

Understanding the Functional Enhancements of DSP/BIOS II and their Utilization in Real-Time DSP Applications

David Dart

Software Development Systems

ABSTRACT

Building upon the software foundation architecture of DSP/BIOS, DSP/BIOS II extends Texas Instruments' system software capabilities and services to enable the development of more advanced and complex applications; and provide support for traditional concurrent-system design paradigms.

This support is a combination of additional software modules and modifications to the kernel to provide preemptive multitasking services common in many embedded microkernels. These additional services enable applications to extend the event-driven, interrupt service paradigms of DSP/BIOS to take advantage of traditional multitasking services, mailboxes, queues, semaphores, and resource protection locks. DSP/BIOS II integrates these traditional concurrent-system kernel services with DSP/BIOS into a fully scalable kernel.

DSP/BIOS II also integrates another device-independent data streaming I/O model to complement the existing DSP/BIOS data pipes, giving developers the opportunity to select the optimum I/O mechanism for their application.

DSP/BIOS II extends the statically configured kernel object programming model of DSP/BIOS with the ability to do dynamic kernel object creations. DSP/BIOS II adds runtime memory management providing dynamic memory allocations and de-allocations. Applications can now manage and configure resources dynamically, enabling developers to build self-configuring and more complex applications whose mix of functions changes over its lifetime.

Contents

1	DSP/BIOS II Overview	3
2	DSP/BIOS II Kernel API Overview	3
3	DSP/BIOS II Kernel Extensions	7
	3.1 Execution Threads in DSP/BIOS II	8
	3.2 Extending DSP/BIOS II Execution Thread Model with Synchronized Tasks	9
	3.3 Semaphores Synchronize Events and Operations	10
	3.4 Data Transfers in DSP/BIOS II	11
	3.4.1 Data Pipes Offer Universal I/O	11
	3.4.2 Data Streams Provide Abstraction and Flexibility	12
	3.4.3 DEV Module Provides Device Independence to SIO	14
	3.5 Creation and Configuration	15
	3.5.1 Dynamic Memory Management	15
4	Epilog	16

List of Figures

Figure 1.	DSP/BIOS II Kernel Execution Threads	7
Figure 2.	DSP/BIOS II Thread Priorities	8
Figure 3.	DSP/BIOS II Synchronized Task States	10
Figure 4.	DSP/BS II Streaming I/O Data Streams	12
Figure 5.	Streaming I/O Device-Dependent (DEV) and Independent (SIO) Modules	13
Figure 6.	Stacking Device Drivers	14

1 DSP/BIOS II Overview

When Texas Instruments introduced DSP/BIOS with Code Composer Studio IDE (CCS), developers secured the opportunity to leverage decades of real-time DSP software experience. By building applications on top of the powerful, yet diminutive system services afforded by DSP/BIOS, developers focus their efforts on their applications rather than the underlying infrastructure. Furthermore, with the underlying infrastructure instrumented for real-time trace and analysis, developers use these services to develop applications that are more reliable.

DSP/BIOS II is an evolutionary next step in the DSP software infrastructure provided by Texas Instruments. It builds on the functionality and services previously provided in DSP/BIOS. The additional features of DSP/BIOS II enable both traditional and new programming methods and paradigms. With the introduction of DSP/BIOS II, the kernel and I/O services extend to encompass and enable a greater variety of application demands.

DSP/BIOS II additionally provides a host-side utility to assist developers during application debug. The DSP/BIOS II Kernel Object Viewing CCS plug-in provides developers visibility into kernel objects during program execution. Developers can use this tool to view the current state of the executions threads or the properties of kernel objects used within the application.

Some DSP applications may require processing commensurate with the data presented. An application may need to create and configure a processing thread on-the-fly to accommodate the incoming data. For example, voice applications in Internet Telephony often need to negotiate and dynamically attach the correct Vocoder to the incoming voice channel. Applications of this class take on a dynamic nature. Both the instantiation and configuration of the data path are dynamic. To accommodate dynamic system requirements, DSP/BIOS II adds support for dynamic real-time kernel object creation and configuration. With DSP/BIOS II, developers choose when to create execution threads and data I/O channels. For optimal performance, applications might create most kernel objects statically up front with the configuration tool, and create others dynamically on the fly, commensurate with the application's runtime demands.

This document describes the new extensions in DSP/BIOS II and compares them to the pre-existing services of DSP/BIOS. In the first section of this document, we present the feature-function matrix to provide an overview of the module extensions and new modules added in DSP/BIOS II. The balance of the paper discusses these extensions in more detail.

2 DSP/BIOS II Kernel API Overview

DSP/BIOS II provides real-time system services and kernel components for Texas Instruments C5000 and C6000 DSP applications. These services and components are available through a combination of design-time declarations using the DSP/BIOS II Configuration Tool, and a runtime library.

Some DSP/BIOS II components require declaration and configuration within the configuration tool and have no Application Programming Interfaces (APIs) exposed. Other services expose a set of APIs to allow developers to access them within their application. The DSP/BIOS II kernel implements run-time services that the target application program invokes through the DSP/BIOS II APIs.

Individual DSP/BIOS II modules in general will manage one or more instances of a related class of kernel objects, and will rely upon global parameter values to control their overall behavior.

We group the DSP/BIOS II modules into six functional classifications (see Table 1):

- Kernel Execution Treads
- Hardware Abstraction
- Device – Independent I/O
- Inter – Thread Communication and Synchronization
- Real-Time Analysis
- Miscellaneous System Services

The DSP/BIOS II **Kernel Execution Thread** modules provide developers with four distinct classes of execution threads. Each thread class provides different execution, preemption and suspension characteristics. DSP/BIOS II supports the two high-priority interrupt threads and the background idle processing loop previously provided in DSP/BIOS. The **HWI**, **SWI**, and **IDL** modules manage these kernel execution threads. In addition, DSP/BIOS II provides a new multitasking thread class, Synchronized Tasks, capable of synchronously suspending and resuming their execution at any point in their execution. Tasks are also capable of arbitrarily adjusting their own (or any other task's) priority during program execution. Synchronized tasks form the basis of traditional concurrent system design. The **TSK** module manages this kernel execution thread.

DSP/BIOS II provides logical interfaces to access and configure certain hardware components independent of the hardware implementation. The **Hardware Abstraction** modules simplify migrating applications across supported ISAs by abstracting the device-dependent components like the system memory map, on-chip timer, and hardware interrupts. As in previous versions of DSP/BIOS, the DSP/BIOS II configuration tool provides a visual editor to define the system memory map, interrupt vector table, and program the on-chip timer. DSP/BIOS II offers new APIs that provide dynamic allocation and freeing of memory during runtime exposed in the **MEM** module.

The **Device-Independent I/O** modules provide services to perform data transfers, whether transferring data between the DSP and a peripheral or between multiple execution threads. DSP/BIOS II continues to support the data pipes previously provided in DSP/BIOS. In addition, DSP/BIOS II supports an alternative I/O model known as data streams. Data pipes are small and fast universal components that transfer data between a reader and a writer thread. Data streams add greater flexibility in the buffering schemes to enable broader application requirements. Data streams rely on one or more underlying device drivers. The device driver model encapsulates the *device-dependent* properties and methods. Device drivers are also capable of performing processing operations on data passing through them through a mechanism known as stacking. Stacking device drivers offer the ability to pipeline processing operations such as data type conversions, scaling or filtering in the data path. The **PIP** and **SIO** modules manage data transfer services in target applications. **SIO** also has an accompanying device driver module **DEV**, which sources or sinks data with **SIO**.

The new **Inter-Thread Communication and Synchronization** modules support the multitasking services provided in DSP/BIOS II. Primary among these are semaphores. Tasks use semaphores to synchronize resource access. Typical examples of resource synchronization include exchanging buffers of data with peripherals, or arbitrating access to shared memory. Semaphores also control synchronizing the execution of multiple tasks. The **SEM** module manages semaphores. The **LCK** module provides shared resource arbitration or mutual exclusion.

Data queues provide inter-thread communication or messaging, without synchronization. Queues are available for all DSP/BIOS II threads. Mailboxes are similar to data queues, except they provide synchronization, and are ideally suited for inter-task communication. The **QUE** module manages data queues, and the **MBX** manager manages the mailbox module.

The **Real-Time Analysis** (RTA) modules enable developers to take advantage of real-time interaction and diagnostics with the DSP application while it is executing in real time. These modules provide APIs to programmatically instrument user-developed components for real-time interaction and diagnostics. The **LOG**, **STS**, and **TRC** modules manage these functions. Critical to real-time analysis is the ability to transfer data between the host and the target. As with previous versions of DSP/BIOS, DSP/BIOS II also provides APIs to manage the transfers. The **HST** and **RTDX** modules manage these functions.

DSP/BIOS II adds new services for program error handling and termination patterned after some standard C-library services grouped under **Miscellaneous System Services** modules. The **SYS** module manages these functions. **SYS** also includes variations on `printf()`.

Atomic functions are non-preemptive operations that execute with interrupts disabled. Operations include AND, OR, INC, DEC, SET and CLEAR. The **ATM** module manages these functions.

Table 1. DSP/BIOS II Feature and Function Matrix

Function / Feature	DSP BIOS	DSP BIOS II	Module	Description	Object Creations		Language Support	
					STATIC Requires the Configuration Tool	DYNAMIC Allowed within the program	C routines	ASM routines
Real Time Analysis								
Event Logging	X	X	LOG	Message Log Manager	X		X	X
Statistics Accumulation	X	X	STS	Statistics Accumulator Manager	X		X	X
Trace Control	X	X	TRC	Trace Manager	X		X	X
File Streaming	X	X	HST	Host I/O Manager	X		X	X
Real Time Data Exchange	X	X	RTDX	Target to Host Communication Manager	X		X	X
Hardware Abstraction								
On-Chip Timer	X	X	CLK	System Clock Manager	X		X	X
Hardware Interrupts	X	X	HWI	Hardware Interrupt Manager	X			X
Static Memory Management	X	X	MEM *	Memory Segment Manager	X			
Dynamic Memory Management		X	MEM **	Memory Segment Manager		X	X	
Kernel Execution Threads								
Software Interrupts	X	X	SWI	Software Interrupt Manager	X		X	X
Periodic Functions	X	X	PRD	Periodic Function Manager	X		X	X
Tasks		X	TSK	Multitasking Manager	X		X	
Idle Loop	X	X	IDL	Idle Function and Processing Loop Manager	X		X	X
Device – Independent I/O								
Data Pipes	X	X	PIP	Data Pipe Manager	X		X	X
Data Streams		X	SIO	Stream I/O Manager	X		X	
Inter – Thread Communication and Synchronization								
Semaphores		X	SEM	Semaphore Manager	X		X	X
Resource Locks		X	LCK	Resource Lock Manager	X		X	X
Mailboxes		X	MBX	Mailbox Manager	X		X	X
Queues		X	QUE	Queue Manager	X		X	X
Miscellaneous System Services								
Atomic Functions (optimized and non-preemptive)		X	ATM	Atomic Functions written in Assembly Language	N/A		N/A	X
Error handling and program termination		X	SYS	System Services Manager	N/A		N/A	X

* Using the Configuration Tool, memory segments are defined and named.

** Once named, this module provides allocation and freeing services.

3 DSP/BIOS II Kernel Extensions

Building upon the real-time foundation of DSP/BIOS, DSP/BIOS II provides additional capabilities and support for traditional concurrent-system design paradigms. This support is a combination of additional software modules, and modifications to the kernel to allow both static and dynamic instantiation and configuration of kernel objects.

Chief among the new software modules is the support for an additional level of thread execution — synchronized tasks. Tasks are independent threads of execution running at lower priority than software interrupts, but above the idle loop (see Figure 1). DSP/BIOS II introduces semaphores to synchronize the scheduling of tasks and access to resources. DSP/BIOS II makes use of semaphores internally within a whole new class of kernel objects including mailboxes, resource locks, and data streams.

DSP/BIOS II provides an additional data transfer module using frame-based data streams with semaphore synchronization and programmable buffer management. Integral to this data streaming module is a corresponding device driver module. This device driver module also supports pipelined operations using a mechanism known as *stacking*.

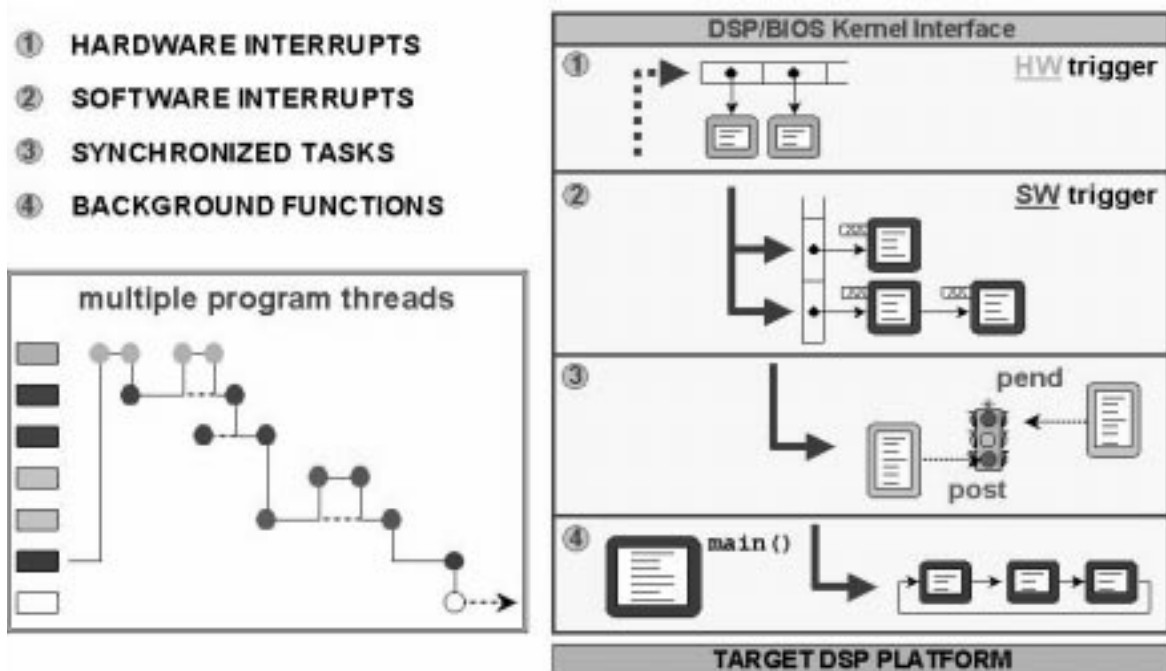


Figure 1. DSP/BIOS II Kernel Execution Threads

New with DSP/BIOS II is support for dynamic creation and configuration of many kernel objects. To enable this, DSP/BIOS II provides additional support for dynamic memory allocation and deallocation. In DSP/BIOS, the configuration tool declared and configured all the kernel objects. DSP/BIOS II additionally allows creation and configuration of many of the kernel objects during program execution. This enables a frontier of applications that need flexible and dynamic services. Applications can dynamically create, configure, or reconfigure the processing threads based on the data presented. For example, telephony applications can create and configure voice channels dynamically with each incoming call, then release them when the call terminates.

3.1 Execution Threads in DSP/BIOS II

Real-time systems deployed on digital signal processors (DSPs) typically require the use of hardware interrupts to trigger processing of external events. Once triggered, the DSPs will invoke an interrupt service routine to perform the required operations needed to service the external event, such as acquiring or delivering data to a peripheral device. These interrupt service routines, or ISRs, will preempt the normal processing operations of the DSP; and once activated, the ISR will execute until it is either finished, or it is preempted by another hardware interrupt if so enabled. The ISR executes in a model known as *run-to-completion*. This execution model does not provide a mechanism to stop execution once started, except by preemption. In DSP/BIOS II, the HWI module manages the hardware interrupts. The HWI module provides mechanisms to enable and disable executing hardware interrupts, locate the interrupt service tables, and locate the HWI dispatcher. HWI executes in the context of the application, and uses the application's stack.

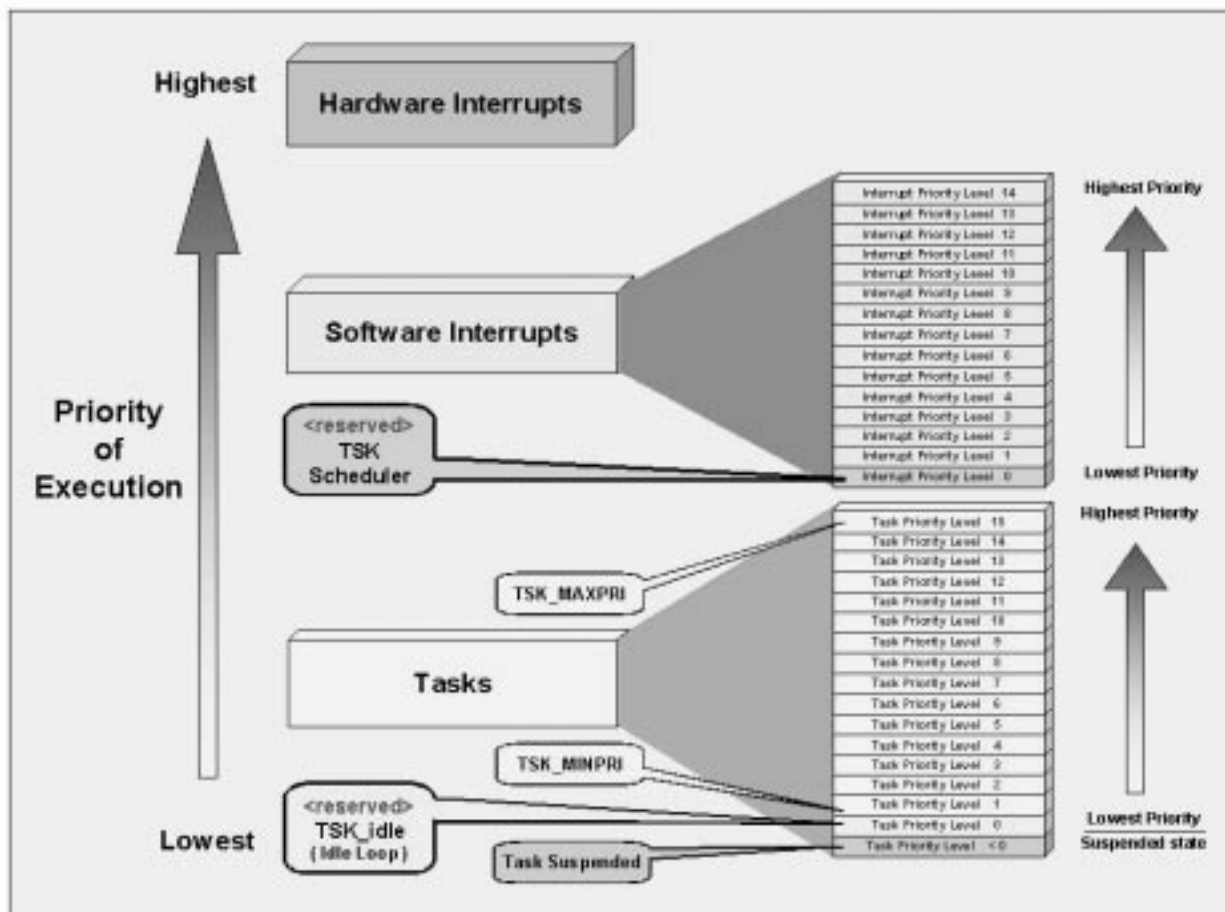


Figure 2. DSP/BIOS II Thread Priorities

Similar to the hardware interrupt is an execution model introduced in previous versions of DSP/BIOS — software interrupts (SWI). As the name implies, this execution model is similar to the hardware interrupt, except all software interrupts execute at a lower priority than hardware interrupts, and a software interrupt cannot interrupt itself. DSP/BIOS II software interrupts are priority-based, providing 14 levels of priority (see Figure 2).

Like the hardware ISR, the software ISR also executes in a run-to-completion mode. Software interrupts can only be preempted by a higher-priority software interrupt or a hardware interrupt. It is quite possible that an application will have multiple software interrupts sharing the same priority level. In this case, software interrupts of the same priority execute on a first-come (posted) first-served basis.

In DSP/BIOS II, the SWI module manages software interrupts. The SWI manager automatically saves processor registers on a stack before executing the software interrupt service routine. SWI executes in the context of the application, and uses the application's stack. Software interrupts execute using the same stack as hardware interrupts so the developer does not need to allocate an additional amount of memory to contain a separate SWI stack.

3.2 Extending DSP/BIOS II Execution Thread Model with Synchronized Tasks

DSP/BIOS II extends the execution thread model of DSP/BIOS with the addition of synchronized task threads (**TSK**). Unlike the hardware and software interrupt model previously described, synchronized tasks in DSP/BIOS II run either until they are complete, preempted, or until they are blocked awaiting a resource or some event to occur. Tasks suspend or terminate their execution themselves by explicitly yielding, sleeping or pending on a resource or event synchronized by a semaphore. In contrast, both hardware and software interrupts cannot yield, sleep, or suspend (block) — they always run to completion unless preempted.

DSP/BIOS II tasks form the basis of traditional concurrent processing by partitioning applications into independent threads of execution synchronized with semaphores or the system clock. Tasks dynamically modify their execution through semaphore synchronization services.

In DSP/BIOS II the TSK module manages and schedules tasks. DSP/BIOS II tasks are priority-based, supporting 15 levels of priority plus a suspended state (negative task priority level). The lowest priority level (0) executes the Idle-loop (see Figure 2). DSP/BIOS II permits both static and dynamic creation of tasks. However, only tasks created statically using the Configuration Tool are visible in the real-time analysis execution graph.

Each TSK task object (shown in Figure 3) is always in one of four possible states of execution:

1. running, which means the task is the one actually executing on the system's processor;
2. ready, which means the task is scheduled for execution subject to processor availability;
3. blocked, or suspended which means the task cannot execute until a particular event occurs within the system; or
4. terminated, which means the task is "terminated" and does not execute again.

Tasks are scheduled for execution according to a priority level assigned to the application.

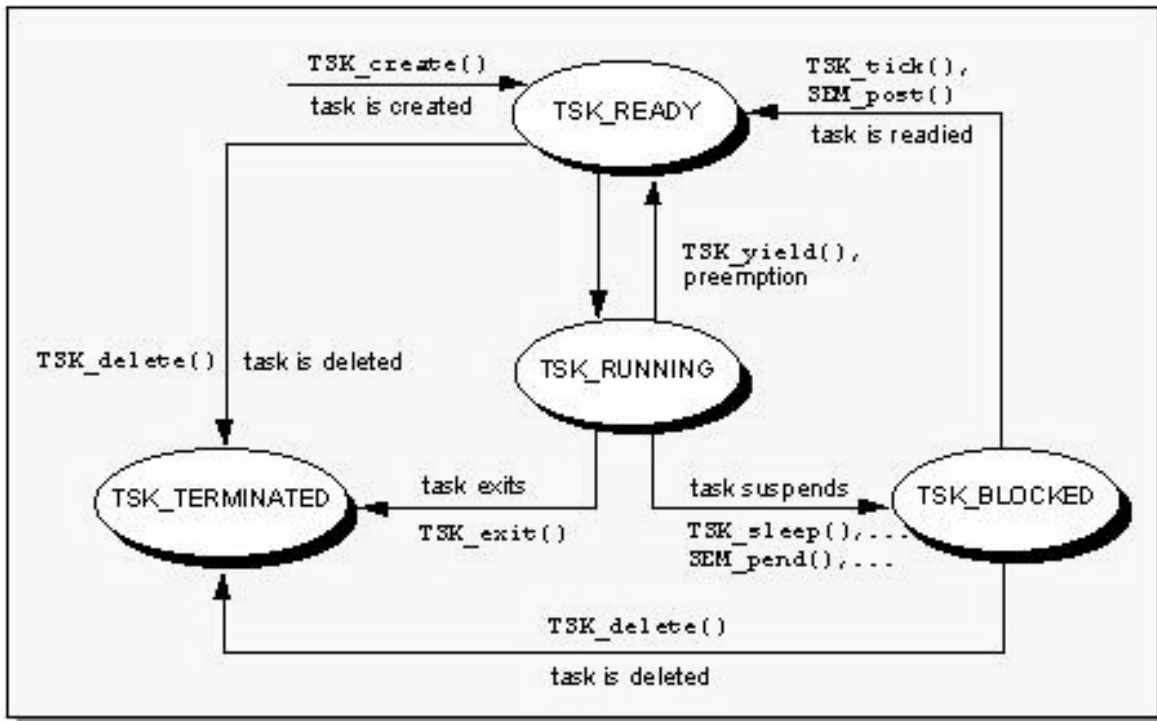


Figure 3. DSP/BIOS II Synchronized Task States

Unlike HWI and SWI objects, tasks require their own configurable stack. This is an important point. The main advantage of SWI over TSK is that SWIs execute on a single stack. Designers leverage this into their designs by partitioning their application into multiple execution threads. In general, applications use SWIs for critical threads and TSKs for the normal multitasking threads. By partitioning to use both SWIs and TSKs, both the application and task stacks are smaller than if only one model is used.

DSP/BIOS II tasks are very similar to task threads in other popular embedded real-time kernels such as PSOS, VxWorks, and Nucleus. Developers familiar with the tasking paradigms of these kernels will find DSP/BIOS II easy to grasp and familiar.

3.3 Semaphores Synchronize Events and Operations

In DSP/BIOS II, the SEM module manages semaphores. DSP/BIOS II implements counting semaphores. Counting semaphores keep an internal count of the number of corresponding resources available. When count is greater than zero, tasks do not block when acquiring a semaphore. This enables developers to implement a variety of synchronization schemes from mutual exclusion to event counting.

PSOS and VxWorks are trademarks of WINDRIVER Systems. Nucleus is a trademark of Advanced Technology Inc.

The semaphore count is initialized to *count* when it is created. In general, *count* is set to the number of resources that the semaphore is synchronizing. Semaphore operations include pending (decrementing the internal counter), and posting (incrementing the internal counter). `SEM_pend()` waits for a semaphore. `SEM_post()` is used to signal a semaphore. If the semaphore count is greater than 0, `SEM_pend()` simply decrements the count and returns. Otherwise, `SEM_pend()` waits for the semaphore to be posted by `SEM_post()`. If a task is waiting for the semaphore, `SEM_post()` removes the task from the semaphore queue and puts it on the ready queue. If no tasks are waiting, `SEM_post()` simply increments the semaphore count and returns.

DSP/BIOS II semaphores include a time-out parameter to enable applications to select between waiting for a resource indefinitely before returning, waiting for a maximum amount of time, or returning immediately based on this value.

Tasks use semaphores to synchronize their execution. Typical examples are tasks waiting for data to become available from a peripheral device, or synchronizing access to shared memory or critical data with other tasks.

DSP/BIOS II provides a dedicated kernel object to synchronize access to a shared resource. Resource locks, managed by the LCK module, arbitrate access to shared resources among several competing tasks. Lock objects are semaphores with an additional property. A task that currently owns the resource may request the resource lock any number of times without blocking. Under similar conditions, a regular semaphore may block the calling `Task` on the second call.

The MBX module manages mailboxes, which use an internal semaphore for access synchronization. Typically, mailboxes provide a utility to transfer messages or data between tasks that execute at different rates or priorities. Content of the mailbox is mostly user-defined to allow maximum flexibility in deployment. Tasks post messages to the mailbox or pend waiting for messages to arrive.

3.4 Data Transfers in DSP/BIOS II

Fundamental to all DSP applications is the need to acquire data, process it, and output the results. DSP applications typically process blocks of data at a time rather than a single datum. So these applications will move continuous blocks of data in from a source, process it, and output the results. Conceptually, this movement of data blocks forms a stream of data flowing in one direction from source to sink. These data streams allow I/O and processing to occur at different rates due to the ability to manage multiple frames asynchronously. That is, while a device is currently filling one buffer with data, the DSP is processing a previously loaded buffer.

DSP/BIOS II provides services to move blocks of data between peripheral devices and applications by two primary mechanisms: data streams and pipes. In both cases, DSP/BIOS II does not copy data from source to destination — this is extremely inefficient. Rather, DSP/BIOS II uses a buffer passing mechanism using address pointers, that minimizes the amount of actual data being moved, and transfers the buffers without any copying.

3.4.1 Data Pipes Offer Universal I/O

DSP/BIOS originally introduced the data pipe mechanism through the PIP module. Data pipes are unidirectional and pass buffers of data, or frames, between any two DSP/BIOS II execution threads. These threads are associated with hardware peripherals (via HWI), software interrupts, tasks, or the background Idle loop. Each data pipe specifically associates one reader and one writer function to access the data pipe.

Data pipes maintain their own memory pool, partitioned into a fixed-number of fixed-length buffers or *frames*. The data pipe *exchanges* these frames between a reader and the data pipe; and between the data pipe and the writer. Applications using data pipes must synchronize the requests for frames (either full or empty) with replacing them back to the pipe. For each frame request, the application must replace the frame back to the pipe before requesting another frame.

The configuration tool is required to create and configure data pipes in DSP/BIOS II. DSP/BIOS II does not supply runtime APIs to create data pipes during program execution.

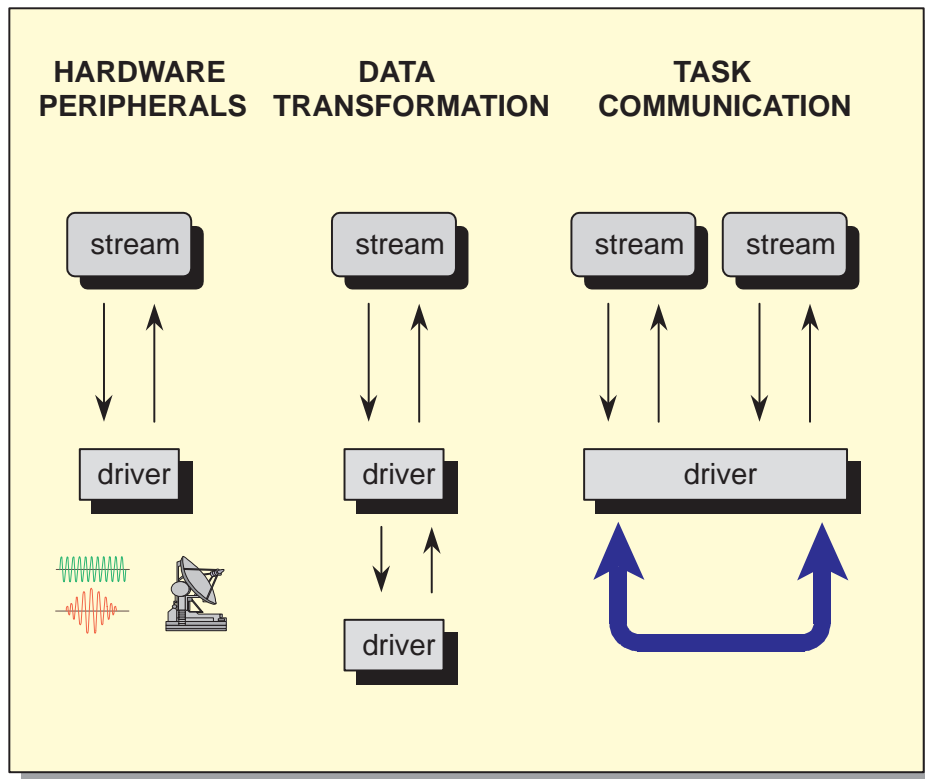


Figure 4. DSP/BIOS II Streaming I/O Data Streams

3.4.2 Data Streams Provide Abstraction and Flexibility

DSP/BIOS II introduces an alternative device-independent I/O model called data streams (see Figure 4). DSP/BIOS II data streams transfer data between the application and any device through a consistent interface exposed through the streaming I/O module, SIO. Manipulations of the data streams include operations to open and close the streaming device, start, stop, and flush the stream, and to provide control.

In DSP/BIOS II, the SIO module manages data streams at the application level. To the application, all devices *appear* the same due to the **device-independent** abstraction inherent in SIO. Applications make calls to `SIO_get()` (or `SIO_reclaim()`) and `SIO_put()` (or `SIO_issue()`), regardless of the underlying device.

Data streams offer flexible buffer management. Data buffers, or frames used by SIO may come from any source including SIO. SIO supports two modes of operation, a standard mode in which SIO creates the specified buffers, or an issue-reclaim mode in which the buffers are supplied to SIO by the application.

In addition, another feature of the data stream model implemented in SIO is selection between a synchronized frame exchange scheme and a scheme that allows the application to independently control issuing and reclaiming frame. In SIO, frames exchange between the calling task and the SIO module. If the SIO channel is reading data from a device, then the application will exchange an empty frame for a filled frame.

The standard SIO model (**SIO_STANDARD**), exchanges data frames, one at a time, with each call to `SIO_get()` or `SIO_put()` automatically. To obtain a new non-empty frame from the device, the application calls `SIO_get()` with one of the arguments being a pointer to an empty frame that the application has finished processing. SIO will exchange frames by assigning the frame pointer passed by the application to the non-empty frame. Meanwhile SIO will consume the empty frame in this exchange process. Important to note is that this exchange process does not copy any data. Through this exchange mechanism, the overhead associated with the exchange is constant, independent of frame size.

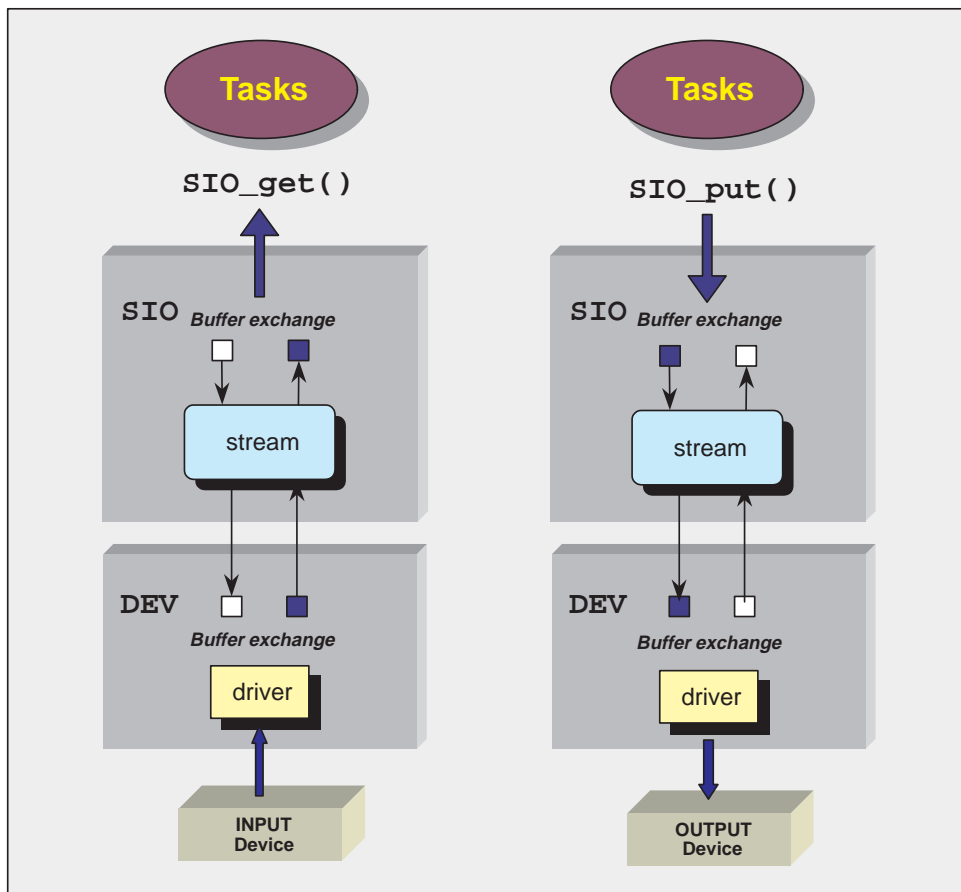


Figure 5. Streaming I/O Device-Dependent (DEV) and Independent (SIO) Modules

An alternative model (**SIO_ISSUERECLAIM**) allows the developer to independently control frame transfers. In the issue-reclaim model, the developer issues any number of frames (`SIO_issue()`) and reclaims them (`SIO_reclaim()`) independently under program control. In both models, SIO preserves the order of the frames exchanged by an internal set of queues, one for frames going to the device, and the other for frames coming from the device.

3.4.3 DEV Module Provides Device Independence to SIO

Complementing and interacting with SIO is the DEV module that manages the **device-dependent** drivers (see Figure 5). DEV provides a device-independent interface to SIO. SIO provides a device-independent interface to the DSP application.

The underlying device driver performs a variety of functions including operations to open/close, start/stop, stream data, and initialize the driver. DSP/BIOS II supports two classes of device drivers, terminating and stacking (see Figure 6).

Terminating drivers perform the classical I/O with peripherals such as CODECS. These device drivers often contain interrupt service routines (ISRs) to interact with the peripheral, however I/O by other means such as polling is equally permissible. Terminating device drivers supply SIO with either a data source or sink.

Stacking drivers are a special class of drivers that perform inline, pipelined processing. Stacking device drivers modify the data passing through them. A typical use of stacking device drivers is to implement operations such as scaling, filtering, or converting between data formats. DSP/BIOS II supports two classes of stacking device drivers: in-place, and copying. Copying drivers allows the input and output buffers to be different sizes. This is useful for stacking devices that perform compression/decompression or packing and unpacking. If encapsulation and reuse of pipeline operations is beneficial, they become candidates for implementation as a stacking device driver. Developers create SIO channels specifying the terminating device driver and any stacking driver(s) to form the complete data path between the task and the device

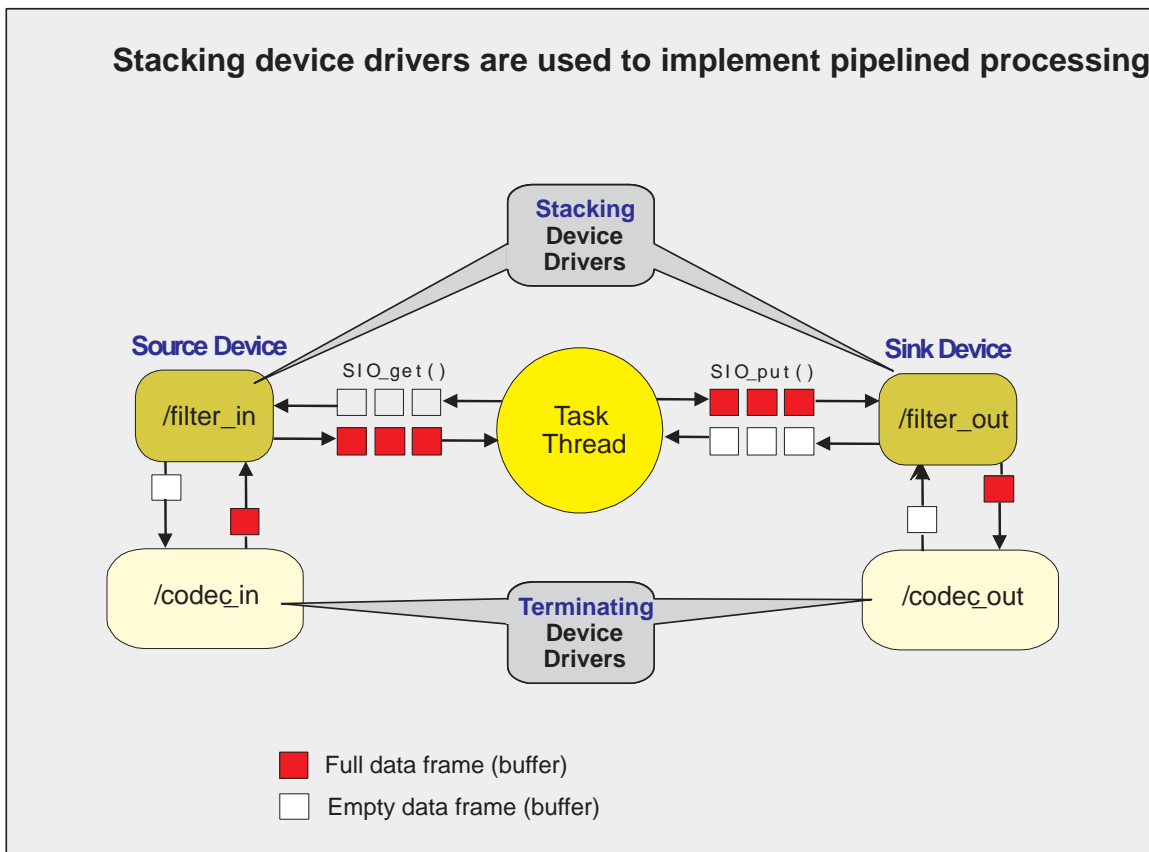


Figure 6. Stacking Device Drivers

3.5 Creation and Configuration

A principle goal behind the creation of DSP/BIOS II is to minimize resource utilization, both memory and CPU, by the system services themselves, freeing these valuable resources for applications. To achieve this, DSP/BIOS II is by design both scalable and minimal in size. Scalability results from the modular construction of the DSP/BIOS II runtime library and APIs. The DSP/BIOS II scalable architecture allows binding only those modules that the application uses to minimize memory utilization. To further minimize resource utilization, the configuration tool allows creation and configuration all DSP/BIOS II objects at design-time. This process is termed **static creation and configuration**. Objects created in the configuration tool combine into an executable image and link with your application. Static creation and configuration minimizes DSP system resource requirements and provides optimum performance.

If the program architecture (or framework), or availability of a kernel object or thread is required for the entire life of the program, then these elements require static declaration and configuration using the configuration tool. In the previous version, DSP/BIOS required the configuration tool for all kernel objects.

Static creation does not limit the ability to design dynamic systems. Developers can implement dynamic systems with statically created objects. To accommodate this in static environments requires preallocating enough resources to handle all the possible requirements. For some systems, this can consume more resources than are actually required at any given time. Overlaying dynamic operation on a static infrastructure in complex systems is challenging, often leading to implementations that are unclear and difficult to maintain.

To achieve application flexibility easily, DSP/BIOS II introduces **dynamic creation and configuration**. DSP/BIOS II permits the dynamic creation, deletion, and configuration of many kernel objects within an application during runtime (see Table 1). Using this paradigm, DSP/BIOS II applications may create tasks to perform operations as needed, and delete them when finished. Similarly, dynamic creation and deletion of data paths between devices and Tasks is achievable. You can delete objects created dynamically. You cannot delete static objects. Only static objects declared using the configuration tool contain instrumentation and are visible to the real-time analysis tools within CCS. Dynamically created objects are not visible to the RTA tools. However all kernel object, static and dynamic, are visible to the Kernel Object Viewing CCS plug-in during application debug.

Typically, tradeoffs between application flexibility, performance, and resource utilization are required. DSP/BIOS II provides developers a choice based on their application requirements.

3.5.1 Dynamic Memory Management

DSP/BIOS II provides dynamic memory management APIs to allow applications to allocate and free memory during runtime. Developers define the memory segments to allocate from using the configuration tool and through runtime APIs. During runtime, the application makes API calls to the MEM module to allocate and free memory. Note that memory fragmenting may occur from frequent allocation and freeing of memory. To minimize memory footprint, DSP/BIOS II does not provide defragmentation services or garbage collection. The developer must program the application to perform these services if they are required.

4 Epilog

With DSP/BIOS II, developers leverage a wealth of system services with the knowledge that Texas Instruments will ensure their availability in the future. These system services offer developers real time components to build their application on.

Advancements in DSP application complexities require more efficient application architectures that fully leverage the resources of DSP systems. This includes both memory and CPU requirements. Some applications are well suited for the static programming model of DSP/BIOS II. Some will require dynamic capabilities to make optimal use of the resources, program performance, and cost. DSP/BIOS II provides a mixture of services and capabilities that the developer can match against their application requirements. Combining the static, global nature of objects created through the configuration tool with the dynamic, localized nature of objects created within the program provides developers with the infrastructure to fully optimize their application's performance and resource utilization.

Building on DSP/BIOS, new modules and components available in DSP/BIOS II add semaphore-based components and traditional multitasking programming paradigms. DSP developers can select which programming model is best suited to their application. Data streams using semaphore synchronization offer developers an alternative method for I/O that can include pipelining functions in the stream. Using stacking device drivers, pipelined operations on data flowing in the stream is configured on the fly as dictated by the data flowing into the application.

IMPORTANT NOTICE

Texas Instruments and its subsidiaries (TI) reserve the right to make changes to their products or to discontinue any product or service without notice, and advise customers to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgment, including those pertaining to warranty, patent infringement, and limitation of liability.

TI warrants performance of its semiconductor products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Customers are responsible for their applications using TI components.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards must be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance or customer product design. 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 of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used. TI's publication of information regarding any third party's products or services does not constitute TI's approval, warranty or endorsement thereof.