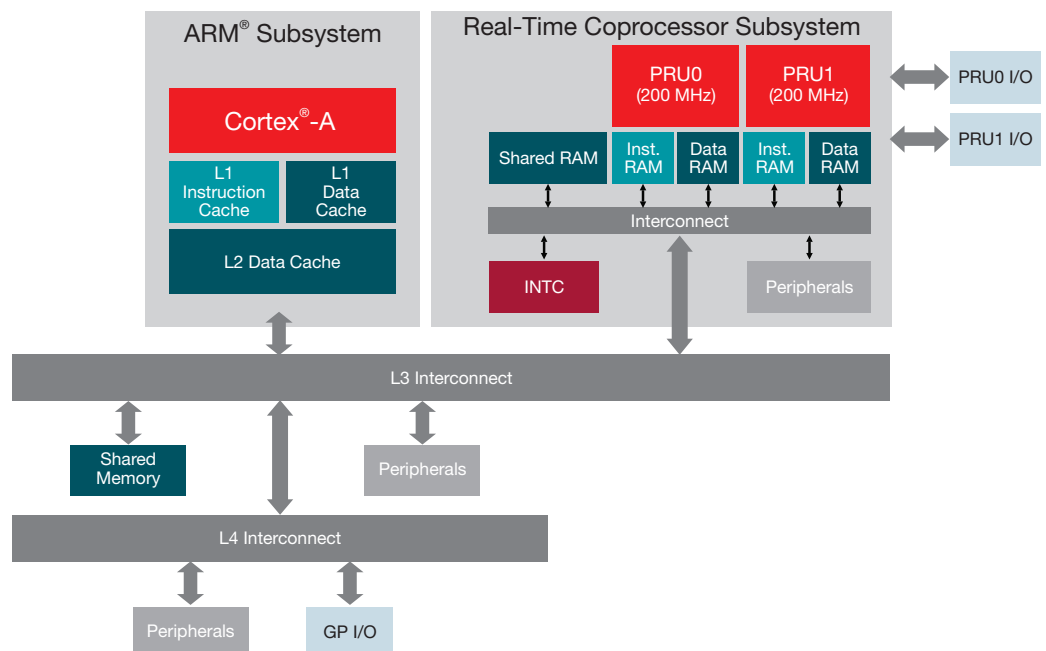


Figure 2. A general-purpose core supplemented with real-time coprocessors

Programmable Real-Time Unit (PRU)

The PRU (Figure 3) which is deployed along with ARM cores in the Sitara AM335x, AM437x and AM5x processors fulfills the role of a low-latency, deterministic real-time subsystem. Each PRU subsystem is made up of two 200-MHz real-time cores (or PRUs), each with a five nanosecond (ns) cycle time per instruction. Since the real-time cores are not equipped with an instruction pipeline, single-cycle instruction execution is ensured. The PRU's small, deterministic instruction set with multiple bit-manipulation instructions is easy to learn and use.

Shared memory as well as instruction and data memory dedicated to each real-time core allows for flexible program execution among all of the real-time and GPP ARM cores that might make up the SoC. One program or task might be better performed on one core while another could be executed faster when the processing load is shared among several PRUs and ARM cores. Direct access from the PRU to the ARM cores' layers of interconnects enables either tightly coupled execution or independent core operations.

Access to all of the system's interconnects allows a PRU to call on any resource in the system when needed for a particular program implementation. In addition, each PRU subsystem comes with its own set of dedicated peripherals to ensure the unit's responsiveness. These peripherals avoid the data traffic on secondary and tertiary interconnects in the system by accessing directly the PRU's real-time cores. Several peripherals, such as Management Data Input/Output (MDIO) and Media Independent Interface (MII), enable real-time Ethernet capabilities. Other PRU subsystem peripherals include enhanced-Capture (eCAP) and a UART interface.

Because of the interrupt controller and fast I/O pins dedicated to each dual-core PRU, the unit is able to closely monitor external events and respond in a predictable period of time. Each PRU has its own set of up to 30 inputs and 32 outputs that directly access external pins on the device package.

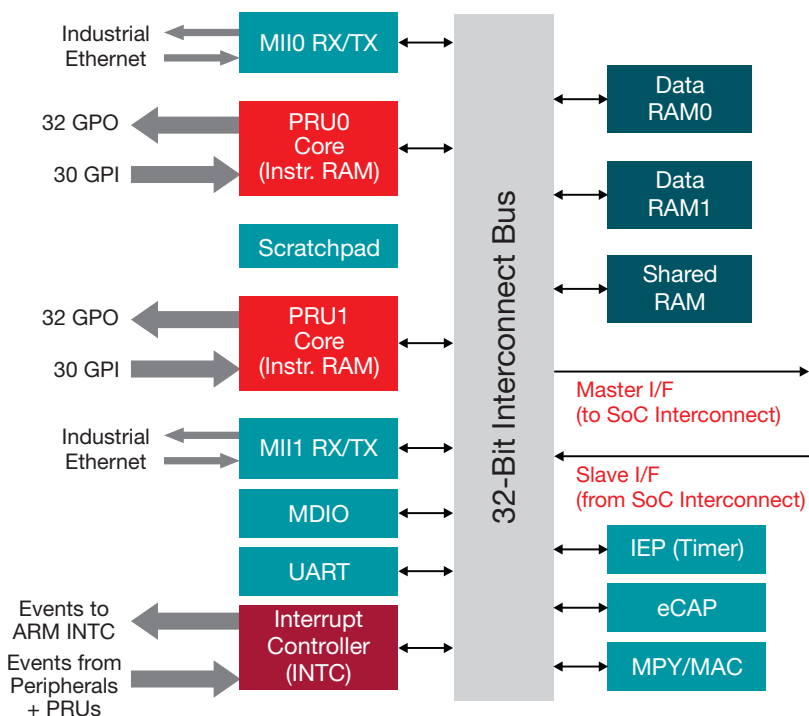


Figure 3. The TI Programmable Real-Time Unit (PRU) subsystem

***The best of both worlds:
Combining ARM and
PRU cores***

The Sitara line of multicore processors features escalating combinations of ARM Cortex-A cores and PRUs. With this architecture (Figure 4), the system designer can select the device with the right combination of high, general-purpose processing and real-time performance to meet the specific needs of the application. ARM cores have all of the resources and instruction support expected for high-performance operating system execution, either for a HLOS or a real-time operating system (RTOS), or both at once.

At the same time, the five nanosecond cycle time of PRU cores as well as their low-latency data transfers and high-speed I/O accesses assure designers that transfers and data modifications will be performed in a predictable period of time. The responsive simplicity of a PRU core make it a natural fit for bit-banged communication interfaces like the Serial Peripheral Interface (SPI), the Inter-IC (I²C) bus interface and others, including many industrial automation protocols.




Figure 4. Base Sitara architecture with ARM cores and PRUs

Some end equipment architectures incorporate a field programmable gate array (FPGA) for real-time processing. Unlike the integrated PRU coprocessor (Figure 4), an FPGA that is external to an ARM-based system processor would necessarily increase the bill of materials (BOM) costs, require additional board space, add complexity to the design and increase the power consumption of the system. With a PRU solution, system developers also benefit from a common software code base which simplifies feature upgrades, and multi-protocol processing across systems or on the same system. Development of successive generations of systems is accelerated by simply migrating the code base to the next model in the product line.

The Sitara hardware structure in Figure 4 supports a software architecture (Figure 5 on the following page) that tightly couples ARM cores and PRUs to ensure seamless and high-speed application processing. Typically, a Linux™ kernel running on an ARM core would include RemoteProc and rpmsg kernel drivers for the PRU subsystem. The RemoteProc is the basic control mechanism, allowing the ARM core to load PRU firmware, enabling PRU processing and other functions. Through the rpmsg driver user space applications and the PRU can pass messages (buffers) back and forth.

Applications

The simple implementation of a PRU combined with its full complement of resources make it a versatile processing engine for a wide range of real-time tasks, subsystems and application modules. Designers have taken advantage of the PRU to deploy straight forward subsystems like stepper motor control units,

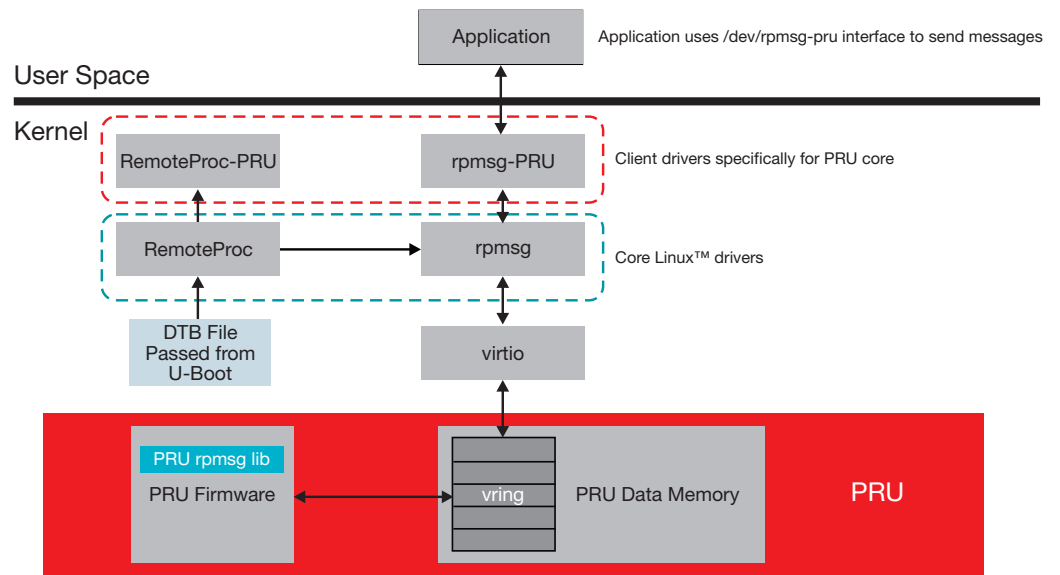


Figure 5. Sitara ARM/PRU software architecture

bit-banging communications processors and sensor interfaces to the more complex tasks such as camera and LCD display interfaces, smart card processing and even very complex applications like Ethernet industrial automation protocol processing.

One of the more interesting applications of the PRU involves a 3D printer (Figure 6). In this case, designers took advantage of the BeagleBone development board with a Sitara AM335x processor featuring an ARM Cortex-A8 core running Linux, the user interface and model processing while the PRU performed the real-time control of five stepper motors. A shared region of memory was reserved for communications between the ARM core and the PRU.



Figure 6. Sitara AM335x with ARM Cortex-A8 core and PRU running a 3D printer

Conclusions

In today's complex application world even the most straight forward general-purpose systems will frequently need the predictability, determinism and low-latency responsiveness that only real-time processing can deliver. Teaming the PRU with high-performance ARM Cortex-A cores in the Sitara line of processors provides the best of both worlds. By offering the best of general-purpose and real-time processor elements in a single package, end equipment is optimized for cost and power consumption, yet flexible for feature upgrades through efficient reprogramming.

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

Sitara is a trademark of Texas Instruments. All other trademarks are the property of their respective owners.

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46, latest issue, and to discontinue any product or service per JESD48, latest issue. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products (also referred to herein as "components") are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its components to the specifications applicable at the time of sale, in accordance with the warranty in TI's terms and conditions of sale of semiconductor products. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by applicable law, testing of all parameters of each component is not necessarily performed.

TI assumes no liability for applications assistance or the design of Buyers' products. Buyers are responsible for their products and applications using TI components. To minimize the risks associated with Buyers' products and applications, Buyers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI components or services are used. Information published by TI regarding third-party products or services does not constitute a license to use such products or services or a warranty or endorsement thereof. Use of such information 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.

Reproduction of significant portions of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI components or services with statements different from or beyond the parameters stated by TI for that component or service voids all express and any implied warranties for the associated TI component or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards which anticipate dangerous consequences of failures, monitor failures and their consequences, lessen the likelihood of failures that might cause harm and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI's goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed a special agreement specifically governing such use.

Only those TI components which TI has specifically designated as military grade or "enhanced plastic" are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components which have **not** been so designated is solely at the Buyer's risk, and that Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components as meeting ISO/TS16949 requirements, mainly for automotive use. In any case of use of non-designated products, TI will not be responsible for any failure to meet ISO/TS16949.

Products

Audio	www.ti.com/audio
Amplifiers	amplifier.ti.com
Data Converters	dataconverter.ti.com
DLP® Products	www.dlp.com
DSP	dsp.ti.com
Clocks and Timers	www.ti.com/clocks
Interface	interface.ti.com
Logic	logic.ti.com
Power Mgmt	power.ti.com
Microcontrollers	microcontroller.ti.com
RFID	www.ti-rfid.com
OMAP Applications Processors	www.ti.com/omap
Wireless Connectivity	www.ti.com/wirelessconnectivity

Applications

Automotive and Transportation	www.ti.com/automotive
Communications and Telecom	www.ti.com/communications
Computers and Peripherals	www.ti.com/computers
Consumer Electronics	www.ti.com/consumer-apps
Energy and Lighting	www.ti.com/energy
Industrial	www.ti.com/industrial
Medical	www.ti.com/medical
Security	www.ti.com/security
Space, Avionics and Defense	www.ti.com/space-avionics-defense
Video and Imaging	www.ti.com/video

TI E2E Community

e2e.ti.com