

Introduction

This chapter discusses how to develop programs for the Nios® II processor based on the Altera® hardware abstraction layer (HAL). This chapter contains the following sections:

- “The Nios II Project Structure” on page 6–2
- “The system.h System Description File” on page 6–4
- “Data Widths and the HAL Type Definitions” on page 6–5
- “UNIX-Style Interface” on page 6–5
- “File System” on page 6–6
- “Using Character-Mode Devices” on page 6–8
- “Using File Subsystems” on page 6–15
- “Using Timer Devices” on page 6–16
- “Using Flash Devices” on page 6–19
- “Using DMA Devices” on page 6–25
- “Using Interrupt Controllers” on page 6–30
- “Reducing Code Footprint” on page 6–30
- “Boot Sequence and Entry Point” on page 6–37
- “Memory Usage” on page 6–39
- “Working with HAL Source Files” on page 6–44

The application program interface (API) for HAL-based systems is readily accessible to software developers who are new to the Nios II processor. Programs based on the HAL use the ANSI C standard library functions and runtime environment, and access hardware resources with the HAL API's generic device models. The HAL API largely conforms to the familiar ANSI C standard library functions, though the ANSI C standard library is separate from the HAL. The close integration of the ANSI C standard library and the HAL makes it possible to develop useful programs that never call the HAL functions directly. For example, you can manipulate character mode devices and files using the ANSI C standard library I/O functions, such as `printf()` and `scanf()`.



This document does not cover the ANSI C standard library. An excellent reference is *The C Programming Language, Second Edition*, by Brian Kernighan and Dennis M. Ritchie (Prentice-Hall).

Nios II Development Flows


The Nios II Embedded Design Suite (EDS) provides two distinct development flows for creating Nios II programs. You can use the Nios II Software Build Tools, or work in the Nios II integrated development environment (IDE). These two approaches use the HAL in the same way.

 In most cases, you should create new projects using either the Nios II Software Build Tools for Eclipse™ or the Software Build Tools command line. The IDE will be discontinued in a future release. IDE support is for the following situations:

- Working with pre-existing Nios II IDE software projects
- Creating new projects for the Nios II C2H compiler
- Debugging with the FS2 console

HAL BSP Settings

Every Nios II board support package (BSP) has settings that determine the BSP's characteristics. For example, HAL BSPs have settings to identify the hardware components associated with standard devices such as `stdout`. Defining and manipulating BSP settings is an important part of Nios II project creation. You manipulate BSP settings with the Nios II BSP Editor, with command-line options, or with Tcl scripts.

 For details about how to control BSP settings, refer to one or more of the following documents:

- For the Nios II Software Build Tools for Eclipse, refer to the *Getting Started with the Graphical User Interface* chapter of the *Nios II Software Developer's Handbook*.
- For the Nios II Software Build Tools command line, refer to the *Nios II Software Build Tools* chapter of the *Nios II Software Developer's Handbook*.

 For detailed descriptions of available BSP settings, refer to the *Nios II Software Build Tools Reference* chapter of the *Nios II Software Developer's Handbook*.

Many HAL settings are reflected in the **system.h** file, which provides a helpful reference for details about your BSP. For information about **system.h**, refer to “*The system.h System Description File*” on page 6-4.


 Do not edit **system.h**. The Nios II EDS provides tools to manipulate system settings.

The Nios II Project Structure

The creation and management of software projects based on the HAL is integrated tightly with the Nios II Software Build Tools. This section discusses the Nios II projects as a basis for understanding the HAL.

Figure 6-1 shows the blocks of a Nios II program with emphasis on how the HAL BSP fits in. The label for each block describes what or who generated that block, and an arrow points to each block's dependency.

Figure 6-1. The Nios II HAL Project Structure



Every HAL-based Nios II program consists of two Nios II projects, as shown in Figure 6-1. Your application-specific code is contained in one project (the user application project), and it depends on a separate BSP project (the HAL BSP).


The application project contains all the code you develop. The executable image for your program ultimately results from building both projects.

With the Nios II Software Build Tools for Eclipse, the tools create the HAL BSP project when you create your application project. In the Nios II Software Build Tools command line flow, you create the BSP using `nios2-bsp` or a related tool.

The HAL BSP project contains all information needed to interface your program to the hardware. The HAL drivers relevant to your SOPC Builder system are incorporated in the BSP project.

The BSP project depends on the SOPC Builder system, defined by a SOPC Information File (`.sopcinfo`). The Nios II Software Build Tools can keep your BSP up-to-date with the SOPC Builder system. This project dependency structure isolates your program from changes to the underlying hardware, and you can develop and debug code without concern about whether your program matches the target hardware.

You can use the Nios II Software Build Tools to update your BSP to match updated hardware. You control whether and when these updates occur.

 For details about how the Software Build Tools keep your BSP up-to-date with your hardware system, refer to “Revising Your BSP” in the *Nios II Software Build Tools* chapter of the *Nios II Software Developer’s Handbook*.

In summary, when your program is based on a HAL BSP, you can always keep it synchronized with the target hardware with a few simple Software Build Tools commands.

The system.h System Description File

The **system.h** file provides a complete software description of the Nios II system hardware. Not all information in **system.h** is useful to you as a programmer, and it is rarely necessary to include it explicitly in your C source files. Nonetheless, **system.h** holds the answer to the question, “What hardware is present in this system?”

The **system.h** file describes each peripheral in the system and provides the following details:

- The hardware configuration of the peripheral
- The base address
- Interrupt request (IRQ) information (if any)
- A symbolic name for the peripheral

The Nios II Software Build Tools generate the **system.h** file for HAL BSP projects. The contents of **system.h** depend on both the hardware configuration and the HAL BSP properties.



Do not edit **system.h**. The Software Build Tools provide facilities to manipulate system settings.

For details about how to control BSP settings, refer to “[HAL BSP Settings](#)” on [page 6-2](#).

The code in [Example 6-1](#) from a **system.h** file shows some of the hardware configuration options this file defines.

Example 6-1. Excerpts from a system.h File

```

/*
 * sys_clk_timer configuration
 *
 */

#define SYS_CLK_TIMER_NAME "/dev/sys_clk_timer"
#define SYS_CLK_TIMER_TYPE "altera_avalon_timer"
#define SYS_CLK_TIMER_BASE 0x00920800
#define SYS_CLK_TIMER_IRQ 0
#define SYS_CLK_TIMER_ALWAYS_RUN 0
#define SYS_CLK_TIMER_FIXED_PERIOD 0

/*
 * jtag_uart configuration
 *
 */

#define JTAG_UART_NAME "/dev/jtag_uart"
#define JTAG_UART_TYPE "altera_avalon_jtag_uart"
#define JTAG_UART_BASE 0x00920820
#define JTAG_UART_IRQ 1

```

Data Widths and the HAL Type Definitions

For embedded processors such as the Nios II processor, it is often important to know the exact width and precision of data. Because the ANSI C data types do not explicitly define data width, the HAL uses a set of standard type definitions instead. The ANSI C types are supported, but their data widths are dependent on the compiler's convention.

The header file `alt_types.h` defines the HAL type definitions; [Table 6-1](#) shows the HAL type definitions.

Table 6-1. The HAL Type Definitions

Type	Meaning
<code>alt_8</code>	Signed 8-bit integer.
<code>alt_u8</code>	Unsigned 8-bit integer.
<code>alt_16</code>	Signed 16-bit integer.
<code>alt_u16</code>	Unsigned 16-bit integer.
<code>alt_32</code>	Signed 32-bit integer.
<code>alt_u32</code>	Unsigned 32-bit integer.
<code>alt_64</code>	Signed 64-bit integer.
<code>alt_u64</code>	Unsigned 64-bit integer.

[Table 6-2](#) shows the data widths that the Altera-provided GNU toolchain uses.

Table 6-2. GNU Toolchain Data Widths

Type	Meaning
<code>char</code>	8 bits.
<code>short</code>	16 bits.
<code>long</code>	32 bits.
<code>int</code>	32 bits.

UNIX-Style Interface

The HAL API provides a number of UNIX-style functions. The UNIX-style functions provide a familiar development environment for new Nios II programmers, and can ease the task of porting existing code to run in the HAL environment. The HAL uses these functions primarily to provide the system interface for the ANSI C standard library. For example, the functions perform device access required by the C library functions defined in `stdio.h`.

The following list contains all of the available UNIX-style functions:

- `_exit()`
- `close()`
- `fstat()`
- `getpid()`
- `gettimeofday()`

- `ioctl()`
- `isatty()`
- `kill()`
- `lseek()`
- `open()`
- `read()`
- `sbrk()`
- `settimeofday()`
- `stat()`
- `usleep()`
- `wait()`
- `write()`

The most commonly used functions are those that relate to file I/O. Refer to “[File System](#)” on page 6-6.



For details about the use of these functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

File System

The HAL provides infrastructure for UNIX-style file access. You can use this infrastructure to build a file system on any storage devices available in your hardware.




For an example, refer to the *Read-Only Zip File System* chapter of the *Nios II Software Developer's Handbook*.

You can access files in a HAL-based file system by using either the C standard library file I/O functions in the newlib C library (for example `fopen()`, `fclose()`, and `fread()`), or using the UNIX-style file I/O provided by the HAL.

The HAL provides the following UNIX-style functions for file manipulation:

- `close()`
- `fstat()`
- `ioctl()`
- `isatty()`
- `lseek()`
- `open()`
- `read()`
- `stat()`
- `write()`

 For more information about these functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

The HAL registers a file subsystem as a mount point in the global HAL file system. Attempts to access files below that mount point are directed to the file subsystem. For example, if a read-only zip file subsystem (**zipfs**) is mounted as **/mount/zipfs0**, the **zipfs** file subsystem handles calls to `fopen()` for **/mount/zipfs0/myfile**.

There is no concept of a current directory. Software must access all files using absolute paths.

The HAL file infrastructure also allows you to manipulate character mode devices with UNIX-style path names. The HAL registers character mode devices as nodes in the HAL file system. By convention, **system.h** defines the name of a device node as the prefix **/dev/** plus the name assigned to the hardware component in SOPC builder. For example, a UART peripheral **uart1** in SOPC builder is **/dev/uart1** in **system.h**.

The code in [Example 6-2](#) reads characters from a read-only zip file subsystem **rozipfs** that is registered as a node in the HAL file system. The standard header files `stdio.h`, `stddef.h`, and `stdlib.h` are installed with the HAL.

Example 6-2. Reading Characters from a File Subsystem

```
#include <stdio.h>
#include <stddef.h>
#include <stdlib.h>

#define BUF_SIZE (10)


int main(void)
{
    FILE* fp;
    char buffer[BUF_SIZE];

    fp = fopen ("/mount/rozipfs/test", "r"); if (fp == NULL)
    {
        printf ("Cannot open file.\n");
        exit (1);
    }

    fread (buffer, BUF_SIZE, 1, fp);

    fclose (fp);

    return 0;
}
```

 For more information about the use of these functions, refer to the newlib C library documentation installed with the Nios II EDS. On the Windows Start menu, click **Programs > Altera > Nios II > Nios II Documentation**.

Using Character-Mode Devices

A character-mode device is a hardware peripheral that sends and/or receives characters serially. A common example is the UART. Character mode devices are registered as nodes in the HAL file system. In general, a program associates a file descriptor to a device's name, and then writes and reads characters to or from the file using the ANSI C file operations defined in `file.h`. The HAL also supports the concept of standard input, standard output, and standard error, allowing programs to call the `stdio.h` I/O functions.

Standard Input, Standard Output and Standard Error

Using standard input (`stdin`), standard output (`stdout`), and standard error (`stderr`) is the easiest way to implement simple console I/O. The HAL manages `stdin`, `stdout`, and `stderr` behind the scenes, which allows you to send and receive characters through these channels without explicitly managing file descriptors. For example, the HAL directs the output of `printf()` to standard out, and `perror()` to standard error. You associate each channel to a specific hardware device by manipulating BSP settings.

The code in [Example 6-3](#) shows the classic Hello World program. This program sends characters to whatever device is associated with `stdout` when the program is compiled.

Example 6-3. Hello World

```
#include <stdio.h>
int main ()
{
    printf ("Hello world!");
    return 0;
}
```

When using the UNIX-style API, you can use the file descriptors `stdin`, `stdout`, and `stderr`, defined in `unistd.h`, to access, respectively, the standard in, standard out, and standard error character I/O streams. `unistd.h` is installed with the Nios II EDS as part of the newlib C library package.

General Access to Character Mode Devices

Accessing a character-mode device other than `stdin`, `stdout`, or `stderr` is as easy as opening and writing to a file. The code in [Example 6-4](#) writes a message to a UART called `uart1`.

Example 6-4. Writing Characters to a UART

```
#include <stdio.h>
#include <string.h>

int main (void)
{
    char* msg = "hello world";
    FILE* fp;

    fp = fopen ("/dev/uart1", "w");
    if (fp!=NULL)
    {
        fprintf(fp, "%s",msg);
        fclose (fp);
    }
    return 0;
}
```

C++ Streams

HAL-based systems can use the C++ streams API for manipulating files from C++.

/dev/null

All systems include the device **/dev/null**. Writing to **/dev/null** has no effect, and all data is discarded. **/dev/null** is used for safe I/O redirection during system startup. This device can also be useful for applications that wish to sink unwanted data.

This device is purely a software construct. It does not relate to any physical hardware device in the system.

Lightweight Character-Mode I/O

The HAL offers several methods of reducing the code footprint of character-mode device drivers. For details, refer to [“Reducing Code Footprint” on page 6-30](#).

Altera Logging Functions

The Altera logging functions provide a separate channel for sending logging and debugging information to a character-mode device, supplementing `stdout` and `stderr`. The Altera logging information can be printed in response to several conditions. Altera logging can be enabled and disabled independently of any normal `stdio` output, making it a powerful debugging tool.

When Altera logging is enabled, your software can print extra messages to a specified port with HAL function calls. The logging port, specified in the BSP, can be a UART or a JTAG UART device. In its default configuration, Altera logging prints out boot messages, which trace each step of the boot process.



Avoid setting the Altera logging device to the device used for `stdout` or `stderr`. If Altera logging output is sent to `stdout` or `stderr`, the logging output might appear interleaved with the `stdout` or `stderr` output

Several logging options are available, controlled by C preprocessor symbols. You can also choose to add custom logging messages.



Altera logging changes system behavior. The logging implementation is designed to be as simple as possible, loading characters directly to the transmit register. It can have a negative impact on software performance.

Altera logging functions are conditionally compiled. When logging is disabled, it has no impact on code footprint or performance.



The Altera reduced device drivers do not support Altera logging.

Enabling Altera Logging

The Nios II Software Build Tools have a setting to enable Altera logging. The setting is called **hal.log_port**. It is similar to **hal.stdout**, **hal.stdin**, and **hal.stderr**. To enable Altera logging, you set **hal.log_port** to a JTAG UART or a UART device. The setting allows the HAL to send log messages to the specified device when a logging macro is invoked.

When Altera logging is enabled, the Nios II Software Build Tools define **ALT_LOG_ENABLE** in **public.mk** to enable log messages. The build tools also set the **ALT_LOG_PORT_TYPE** and **ALT_LOG_PORT_BASE** values in **system.h** to point to the specified device.

When Altera logging is enabled without special options, the HAL prints out boot messages to the selected port. For typical software that uses the standard **alt_main.c** (such as the Hello World software example), the messages appear as in [Example 6-5](#).

Example 6-5. Default Boot Logging Output

```
[crt0.S] Inst & Data Cache Initialized.
[crt0.S] Setting up stack and global pointers.
[crt0.S] Clearing BSS
[crt0.S] Calling alt_main.
[alt_main.c] Entering alt_main, calling alt_irq_init.
[alt_main.c] Done alt_irq_init, calling alt_os_init.
[alt_main.c] Done OS Init, calling alt_sem_create.
[alt_main.c] Calling alt_sys_init.
[alt_main.c] Done alt_sys_init. Redirecting IO.
[alt_main.c] Calling C++ constructors.
[alt_main.c] Calling main.
[alt_exit.c] Entering _exit() function.
[alt_exit.c] Exit code from main was 0.
[alt_exit.c] Calling ALT_OS_STOP().
[alt_exit.c] Calling ALT_SIM_HALT().
[alt_exit.c] Spinning forever.
```



A write operation to the Altera logging device stalls in **ALT_LOG_PRINTF()** until the characters are read from the Altera logging device's output buffer. To ensure that the Nios II application completes initialization, run the **nios2-terminal** command from the Nios II Command Shell to accept the Altera logging output.

Extra Logging Options

In addition to the default boot messages, logging options are incorporated in Altera logging. Each option is controlled by a C preprocessor symbol. The details of each option are outlined in [Table 6-3](#).

Table 6-3. Altera Logging Options (Part 1 of 2)

Name	Description	
System clock log	Purpose	Prints out a message from the system clock interrupt handler at a specified interval. This indicates that the system is still running. The default interval is every 1 second.
	Preprocessor symbol	ALT_LOG_SYS_CLK_ON_FLAG_SETTING
	Modifiers	<p>The system clock log has two modifiers, providing two different ways to specify the logging interval.</p> <ul style="list-style-type: none"> ■ ALT_LOG_SYS_CLK_INTERVAL—Specifies the logging interval in system clock ticks. The default is <i><clock ticks per second></i>, that is, one second. ■ ALT_LOG_SYS_CLK_INTERVAL_MULTIPLIER—Specifies the logging interval in seconds. The default is 1. When you modify ALT_LOG_SYS_CLK_INTERVAL_MULTIPLIER, ALT_LOG_SYS_CLK_INTERVAL is recalculated.
	Sample Output	<pre>System Clock On 0 System Clock On 1</pre>
Write echo	Purpose	Every time <code>alt_write()</code> is called (normally, whenever characters are sent to <code>stdout</code>), the first <i><n></i> characters are echoed to a logging message. The message starts with the string "Write Echo:". <i><n></i> is specified with ALT_LOG_WRITE_ECHO_LEN. The default is 15 characters.
	Preprocessor symbol	ALT_LOG_WRITE_ON_FLAG_SETTING
	Modifiers	ALT_LOG_WRITE_ECHO_LEN—Number of characters to echo. Default is 15.
	Sample Output	Write Echo: Hello from Nio
JTAG startup log	Purpose	<p>At JTAG UART driver initialization, print out a line with the number of characters in the software transmit buffer followed by the JTAG UART control register contents. The number of characters, prefaced by the string "SW CirBuf", might be negative, because it is computed as (<i><tail_pointer></i> - <i><head_pointer></i>) on a circular buffer.</p> <p>For more information about the JTAG UART control register fields, refer to the Off-Chip Interface Peripherals section in <i>Volume 5: Embedded Peripherals of the Quartus® II Handbook</i>.</p>
	Preprocessor symbol	ALT_LOG_JTAG_UART_STARTUP_INFO_ON_FLAG_SETTING
	Modifiers	None
	Sample Output	<pre>JTAG Startup Info: SW CirBuf = 0, HW FIFO wspace=64 AC=0 WI=0 RI=0 WE=0 RE=1</pre>

Table 6-3. Altera Logging Options (Part 2 of 2)

Name	Description	
JTAG interval log	Purpose	Creates an alarm object to print out the same JTAG UART information as the JTAG startup log, but at a repeated interval. Default interval is 0.1 second, or 10 messages a second.
	Preprocessor symbol	ALT_LOG_JTAG_UART_ALARM_ON_FLAG_SETTING
	Modifiers	The JTAG interval log has two modifiers, providing two different ways to specify the logging interval. <ul style="list-style-type: none"> ■ ALT_LOG_JTAG_UART_TICKS—Logging interval in ticks. Default is <i><ticks_per_second></i> / 10. ■ ALT_LOG_JTAG_UART_TICKS_DIVISOR—Specifies the number of logs per second. The default is 10. When you modify ALT_LOG_JTAG_UART_TICKS_DIVISOR, ALT_LOG_JTAG_UART_TICKS is recalculated.
	Sample Output	JTAG Alarm: SW CirBuf = 0, HW FIFO wspace=45 AC=0 WI=0 RI=0 WE=0 RE=1
JTAG interrupt service routine (ISR) log	Purpose	Prints out a message every time the JTAG UART near-empty interrupt triggers. Message contains the same JTAG UART information as in the JTAG startup log.
	Preprocessor symbol	ALT_LOG_JTAG_UART_ISR_ON_FLAG_SETTING
	Modifiers	None
	Sample Output	JTAG IRQ: SW CirBuf = -20, HW FIFO wspace=64 AC=0 WI=1 RI=0 WE=1 RE=1
Boot log	Purpose	Prints out messages tracing the software boot process. The boot log is turned on by default when Altera logging is enabled.
	Preprocessor symbol	ALT_LOG_BOOT_ON_FLAG_SETTING
	Modifiers	None
	Sample Output	Refer to “Enabling Altera Logging” on page 6-10 .

Setting a preprocessor flag to 1 enables the corresponding option. Any value other than 1 disables the option.

Several options have modifiers, which are additional preprocessor symbols controlling details of how the options work. For example, the system clock log’s modifiers control the logging interval. Option modifiers are also listed in [Table 6-3](#). An option’s modifiers are meaningful only when the option is enabled.

Logging Levels

An additional preprocessor symbol, ALT_LOG_FLAGS, can be set to provide some grouping for the extra logging options. ALT_LOG_FLAGS implements logging levels based on performance impact. With higher logging levels, the Altera logging options take more processor time. ALT_LOG_FLAGS levels are defined in [Table 6-4](#).

Table 6-4. Altera Logging Levels

Logging Level	Logging
0	Boot log (default)
1	Level 0 plus system clock log and JTAG startup log
2	Level 1 plus JTAG interval log and write echo
3	Level 2 plus JTAG ISR log
-1	Silent mode—No Altera logging

Note to Table 6-4:

- (1) You can use logging level -1 to turn off logging without changing the program footprint. The logging code is still present in your executable image, as determined by other logging options chosen. This is useful when you wish to switch the log output on or off without disturbing the memory map.

Because each logging option is controlled by an independent preprocessor symbol, individual options in the logging levels can be overridden.

Example: Creating a BSP with Logging

Example 6-6 creates a HAL BSP with Altera logging enabled and the following options in addition to the default boot log:

- System clock log
- JTAG startup log
- JTAG interval log, logging twice a second
- No write echo

Example 6-6. BSP With Logging

```
nios2-bsp hal my_bsp ../my_hardware.sopcinfo \
--set hal.log_port uart1 \
--set hal.make.bsp_cflags_user_flags \
-DALT_LOG_FLAGS=2 \
-DALT_LOG_WRITE_ON_FLAG_SETTING=0 \
-DALT_LOG_JTAG_UART_TICKS_DIVISOR=2
```

The `-DALT_LOG_FLAGS=2` argument adds `-DALT_LOG_FLAGS=2` to the `ALT_CPP_FLAGS` make variable in `public.mk`.

Custom Logging Messages

You can add custom messages that are sent to the Altera logging device. To define a custom message, include the header file `alt_log_printf.h` in your C source file as follows:

```
#include "sys/alt_log_printf.h"
```

Then use the following macro function:

```
ALT_LOG_PRINTF(const char *format, ...)
```

This C preprocessor macro is a pared-down version of `printf()`. The format argument supports most `printf()` options. It supports `%c`, `%d`, `%I`, `%o`, `%s`, `%u`, `%x`, and `%X`, as well as some precision and spacing modifiers, such as `%-9.3o`. It does not support floating point formats, such as `%f` or `%g`. This function is not compiled if Altera logging is not enabled.

If you want your custom logging message be controlled by Altera logging preprocessor options, use the appropriate Altera logging option preprocessor flags from [Table 6-4](#), or [Table 6-3 on page 6-11](#). [Example 6-7](#) illustrates two ways to implement logging options with custom logging messages.

Example 6-7. Using Preprocessor Flags

```
/* The following example prints "Level 2 logging message" if
   logging is set to level 2 or higher */
#if ( ALT_LOG_FLAGS >= 2 )
    ALT_LOG_PRINTF ( "Level 2 logging message" );
#endif

/* The following example prints "Boot logging message" if boot logging
   is turned on */
#if ( ALT_LOG_BOOT_ON_FLAG_SETTING == 1)
    ALT_LOG_PRINTF ( "Boot logging message" );
#endif
```

Altera Logging Files

[Table 6-5](#) lists HAL source files which implement Altera logging functions.

Table 6-5. HAL Implementation Files for Altera Logging

Location (1)	File Name
<code>components/altera_hal/HAL/inc/sys/</code>	<code>alt_log_printf.h</code>
<code>components/altera_hal/HAL/src/</code>	<code>alt_log_printf.c</code>
<code>components/altera_nios2/HAL/src/</code>	<code>alt_log_macro.S</code>

Note to [Table 6-5](#):

(1) All file locations are relative to `$SOPC_KIT_NIOS2`.

[Table 6-6](#) lists HAL source files which use Altera logging functions. These files implement the logging options listed in [table Table 6-3 on page 6-11](#). They also serve as examples of logging usage.

Table 6-6. HAL Example Files for Altera Logging

Location (1)	File Name
components/altera_avalon_jtag_uart/HAL/src/	altera_avalon_jtag_uart.c
components/altera_avalon_timer/HAL/src/	altera_avalon_timer_sc.c
components/altera_hal/HAL/src/	alt_exit.c
components/altera_hal/HAL/src/	alt_main.c
components/altera_hal/HAL/src/	alt_write.c
components/altera_nios2/HAL/src/	crt0.S

Note to Table 6-6:

(1) All file locations are relative to \$SOPC_KIT_NIOS2.

Using File Subsystems

The HAL generic device model for file subsystems allows access to data stored in an associated storage device using the C standard library file I/O functions. For example, the Altera read-only zip file system provides read-only access to a file system stored in flash memory.

A file subsystem is responsible for managing all file I/O access beneath a given mount point. For example, if a file subsystem is registered with the mount point `/mnt/rozipfs`, all file access beneath this directory, such as `fopen("/mnt/rozipfs/myfile", "r")`, is directed to that file subsystem.

As with character mode devices, you can manipulate files in a file subsystem using the C file I/O functions defined in `file.h`, such as `fopen()` and `fread()`.



For more information about the use of file I/O functions, refer to the newlib C library documentation installed with the Nios II EDS. On the Windows Start menu, click **Programs > Altera > Nios II <version> > Nios II EDS <version> Documentation**.

Host-Based File System

The host-based file system enables programs executing on a target board to read and write files stored on the host computer. The Nios II Studio transmits file data over the Altera download cable. Your program accesses the host based file system using the ANSI C standard library I/O functions, such as `fopen()` and `fread()`. The host-based file system is a software package which you add to your BSP.

The following features and restrictions apply to the host based file system:

- The host-based file system makes the Nios II C/C++ application project directory and its subdirectories available to the hardware abstraction layer (HAL) file system on the target hardware.
- The target processor can access any file in the project directory. Be careful not to corrupt project source files.

- The host-based file system only operates while debugging a project. It cannot be used for run sessions.
- Host file data travels between host and target serially through the Altera download cable, and therefore file access time is relatively slow. Depending on your host and target system configurations, it can take several milliseconds per call to the host. For higher performance, use buffered I/O function such as `fread()` and `fwrite()`, and increase the buffer size for large files.

You configure the host-based file system using the Nios II BSP Editor. The host-based file system has one setting: the mount point, which specifies the mount point within the HAL file system. For example, if you name the mount point `/mnt/host` and the project directory on you host computer is `/software/project1`, in a HAL-based program, the following code opens the file `/software/project1/datafile.dat`:

```
fopen( "/mnt/host/datafile.dat", "r" );
```

Using Timer Devices

Timer devices are hardware peripherals that count clock ticks and can generate periodic interrupt requests. You can use a timer device to provide a number of time-related facilities, such as the HAL system clock, alarms, the time-of-day, and time measurement. To use the timer facilities, the Nios II processor system must include a timer peripheral in hardware.

The HAL API provides two types of timer device drivers:

- System clock driver—Supports alarms, such as you would use in a scheduler.
- Timestamp driver—Supports high-resolution time measurement.

An individual timer peripheral can behave as either a system clock or a timestamp, but not both.



The HAL-specific API functions for accessing timer devices are defined in `sys/alt_alarm.h` and `sys/alt_timestamp.h`.

System Clock Driver

The HAL system clock driver provides a periodic heartbeat, causing the system clock to increment on each beat. Software can use the system clock facilities to execute functions at specified times, and to obtain timing information. You select a specific hardware timer peripheral as the system clock device by manipulating BSP settings.


For details about how to control BSP settings, refer to [“HAL BSP Settings” on page 6-2](#).

The HAL provides implementations of the following standard UNIX functions: `gettimeofday()`, `settimeofday()`, and `times()`. The times returned by these functions are based on the HAL system clock.

The system clock measures time in clock ticks. For embedded engineers who deal with both hardware and software, do not confuse the HAL system clock with the clock signal driving the Nios II processor hardware. The period of a HAL system clock tick is generally much longer than the hardware system clock. `system.h` defines the clock tick frequency.

At runtime, you can obtain the current value of the system clock by calling the `alt_nticks()` function. This function returns the elapsed time in system clock ticks since reset. You can get the system clock rate, in ticks per second, by calling the function `alt_ticks_per_second()`. The HAL timer driver initializes the tick frequency when it creates the instance of the system clock.

The standard UNIX function `gettimeofday()` is available to obtain the current time. You must first calibrate the time of day by calling `settimeofday()`. In addition, you can use the `times()` function to obtain information about the number of elapsed ticks. The prototypes for these functions appear in **times.h**.

 For more information about the use of these functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

Alarms

You can register functions to be executed at a specified time using the HAL alarm facility. A software program registers an alarm by calling the function `alt_alarm_start()`:

```
int alt_alarm_start (alt_alarm* alarm,
                    alt_u32   nticks,
                    alt_u32   (*callback) (void* context),
                    void*     context);
```


The function `callback()` is called after `nticks` have elapsed. The input argument `context` is passed as the input argument to `callback()` when the call occurs. The HAL does not use the `context` parameter. It is only used as a parameter to the `callback()` function.

Your code must allocate the `alt_alarm` structure, pointed to by the input argument `alarm`. This data structure must have a lifetime that is at least as long as that of the alarm. The best way to allocate this structure is to declare it as a static or global. `alt_alarm_start()` initializes `*alarm`.

The callback function can reset the alarm. The return value of the registered callback function is the number of ticks until the next call to `callback`. A return value of zero indicates that the alarm should be stopped. You can manually cancel an alarm by calling `alt_alarm_stop()`.

One alarm is created for each call to `alt_alarm_start()`. Multiple alarms can run simultaneously.

Alarm callback functions execute in an exception context. This imposes functional restrictions which you must observe when writing an alarm callback.

 For more information about the use of these functions, refer to the *Exception Handling* chapter of the *Nios II Software Developer's Handbook*.

The code fragment in **Example 6-8** demonstrates registering an alarm for a periodic callback every second.

Example 6-8. Using a Periodic Alarm Callback Function

```

#include <stddef.h>
#include <stdio.h>
#include "sys/alt_alarm.h"
#include "alt_types.h"

/*
 * The callback function.
 */

alt_u32 my_alarm_callback (void* context)
{
    /* This function is called once per second */
    return alt_ticks_per_second();
}

...

/* The alt_alarm must persist for the duration of the alarm. */
static alt_alarm alarm;

...

if (alt_alarm_start (&alarm,
                    alt_ticks_per_second(),
                    my_alarm_callback,
                    NULL) < 0)
{
    printf ("No system clock available\n");
}

```

Timestamp Driver

Sometimes you want to measure time intervals with a degree of accuracy greater than that provided by HAL system clock ticks. The HAL provides high resolution timing functions using a timestamp driver. A timestamp driver provides a monotonically increasing counter that you can sample to obtain timing information. The HAL only supports one timestamp driver in the system.


You specify a hardware timer peripheral as the timestamp device by manipulating BSP settings. The Altera-provided timestamp driver uses the timer that you specify.

If a timestamp driver is present, the following functions are available:

- `alt_timestamp_start()`
- `alt_timestamp()`

Calling `alt_timestamp_start()` starts the counter running. Subsequent calls to `alt_timestamp()` return the current value of the timestamp counter. Calling `alt_timestamp_start()` again resets the counter to zero. The behavior of the timestamp driver is undefined when the counter reaches $(2^{32} - 1)$.

You can obtain the rate at which the timestamp counter increments by calling the function `alt_timestamp_freq()`. This rate is typically the hardware frequency of the Nios II processor system—usually millions of cycles per second. The timestamp drivers are defined in the `alt_timestamp.h` header file.

 For more information about the use of these functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

The code fragment in [Example 6-9](#) shows how you can use the timestamp facility to measure code execution time.

Example 6-9. Using the Timestamp to Measure Code Execution Time

```
#include <stdio.h>
#include "sys/alt_timestamp.h"
#include "alt_types.h"

int main (void)
{
    alt_u32 time1;
    alt_u32 time2;
    alt_u32 time3;

    if (alt_timestamp_start() < 0)
    {
        printf ("No timestamp device available\n");
    }
    else
    {
        time1 = alt_timestamp();
        func1(); /* first function to monitor */
        time2 = alt_timestamp();
        func2(); /* second function to monitor */
        time3 = alt_timestamp();

        printf ("time in func1 = %u ticks\n",
            (unsigned int) (time2 - time1));
        printf ("time in func2 = %u ticks\n",
            (unsigned int) (time3 - time2));
        printf ("Number of ticks per second = %u\n",
            (unsigned int)alt_timestamp_freq());
    }
    return 0;
}
```

Using Flash Devices


The HAL provides a generic device model for nonvolatile flash memory devices. Flash memories use special programming protocols to store data. The HAL API provides functions to write data to flash memory. For example, you can use these functions to implement a flash-based file subsystem.

The HAL API also provides functions to read flash, although it is generally not necessary. For most flash devices, programs can treat the flash memory space as simple memory when reading, and do not need to call special HAL API functions. If the flash device has a special protocol for reading data, such as the Altera erasable programmable configurable serial (EPCS) configuration device, you must use the HAL API to both read and write data.

This section describes the HAL API for the flash device model. The following two APIs provide two different levels of access to the flash:

- Simple flash access—Functions that write buffers to flash and read them back at the block level. In writing, if the buffer is less than a full block, these functions erase preexisting flash data above and below the newly written data.
- Fine-grained flash access—Functions that write buffers to flash and read them back at the buffer level. In writing, if the buffer is less than a full block, these functions preserve preexisting flash data above and below the newly written data. This functionality is generally required for managing a file subsystem.

The API functions for accessing flash devices are defined in `sys/alt_flash.h`.

 For more information about the use of these functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*. You can get details about the Common Flash Interface, including the organization of common flash interface (CFI) erase regions and blocks, from JEDEC (www.jedec.org). You can find the CFI standard by searching for document JESD68.

Simple Flash Access

This interface consists of the functions `alt_flash_open_dev()`, `alt_write_flash()`, `alt_read_flash()`, and `alt_flash_close_dev()`. The code “Using the Simple Flash API Functions” on page 6-22 shows the use of all of these functions in one code example. You open a flash device by calling `alt_flash_open_dev()`, which returns a file handle to a flash device. This function takes a single argument that is the name of the flash device, as defined in `system.h`.

After you obtain a handle, you can use the `alt_write_flash()` function to write data to the flash device. The prototype is:

```
int alt_write_flash( alt_flash_fd* fd,
                   int           offset,
                   const void*   src_addr,
                   int           length )
```

A call to this function writes to the flash device identified by the handle `fd`. The driver writes the data starting at `offset` bytes from the base of the flash device. The data written comes from the address pointed to by `src_addr`, and the amount of data written is `length`.

There is also an `alt_read_flash()` function to read data from the flash device. The prototype is:

```
int alt_read_flash( alt_flash_fd* fd,
                  int           offset,
                  void*         dest_addr,
                  int           length )
```

A call to `alt_read_flash()` reads from the flash device with the handle `fd`, `offset` bytes from the beginning of the flash device. The function writes the data to location pointed to by `dest_addr`, and the amount of data read is `length`. For most flash devices, you can access the contents as standard memory, making it unnecessary to use `alt_read_flash()`.

The function `alt_flash_close_dev()` takes a file handle and closes the device. The prototype for this function is:

```
void alt_flash_close_dev(alt_flash_fd* fd )
```

The code in [Example 6-10](#) shows the use of simple flash API functions to access a flash device named `/dev/ext_flash`, as defined in `system.h`.

Block Erasure or Corruption

Generally, flash memory is divided into blocks. `alt_write_flash()` might need to erase the contents of a block before it can write data to it. In this case, it makes no attempt to preserve the existing contents of the block. This action can lead to unexpected data corruption (erasure), if you are performing writes that do not fall on block boundaries. If you wish to preserve existing flash memory contents, use the fine-grained flash functions. These are discussed in the following section.

[Table 6-7 on page 6-23](#) shows how you can cause unexpected data corruption by writing using the simple flash access functions. [Table 6-7](#) shows the example of an 8 Kbyte flash memory comprising two 4 Kbyte blocks. First write 5 Kbytes of all `0xAA` to flash memory at address `0x0000`, and then write 2 Kbytes of all `0xBB` to address `0x1400`. After the first write succeeds (at time $t(2)$), the flash memory contains 5 Kbyte of `0xAA`, and the rest is empty (that is, `0xFF`). Then the second write begins, but before writing to the second block, the block is erased. At this point, $t(3)$, the flash contains 4 Kbyte of `0xAA` and 4 Kbyte of `0xFF`. After the second write finishes, at time $t(4)$, the 2 Kbyte of `0xFF` at address `0x1000` is corrupted.

Fine-Grained Flash Access

Three additional functions provide complete control for writing flash contents at the highest granularity:

- `alt_get_flash_info()`
- `alt_erase_flash_block()`
- `alt_write_flash_block()`

By the nature of flash memory, you cannot erase a single address in a block. You must erase (that is, set to all ones) an entire block at a time. Writing to flash memory can only change bits from 1 to 0; to change any bit from 0 to 1, you must erase the entire block along with it.

Therefore, to alter a specific location in a block while leaving the surrounding contents unchanged, you must read out the entire contents of the block to a buffer, alter the value(s) in the buffer, erase the flash block, and finally write the whole block-sized buffer back to flash memory. The fine-grained flash access functions automate this process at the flash block level.

Example 6-10. Using the Simple Flash API Functions

```

#include <stdio.h>
#include <string.h>
#include "sys/alt_flash.h"
#define BUF_SIZE 1024

int main ()
{
    alt_flash_fd* fd;
    int          ret_code;
    char         source[BUF_SIZE];
    char         dest[BUF_SIZE];

    /* Initialize the source buffer to all 0xAA */
    memset(source, 0xAA, BUF_SIZE);

    fd = alt_flash_open_dev("/dev/ext_flash");
    if (fd!=NULL)
    {
        ret_code = alt_write_flash(fd, 0, source, BUF_SIZE);
        if (ret_code==0)
        {
            ret_code = alt_read_flash(fd, 0, dest, BUF_SIZE);
            if (ret_code==0)
            {
                /*
                 * Success.
                 * At this point, the flash is all 0xAA and we
                 * should have read that all back to dest
                 */
            }
        }
        alt_flash_close_dev(fd);
    }
    else
    {
        printf("Can't open flash device\n");
    }
    return 0;
}

```

`alt_get_flash_info()` gets the number of erase regions, the number of erase blocks in each region, and the size of each erase block. The function prototype is as follows:

```

int alt_get_flash_info (
    alt_flash_fd* fd,
    flash_region** info,
    int*          number_of_regions )

```

If the call is successful, on return the address pointed to by `number_of_regions` contains the number of erase regions in the flash memory, and `*info` points to an array of `flash_region` structures. This array is part of the file descriptor.

Table 6-7. Example of Writing Flash and Causing Unexpected Data Corruption

Address	Block	Time t(0)	Time t(1)	Time t(2)	Time t(3)	Time t(4)
		Before First Write	First Write		Second Write	
			After Erasing Block(s)	After Writing Data 1	After Erasing Block(s)	After Writing Data 2
0x0000	1	??	FF	AA	AA	AA
0x0400	1	??	FF	AA	AA	AA
0x0800	1	??	FF	AA	AA	AA
0x0C00	1	??	FF	AA	AA	AA
0x1000	2	??	FF	AA	FF	FF (1)
0x1400	2	??	FF	FF	FF	BB
0x1800	2	??	FF	FF	FF	BB
0x1C00	2	??	FF	FF	FF	FF

Note to Table 6-7:

(1) Unintentionally cleared to FF during erasure for second write.

The `flash_region` structure is defined in `sys/alt_flash_types.h`. The data structure is defined as follows:

```
typedef struct flash_region
{
    int offset;          /* Offset of this region from start of the flash */
    int region_size;    /* Size of this erase region */
    int number_of_blocks; /* Number of blocks in this region */
    int block_size;     /* Size of each block in this erase region */
}flash_region;
```

With the information obtained by calling `alt_get_flash_info()`, you are in a position to erase or program individual blocks of the flash device.

`alt_erase_flash()` erases a single block in the flash memory. The function prototype is as follows:


```
int alt_erase_flash_block ( alt_flash_fd* fd, int offset, int length )
```

The flash memory is identified by the handle `fd`. The block is identified as being `offset` bytes from the beginning of the flash memory, and the block size is passed in `length`.

`alt_write_flash_block()` writes to a single block in the flash memory. The prototype is:

```
int alt_write_flash_block( alt_flash_fd* fd,
                          int           block_offset,
                          int           data_offset,
                          const void   *data,
                          int           length)
```

This function writes to the flash memory identified by the handle `fd`. It writes to the block located `block_offset` bytes from the start of the flash device. The function writes `length` bytes of data from the location pointed to by `data` to the location `data_offset` bytes from the start of the flash device.

 These program and erase functions do not perform address checking, and do not verify whether a write operation spans into the next block. You must pass in valid information about the blocks to program or erase.

The code in [Example 6-11](#) on [page 6-24](#) demonstrates the use of the fine-grained flash access functions.

Example 6-11. Using the Fine-Grained Flash Access API Functions

```
#include <string.h>
#include "sys/alt_flash.h"
#include "stdtypes.h"
#include "system.h"#define BUF_SIZE 100

int main (void)
{
    flash_region* regions;
    alt_flash_fd* fd;
    int          number_of_regions;
    int          ret_code;
    char        write_data[BUF_SIZE];

    /* Set write_data to all 0xa */
    memset(write_data, 0xA, BUF_SIZE);

    fd = alt_flash_open_dev(EXT_FLASH_NAME);

    if (fd)
    {
        ret_code = alt_get_flash_info(fd, &regions, &number_of_regions);

        if (number_of_regions && (regions->offset == 0))
        {
            /* Erase the first block */
            ret_code = alt_erase_flash_block(fd,
                                           regions->offset,
                                           regions->block_size);

            if (ret_code == 0)
            {
                /*
                 * Write BUF_SIZE bytes from write_data 100 bytes to
                 * the first block of the flash
                 */
                ret_code = alt_write_flash_block (
                    fd,
                    regions->offset,
                    regions->offset+0x100,
                    write_data,
                    BUF_SIZE );
            }
        }
    }
    return 0;
}
```


Using DMA Devices

The HAL provides a device abstraction model for direct memory access (DMA) devices. These are peripherals that perform bulk data transactions from a data source to a destination. Sources and destinations can be memory or another device, such as an Ethernet connection.

In the HAL DMA device model, there are two categories of DMA transactions: transmit and receive. The HAL provides two device drivers to implement transmit channels and receive channels. A transmit channel takes data in a source buffer and transmits it to a destination device. A receive channel receives data from a device and deposits it in a destination buffer. Depending on the implementation of the underlying hardware, software might have access to only one of these two endpoints.

Figure 6-2 shows the three basic types of DMA transactions. Copying data from memory to memory involves both receive and transmit DMA channels simultaneously.

Figure 6-2. Three Basic Types of DMA Transactions



The API for access to DMA devices is defined in `sys/alt_dma.h`.

For more information about the use of these functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

DMA devices operate on the contents of physical memory, therefore when reading and writing data you must consider cache interactions.

For more information about cache memory, refer to the *Cache and Tightly-Coupled Memory* chapter of the *Nios II Software Developer's Handbook*.

DMA Transmit Channels

DMA transmit requests are queued using a DMA transmit device handle. To obtain a handle, use the function `alt_dma_txchan_open()`. This function takes a single argument, the name of a device to use, as defined in `system.h`.

The code in [Example 6-12](#) shows how to obtain a handle for a DMA transmit device `dma_0`.

Example 6-12. Obtaining a File Handle for a DMA Device

```
#include <stddef.h>
#include "sys/alt_dma.h"

int main (void)
{
    alt_dma_txchan tx;

    tx = alt_dma_txchan_open ("/dev/dma_0");
    if (tx == NULL)
    {
        /* Error */
    }
    else
    {
        /* Success */
    }
    return 0;
}
```


You can use this handle to post a transmit request using `alt_dma_txchan_send()`. The prototype is:

```
typedef void (alt_txchan_done)(void* handle);

int alt_dma_txchan_send (alt_dma_txchan dma,
                        const void* from,
                        alt_u32 length,
                        alt_txchan_done* done,
                        void* handle);
```

Calling `alt_dma_txchan_send()` posts a transmit request to channel `dma`. Argument `length` specifies the number of bytes of data to transmit, and argument `from` specifies the source address. The function returns before the full DMA transaction completes. The return value indicates whether the request is successfully queued. A negative return value indicates that the request failed. When the transaction completes, the user-supplied function `done` is called with argument `handle` to provide notification.

Two additional functions are provided for manipulating DMA transmit channels: `alt_dma_txchan_space()`, and `alt_dma_txchan_ioctl()`. The `alt_dma_txchan_space()` function returns the number of additional transmit requests that can be queued to the device. The `alt_dma_txchan_ioctl()` function performs device-specific manipulation of the transmit device.

 If you are using the Avalon Memory-Mapped® (Avalon-MM) DMA device to transmit to hardware (not memory-to-memory transfer), call the `alt_dma_txchan_ioctl()` function with the request argument set to `ALT_DMA_TX_ONLY_ON`.

 For further information, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

DMA Receive Channels

DMA receive channels operate similarly to DMA transmit channels. Software can obtain a handle for a DMA receive channel using the `alt_dma_rxchan_open()` function. You can then use the `alt_dma_rxchan_prepare()` function to post receive requests. The prototype for `alt_dma_rxchan_prepare()` is:


```
typedef void (alt_rxchan_done)(void* handle, void* data);

int alt_dma_rxchan_prepare (alt_dma_rxchan  dma,
                           void*          data,
                           alt_u32        length,
                           alt_rxchan_done* done,
                           void*          handle);
```

A call to this function posts a receive request to channel `dma`, for up to `length` bytes of data to be placed at address `data`. This function returns before the DMA transaction completes. The return value indicates whether the request is successfully queued. A negative return value indicates that the request failed. When the transaction completes, the user-supplied function `done()` is called with argument `handle` to provide notification and a pointer to the receive data.

Certain errors can prevent the DMA transfer from completing. Typically this is caused by a catastrophic hardware failure; for example, if a component involved in the transfer fails to respond to a read or write request. If the DMA transfer does not complete (that is, less than `length` bytes are transferred), function `done()` is never called.

Two additional functions are provided for manipulating DMA receive channels: `alt_dma_rxchan_depth()` and `alt_dma_rxchan_ioctl()`.

 If you are using the Avalon-MM DMA device to receive from hardware (not memory-to-memory transfer), call the `alt_dma_rxchan_ioctl()` function with the request argument set to `ALT_DMA_RX_ONLY_ON`.

`alt_dma_rxchan_depth()` returns the maximum number of receive requests that can be queued to the device. `alt_dma_rxchan_ioctl()` performs device-specific manipulation of the receive device.

 For further details, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

The code in [Example 6-13](#) shows a complete example application that posts a DMA receive request, and blocks in `main()` until the transaction completes.

Example 6-13. A DMA Transaction on a Receive Channel

```

#include <stdio.h>
#include <stddef.h>
#include <stdlib.h>
#include "sys/alt_dma.h"
#include "alt_types.h"

/* flag used to indicate the transaction is complete */
volatile int dma_complete = 0;

/* function that is called when the transaction completes */
void dma_done (void* handle, void* data)
{
    dma_complete = 1;
}

int main (void)
{
    alt_u8 buffer[1024];
    alt_dma_rxchan rx;

    /* Obtain a handle for the device */
    if ((rx = alt_dma_rxchan_open ("/dev/dma_0")) == NULL)
    {
        printf ("Error: failed to open device\n");
        exit (1);
    }
    else
    {
        /* Post the receive request */
        if (alt_dma_rxchan_prepare (rx, buffer, 1024, dma_done, NULL) < 0)
        {
            printf ("Error: failed to post receive request\n");
            exit (1);
        }

        /* Wait for the transaction to complete */
        while (!dma_complete);
        printf ("Transaction complete\n");
        alt_dma_rxchan_close (rx);
    }
    return 0;
}

```

Memory-to-Memory DMA Transactions

Copying data from one memory buffer to another buffer involves both receive and transmit DMA drivers. The code in [Example 6-14](#) shows the process of queuing up a receive request followed by a transmit request to achieve a memory-to-memory DMA transaction.

Example 6-14. Copying Data from Memory to Memory (Part 1 of 2)

```
#include <stdio.h>
#include <stdlib.h>

#include "sys/alt_dma.h"
#include "system.h"

static volatile int rx_done = 0;

/*
 * Callback function that obtains notification that the data
 * is received.
 */

static void done (void* handle, void* data)
{
    rx_done++;
}

/*
 *
 */

int main (int argc, char* argv[], char* envp[])
{
    int rc;

    alt_dma_txchan txchan;
    alt_dma_rxchan rxchan;

    void* tx_data = (void*) 0x901000; /* pointer to data to send */
    void* rx_buffer = (void*) 0x902000; /* pointer to rx buffer */

    /* Create the transmit channel */

    if ((txchan = alt_dma_txchan_open("/dev/dma_0")) == NULL)
    {
        printf ("Failed to open transmit channel\n");
        exit (1);
    }

    /* Create the receive channel */

    if ((rxchan = alt_dma_rxchan_open("/dev/dma_0")) == NULL)
    {
        printf ("Failed to open receive channel\n");
        exit (1);
    }

    /* Post the transmit request */

    if ((rc = alt_dma_txchan_send (txchan,
                                   tx_data,
                                   128,
                                   NULL,
                                   NULL)) < 0)
    {
        printf ("Failed to post transmit request, reason = %i\n", rc);
        exit (1);
    }

    /* Continued... */
}
```

Example 6-14. Copying Data from Memory to Memory (Part 2 of 2)

```
/* Post the receive request */

if ((rc = alt_dma_rxchan_prepare (rxchan,
                                rx_buffer,
                                128,
                                done,
                                NULL)) < 0)
{
    printf ("Failed to post read request, reason = %i\n", rc);
    exit (1);
}

/* wait for transfer to complete */

while (!rx_done);

printf ("Transfer successful!\n");

return 0;
}
```

Using Interrupt Controllers

The HAL supports two types of interrupt controllers:

- The Nios II internal interrupt controller
- An external interrupt controller component



For information about working with interrupt controllers, refer to the *Exception Handling* chapter of the *Nios II Software Developer's Handbook*.

Reducing Code Footprint

Code size is always a concern for embedded systems developers, because there is a cost associated with the memory device that stores code. The ability to control and reduce code size is important in controlling this cost.

The HAL environment is designed to include only those features that you request, minimizing the total code footprint. If your Nios II hardware system contains exactly the peripherals used by your program, the HAL contains only the drivers necessary to control the hardware.


The following sections describe options to consider when you need to further reduce code size. The **hello_world_small** example project demonstrates the use of some of these options to reduce code size to the absolute minimum.

Implementing the options in the following sections entails making changes to BSP settings. For detailed information about manipulating BSP settings, refer to “[HAL BSP Settings](#)” on page 6-2.

Enable Compiler Optimizations

To enable compiler optimizations, use the `-O3` compiler optimization level for the `nios2-elf-gcc` compiler. You can specify this command-line option through a BSP setting.

With this option turned on, the Nios II compiler compiles code with the maximum optimization available, for both size and speed.

 You must set this option for both the BSP and the application project.

Use Reduced Device Drivers

Some devices provide two driver variants, a fast variant and a small variant. The feature sets provided by these two variants are device specific. The fast variant is full-featured, and the small variant provides a reduced code footprint.

By default the HAL always uses the fast driver variants. You can select the reduced device driver for all hardware components, or for an individual component, through HAL BSP settings.

Table 6-8 lists the Altera Nios II peripherals that currently provide small footprint drivers. The small footprint option might also affect other peripherals. Refer to each peripheral's data sheet for complete details of its driver's small footprint behavior.

Table 6-8. Altera Peripherals Offering Small Footprint Drivers

Peripheral	Small Footprint Behavior
UART	Polled operation, rather than IRQ-driven
JTAG UART	Polled operation, rather than IRQ-driven
Common flash interface controller	Driver excluded in small footprint mode
LCD module controller	Driver excluded in small footprint mode
EPCS serial configuration device	Driver excluded in small footprint mode

Reduce the File Descriptor Pool

The file descriptors that access character mode devices and files are allocated from a file descriptor pool. You can change the size of the file descriptor pool through a BSP setting. The default is 32.

Use `/dev/null`

At boot time, standard input, standard output, and standard error are all directed towards the null device, that is, `/dev/null`. This direction ensures that calls to `printf()` during driver initialization do nothing and therefore are harmless. After all drivers are installed, these streams are redirected to the channels configured in the HAL. The footprint of the code that performs this redirection is small, but you can eliminate it entirely by selecting `null` for `stdin`, `stdout`, and `stderr`. This selection assumes that you want to discard all data transmitted on standard out or standard error, and your program never receives input through `stdin`. You can control the assignment of `stdin`, `stdout`, and `stderr` channels by manipulating BSP settings.

Use a Smaller File I/O Library

Use the Small newlib C Library

The full newlib ANSI C standard library is often unnecessary for embedded systems. The GNU Compiler Collection (GCC) provides a reduced implementation of the newlib ANSI C standard library, omitting features of newlib that are often superfluous for embedded systems. The small newlib implementation requires a smaller code footprint. When you use `nios2-elf-gcc` at the command line, the `-msmallc` command-line option enables the small C library.

You can select the small newlib library through BSP settings. [Table 6-9](#) summarizes the limitations of the Nios II small newlib C library implementation.

Table 6-9. Limitations of the Nios II Small newlib C Library (Part 1 of 2)


Limitation	Functions Affected
No floating-point support for <code>printf()</code> family of routines. The functions listed are implemented, but <code>%f</code> and <code>%g</code> options are not supported. (1)	<code>asprintf()</code> <code>fprintf()</code> <code>fprintf()</code> <code>iprintf()</code> <code>printf()</code> <code>siprintf()</code> <code>snprintf()</code> <code>sprintf()</code>
No floating-point support for <code>vprintf()</code> family of routines. The functions listed are implemented, but <code>%f</code> and <code>%g</code> options are not supported.	<code>vasprintf()</code> <code>vfprintf()</code> <code>vfprintf()</code> <code>vprintf()</code> <code>vsprintf()</code> <code>vsnprintf()</code> <code>vsprintf()</code>
No support for <code>scanf()</code> family of routines. The functions listed are not supported.	<code>fscanf()</code> <code>scanf()</code> <code>sscanf()</code> <code>vfscanf()</code> <code>vscanf()</code> <code>vsscanf()</code>
No support for seeking. The functions listed are not supported.	<code>fseek()</code> <code>ftell()</code>
No support for opening/closing <code>FILE *</code> . Only pre-opened <code>stdout</code> , <code>stderr</code> , and <code>stdin</code> are available. The functions listed are not supported.	<code>fopen()</code> <code>fclose()</code> <code>fdopen()</code> <code>fcloseall()</code> <code>fileno()</code>


Table 6-9. Limitations of the Nios II Small newlib C Library (Part 2 of 2)


Limitation	Functions Affected
No buffering of <code>stdio.h</code> output routines.	functions supported with no buffering: <code>fiprintf()</code> <code>fputc()</code> <code>fputs()</code> <code>perror()</code> <code>putc()</code> <code>putchar()</code> <code>puts()</code> <code>printf()</code> functions not supported: <code>setbuf()</code> <code>setvbuf()</code>
No <code>stdio.h</code> input routines. The functions listed are not supported.	<code>fgetc()</code> <code>gets()</code> <code>fscanf()</code> <code>getc()</code> <code>getchar()</code> <code>gets()</code> <code>getw()</code> <code>scanf()</code>
No support for locale.	<code>setlocale()</code> <code>localeconv()</code>
No support for C++, because the functions listed in this table are not supported.	

Note to Table 6-9:

(1) These functions are a Nios II extension. GCC does not implement them in the small newlib C library.

 The small newlib C library does not support MicroC/OS-II.

 For details about the GCC small newlib C library, refer to the newlib documentation installed with the Nios II EDS. On the Windows **Start** menu, click **Programs > Altera > Nios II > Nios II Documentation**.

 The Nios II implementation of the small newlib C library differs slightly from GCC. [Table 6-9](#) provides details about the differences.

Use UNIX-Style File I/O

If you need to reduce the code footprint further, you can omit the newlib C library, and use the UNIX-style API. For details, refer to [“UNIX-Style Interface” on page 6-5](#).

The Nios II EDS provides ANSI C file I/O, in the newlib C library, because there is a per-access performance overhead associated with accessing devices and files using the UNIX-style file I/O functions. The ANSI C file I/O provides buffered access, thereby reducing the total number of hardware I/O accesses performed. Also the ANSI C API is more flexible and therefore easier to use. However, these benefits are gained at the expense of code footprint.

Emulate ANSI C Functions

If you choose to omit the full implementation of newlib, but you need a limited number of ANSI-style functions, you can implement them easily using UNIX-style functions. The code in [Example 6-15](#) shows a simple, unbuffered implementation of `getchar()`.

Example 6-15. Unbuffered `getchar()`

```
/* getchar: unbuffered single character input */
int getchar ( void )
{
    char c;
    return ( read ( 0, &c, 1 ) == 1 ) ? ( unsigned char ) c : EOF;
}
```



This example is from *The C Programming Language, Second Edition*, by Brian W. Kernighan and Dennis M. Ritchie. This standard textbook contains many other useful functions.

Use the Lightweight Device Driver API

The lightweight device driver API allows you to minimize the overhead of accessing device drivers. It has no direct effect on the size of the drivers themselves, but lets you eliminate driver API features which you might not need, reducing the overall size of the HAL code.

The lightweight device driver API is available for character-mode devices. The following device drivers support the lightweight device driver API:

- JTAG UART
- UART
- Optrex 16207 LCD

For these devices, the lightweight device driver API conserves code space by eliminating the dynamic file descriptor table and replacing it with three static file descriptors, corresponding to `stdin`, `stdout`, and `stderr`. Library functions related to opening, closing, and manipulating file descriptors are unavailable, but all other library functionality is available. You can refer to `stdin`, `stdout`, and `stderr` as you would to any other file descriptor. You can also refer to the following predefined file numbers:

```
#define STDIN 0
#define STDOUT 1
#define STDERR 2
```

This option is appropriate if your program has a limited need for file I/O. The Altera host-based file system and the Altera read-only zip file system are not available with the reduced device driver API. You can select the reduced device drivers through BSP settings.

By default, the lightweight device driver API is disabled.



For further details about the lightweight device driver API, refer to the *Developing Device Drivers for the Hardware Abstraction Layer* chapter of the *Nios II Software Developer's Handbook*.

Use the Minimal Character-Mode API

If you can limit your use of character-mode I/O to very simple features, you can reduce code footprint by using the minimal character-mode API. This API includes the following functions:

- `alt_printf()`
- `alt_putchar()`
- `alt_putstr()`
- `alt_getchar()`

These functions are appropriate if your program only needs to accept command strings and send simple text messages. Some of them are helpful only in conjunction with the lightweight device driver API, discussed in “Use the Lightweight Device Driver API” on page 6-34.

To use the minimal character-mode API, include the header file `sys/alt_stdio.h`.

The following sections outline the effects of the functions on code footprint.

alt_printf()

This function is similar to `printf()`, but supports only the `%c`, `%s`, `%x`, and `%%` substitution strings. `alt_printf()` takes up substantially less code space than `printf()`, regardless whether you select the lightweight device driver API. `alt_printf()` occupies less than 1Kbyte with compiler optimization level `-O2`.

alt_putchar()


Equivalent to `putchar()`. In conjunction with the lightweight device driver API, this function further reduces code footprint. In the absence of the lightweight API, it calls `putchar()`.

alt_putstr()

Similar to `puts()`, except that it does not append a newline character to the string. In conjunction with the lightweight device driver API, this function further reduces code footprint. In the absence of the lightweight API, it calls `puts()`.

alt_getchar()

Equivalent to `getchar()`. In conjunction with the lightweight device driver API, this function further reduces code footprint. In the absence of the lightweight API, it calls `getchar()`.

 For further details about the minimal character-mode functions, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.


Eliminate Unused Device Drivers

If a hardware device is present in the system, by default the Nios II development flows assume the device needs drivers, and configure the HAL BSP accordingly. If the HAL can find an appropriate driver, it creates an instance of this driver. If your program never actually accesses the device, resources are being used unnecessarily to initialize the device driver.

If the hardware includes a device that your program never uses, consider removing the device from the hardware. This reduces both code footprint and FPGA resource usage.

However, there are cases when a device must be present, but runtime software does not require a driver. The most common example is flash memory. The user program might boot from flash, but not use it at runtime; thus, it does not need a flash driver.

You can selectively omit any individual driver, select a specific driver version, or substitute your own driver.

 For further information about controlling driver configurations, refer to the *Nios II Software Build Tools* chapter of the *Nios II Software Developer's Handbook*.

Another way to control the device driver initialization process is to use the free-standing environment. For details, refer to “[Boot Sequence and Entry Point](#)” on [page 6-37](#).

Eliminate Unneeded Exit Code


The HAL calls the `exit()` function at system shutdown to provide a clean exit from the program. `exit()` flushes all of the C library internal I/O buffers and calls any C++ functions registered with `atexit()`. In particular, `exit()` is called on return from `main()`. Two HAL options allow you to minimize or eliminate this exit code.

Eliminate Clean Exit

To avoid the overhead associated with providing a clean exit, your program can use the function `_exit()` in place of `exit()`. This function does not require you to change source code. You can select the `_exit()` function through a BSP setting.

Eliminate All Exit Code

Many embedded systems never exit at all. In such cases, exit code is unnecessary. You can eliminate all exit code through a BSP setting.

 If you enable this option, ensure that your `main()` function (or `alt_main()` function) does not return.

Turn off C++ Support

By default, the HAL provides support for C++ programs, including default constructors and destructors. You can disable C++ support through a BSP setting.

Boot Sequence and Entry Point

Normally, your program's entry point is the function `main()`. There is an alternate entry point, `alt_main()`, that you can use to gain greater control of the boot sequence. The difference between entering at `main()` and entering at `alt_main()` is the difference between hosted and free-standing applications.

Hosted versus Free-Standing Applications

The ANSI C standard defines a hosted application as one that calls `main()` to begin execution. At the start of `main()`, a hosted application presumes the runtime environment and all system services are initialized and ready to use. This is true in the HAL environment. If you are new to Nios II programming, the HAL's hosted environment helps you come up to speed more easily, because you need not consider what devices exist in the system or how to initialize each one. The HAL initializes the whole system.

The ANSI C standard also provides for an alternate entry point that avoids automatic initialization, and assumes that the Nios II programmer initializes any needed hardware explicitly. The `alt_main()` function provides a free-standing environment, giving you complete control over the initialization of the system. The free-standing environment places on the programmer the responsibility to initialize any system features used in the program. For example, calls to `printf()` do not function correctly in the free-standing environment, unless `alt_main()` first instantiates a character-mode device driver, and redirects `stdout` to the device.



Using the free-standing environment increases the complexity of writing Nios II programs, because you assume responsibility for initializing the system. If your main interest is to reduce code footprint, you should use the suggestions described in ["Reducing Code Footprint" on page 6-30](#). It is easier to reduce the HAL BSP footprint by using BSP settings, than to use the free-standing mode.

The Nios II EDS provides examples of both free-standing and hosted programs.

Boot Sequence for HAL-Based Programs

The HAL provides system initialization code in the C runtime library (`cr0.S`). This code performs the following boot sequence:

- Flushes the instruction and data cache.
- Configures the stack pointer.
- Configures the global pointer register.
- Initializes the BSS region to zeroes using the linker-supplied symbols `__bss_start` and `__bss_end`. These are pointers to the beginning and the end of the BSS region.
- If there is no boot loader present in the system, copies to RAM any linker section whose run address is in RAM, such as `.rwdata`, `.rodata`, and `.exceptions`. Refer to ["Global Pointer Register" on page 6-42](#).
- Calls `alt_main()`.

The HAL provides a default implementation of the `alt_main()` function, which performs the following steps:

- Calls the `alt_irq_init()` function, located in **alt_sys_init.c**. `alt_irq_init()` **initializes the hardware interrupt controller**. The Nios II development flow creates the file **alt_sys_init.c** for each HAL BSP.
- Calls `ALT_OS_INIT()` to perform any necessary operating system specific initialization. For a system that does not include an operating system (OS) scheduler, this macro has no effect.
- If you are using the HAL with an operating system, initializes the `alt_fd_list_lock` semaphore, which controls access to the HAL file systems.
- Enables interrupts.
- Calls the `alt_sys_init()` function, also located in **alt_sys_init.c**. `alt_sys_init()` initializes all device drivers and software packages in the system.
- Redirects the C standard I/O channels (`stdin`, `stdout`, and `stderr`) to use the appropriate devices.
- Calls the C++ constructors, using the `_do_ctors()` function.
- Registers the C++ destructors to be called at system shutdown.
- Calls `main()`.
- Calls `exit()`, passing the return code of `main()` as the input argument for `exit()`.

alt_main.c, installed with the Nios II EDS, provides this default implementation. The Software Build Tools copy **alt_main.c** to your BSP directory.

Customizing the Boot Sequence

You can provide your own implementation of the start-up sequence by simply defining `alt_main()` in your Nios II project. This gives you complete control of the boot sequence, and allows you to selectively enable HAL services. If your application requires an `alt_main()` entry point, you can copy the default implementation as a starting point and customize it to your needs.


Function `alt_main()` calls function `main()`. After `main()` returns, the default `alt_main()` enters an infinite loop. Alternatively, your custom `alt_main()` might terminate by calling `exit()`. Do not use a `return` statement.

The prototype for `alt_main()` is:

```
void alt_main (void)
```

The HAL build environment includes mechanisms to override default HAL BSP code. This lets you override boot loaders, as well as default device drivers and other system code, with your own implementation.

alt_sys_init.c is a generated file, which you should not modify. However, the Nios II Software Build Tools enable you to control the generated contents of **alt_sys_init.c**. To specify the initialization sequence in **alt_sys_init.c**, you manipulate the `auto_initialize` and `alt_sys_init_priority` properties of each driver, using the `set_sw_property` Tcl command.

 For more information about generated files and how to control the contents of `alt_sys_init.c`, refer to the *Nios II Software Build Tools* chapter of the *Nios II Software Developer's Handbook*. For general information about `alt_sys_init.c`, refer to the *Developing Device Drivers for the Hardware Abstraction Layer* chapter of the *Nios II Software Developer's Handbook*. For details about the `set_sw_property` Tcl command, refer to the *Nios II Software Build Tools Reference* chapter of the *Nios II Software Developer's Handbook*.

Memory Usage

This section describes how the HAL uses memory and arranges code, data, stack, and other logical memory sections, in physical memory.

Memory Sections

By default, HAL-based systems are linked using a generated linker script that is created by the Nios II Software Build Tools. This linker script controls the mapping of code and data to the available memory sections. The autogenerated linker script creates standard code and data sections (`.text`, `.rodata`, `.rwddata`, and `.bss`), plus a section for each physical memory device in the system. For example, if a memory component named `sdram` is defined in the `system.h` file, there is a memory section named `.sdram`. [Figure 6-3](#) shows the organization of a typical HAL link map.

The memory devices that contain the Nios II processor's reset and exception addresses are a special case. The Nios II tools construct the 32-byte `.entry` section starting at the reset address. This section is reserved exclusively for the use of the reset handler. Similarly, the tools construct a `.exceptions` section, starting at the exception address.

In a memory device containing the reset or exception address, the linker creates a normal (nonreserved) memory section above the `.entry` or `.exceptions` section. If there is a region of memory below the `.entry` or `.exceptions` section, it is unavailable to the Nios II software. [Figure 6-3](#) illustrates an unavailable memory region below the `.exceptions` section.

Assigning Code and Data to Memory Partitions

This section describes how to control the placement of program code and data in specific memory sections. In general, the Nios II development flow specifies a sensible default partitioning. However, you might wish to change the partitioning in special situations.

For example, to enhance performance, it is a common technique to place performance-critical code and data in RAM with fast access time. It is also common during the debug phase to reset (that is, boot) the processor from a location in RAM, but then boot from flash memory in the released version of the software. In these cases, you must specify manually which code belongs in which section.

Figure 6-3. Sample HAL Link Map

Physical Memory	HAL Memory Sections
	.entry
ext_flash	.ext_flash
⋮	⋮
sdram	(unused)
	.exceptions
	.text
	.rodata
	.rwdata
	.bss
	.sdram
⋮	⋮
ext_ram	.ext_ram
⋮	⋮
epcs_controller	.epcs_controller

Simple Placement Options

The reset handler code is always placed at the base of the `.reset` partition. The general exception funnel code is always the first code in the section that contains the exception address. By default, the remaining code and data are divided into the following output sections:

- `.text`—All remaining code
- `.rodata`—The read-only data
- `.rwdata`—Read-write data
- `.bss`—Zero-initialized data

You can control the placement of `.text`, `.rodata`, `.rwdata`, and all other memory partitions by manipulating BSP settings. For details about how to control BSP settings, refer to [“HAL BSP Settings” on page 6-2](#).

The Nios II BSP Editor is a very convenient way to manipulate the linker's memory map. The BSP Editor displays memory section and region assignments graphically, allowing you to see overlapping or unused sections of memory. The BSP Editor is available either through the Nios II Software Build Tools for Eclipse, or at the command line of the Nios II Software Build Tools.

 For details, refer to the *Getting Started from the Command Line* chapter of the *Nios II Software Developer's Handbook*.

Advanced Placement Options

In your program source code, you can specify a target memory section for each piece of code. In C or C++, you can use the `section` attribute. This attribute must be placed in a function prototype; you cannot place it in the function declaration itself. The code in [Example 6-16](#) places a variable `foo` in the memory named `ext_ram`, and the function `bar()` in the memory named `sdram`.

Example 6-16. Manually Assigning C Code to a Specific Memory Section


```
/* data should be initialized when using the section attribute */
int foo __attribute__((section(".ext_ram.rwdata"))) = 0;

void bar (void) __attribute__((section(".sdram.txt")));


void bar (void)
{
    foo++;
}
```


In assembly you do this using the `.section` directive. For example, all code after the following line is placed in the memory device named `ext_ram`:

```
.section .ext_ram.txt
```

 The section names `ext_ram` and `sdram` are examples. You need to use section names corresponding to your hardware. When creating section names, use the following extensions:

- `.txt` for code: for example, `.sdram.txt`
- `.rodata` for read-only data: for example, `.cfi_flash.rodata`
- `.rwdata` for read-write data: for example, `.ext_ram.rwdata`

 For details about the use of these features, refer to the GNU compiler and assembler documentation. This documentation is installed with the Nios II EDS. To find it, open the Nios II EDS documentation launchpad, scroll down to **Software Development**, and click **Using the GNU Compiler Collection (GCC)**.

 A powerful way to manipulate the linker memory map is by using the Nios II BSP Editor. With the BSP Editor, you can assign linker sections to specific physical regions, and then review a graphical representation of memory showing unused or overlapping regions. You start the BSP Editor from the Nios II Command Shell. For details about using the BSP Editor, refer to the editor's tool tips.

Placement of the Heap and Stack

By default, the heap and stack are placed in the same memory partition as the `.rwdata` section. The stack grows downwards (toward lower addresses) from the end of the section. The heap grows upwards from the last used memory in the `.rwdata` section. You can control the placement of the heap and stack by manipulating BSP settings.

By default, the HAL performs no stack or heap checking. This makes function calls and memory allocation faster, but it means that `malloc()` (in C) and `new` (in C++) are unable to detect heap exhaustion. You can enable run-time stack checking by manipulating BSP settings. With stack checking on, `malloc()` and `new()` can detect heap exhaustion.

To specify the heap size limit, set the preprocessor symbol `ALT_MAX_HEAP_BYTES` to the maximum heap size in decimal. For example, the preprocessor argument `-DALT_MAX_HEAP_SIZE=1048576` sets the heap size limit to `0x100000`. You can specify this command-line option through a BSP setting. For more information about manipulating BSP settings, refer to “HAL BSP Settings” on page 6-2.

Stack checking has performance costs. If you choose to leave stack checking turned off, you must code your program so as to ensure that it operates within the limits of available heap and stack memory.



Refer to the *Nios II Software Build Tools* chapter of the *Nios II Software Developer's Handbook* for details about selecting stack and heap placement, and setting up stack checking.

For details about how to control BSP settings, refer to “HAL BSP Settings” on page 6-2.


Global Pointer Register

The global pointer register enables fast access to global data structures in Nios II programs. The Nios II compiler implements the global pointer, and determines which data structures to access with it. You do not need to do anything unless you want to change the default compiler behavior.

The global pointer register can access a single contiguous region of 64 Kbytes. To avoid overflowing this region, the compiler only uses the global pointer with small global data structures. A data structure is considered “small” if its size is less than a specified threshold. By default, this threshold is 8 bytes.

The small data structures are allocated to the small global data sections, `.sdata`, `.sdata2`, `.sbss`, and `.sbss2`. The small global data sections are subsections of the `.rwdata` and `.bss` sections. They are located together, as shown in Figure 6-4, to enable the global pointer to access them.

Figure 6-4. Small Global Data sections



If the total size of the small global data structures is more than 64 Kbytes, these data structures overflow the global pointer region. The linker produces an error message saying "Unable to reach *<variable name>* ... from the global pointer ... because the offset ... is out of the allowed range, -32768 to 32767."

You can fix this with the `-G` compiler option. This option sets the threshold size. For example, `-G 4` restricts global pointer usage to data structures 4 bytes long or smaller. Reducing the global pointer threshold reduces the size of the small global data sections.

The `-G` option's numeric argument is in decimal. You can specify this compiler option through a project setting. For information about manipulating project settings, refer to "HAL BSP Settings" on page 6-2.



You must set this option to the same value for both the BSP and the application project.


Boot Modes

The processor's boot memory is the memory that contains the reset vector. This device might be an external flash or an Altera EPCS serial configuration device, or it might be an on-chip RAM. Regardless of the nature of the boot memory, HAL-based systems are constructed so that all program and data sections are initially stored in it. The HAL provides a small boot loader program that copies these sections to their run time locations at boot time. You can specify run time locations for program and data memory by manipulating BSP settings.

If the runtime location of the `.text` section is outside of the boot memory, the Altera flash programmer places a boot loader at the reset address. This boot loader is responsible for loading all program and data sections before the call to `_start`. When booting from an EPCS device, this loader function is provided by the hardware.

However, if the runtime location of the `.text` section is in the boot memory, the system does not need a separate loader. Instead the `_reset` entry point in the HAL executable program is called directly. The function `_reset` initializes the instruction cache and then calls `_start`. This initialization sequence lets you develop applications that boot and execute directly from flash memory.

When running in this mode, the HAL executable program must take responsibility for loading any sections that require loading to RAM. The `.rdata`, `.rodata`, and `.exceptions` sections are loaded before the call to `alt_main()`, as required. This loading is performed by the function `alt_load()`. To load any additional sections, use the `alt_load_section()` function.

 For more information about `alt_load_section()`, refer to the *HAL API Reference* chapter of the *Nios II Software Developer's Handbook*.

Working with HAL Source Files

You might wish to view files in the HAL, especially header files, for reference. This section describes how to find and use HAL source files.


Finding HAL Files

You determine the location of HAL source files when you create the BSP. HAL source files (and other BSP files) are copied to the BSP directory.

 For details, refer to the *Nios II Software Build Tools Reference* chapter of the *Nios II Software Developer's Handbook*.

Overriding HAL Functions

HAL source files are copied to your BSP directory when you create your BSP. If you regenerate a BSP, any HAL source files that differ from the installation files are copied. Avoid modifying BSP files. To override default HAL code, use BSP settings, or custom device drivers or software packages.

 For information about what happens when you regenerate a BSP, refer to "Revising your BSP" in the *Nios II Software Build Tools* chapter of the *Nios II Software Developer's Handbook*.



Avoid modifying HAL source files. If you modify a HAL source file, you cannot regenerate the BSP without losing your changes. This makes it difficult to keep the BSP coordinated with changes to the underlying SOPC Builder system.



For more information, refer to “Nios II Software Projects” in the *Nios II Software Build Tools* chapter of the *Nios II Software Developer’s Handbook*.

Referenced Documents

This chapter references the following documents:

- *Getting Started with the Graphical User Interface* chapter of the *Nios II Software Developer’s Handbook*
- *Getting Started from the Command Line* chapter of the *Nios II Software Developer’s Handbook*
- *Nios II Software Build Tools* chapter of the *Nios II Software Developer’s Handbook*
- *Developing Device Drivers for the Hardware Abstraction Layer* chapter of the *Nios II Software Developer’s Handbook*
- *Exception Handling* chapter of the *Nios II Software Developer’s Handbook*
- *Cache and Tightly-Coupled Memory* chapter of the *Nios II Software Developer’s Handbook*
- *Read-Only Zip File System* chapter of the *Nios II Software Developer’s Handbook*
- *HAL API Reference* chapter of the *Nios II Software Developer’s Handbook*
- *Nios II Software Build Tools Reference* chapter of the *Nios II Software Developer’s Handbook*
- *Off-Chip Interface Peripherals* section in *Volume 5: Embedded Peripherals* of the *Quartus II Handbook*
- *The C Programming Language, Second Edition*, by Brian Kernighan and Dennis M. Ritchie (Prentice-Hall)
- *GNU documentation* on the Nios II EDS documentation launchpad

Document Revision History

Table 6–10 shows the revision history for this document.

Table 6-10. Document Revision History

Date & Document Version	Changes Made	Summary of Changes
November 2009 v9.1.0	<ul style="list-style-type: none"> ■ Described <code>alt_irq_init()</code> function ■ Inserted host-based file system description ■ Removed IDE-specific information ■ Updated information about overriding HAL functions 	<ul style="list-style-type: none"> ■ Introduced external interrupt controller ■ BSP generation file-copy behavior changed
March 2009 v9.0.0	<ul style="list-style-type: none"> ■ Reorganized and updated information and terminology to clarify role of Nios II Software Build Tools. ■ Add documentation for Altera logging. ■ Corrected minor typographical errors. 	Altera logging
May 2008 v8.0.0	No change from previous release.	
October 2007 v7.2.0	<ul style="list-style-type: none"> ■ Added documentation for HAL program development with the Nios II Software Build Tools. ■ Additional documentation of alarms functions ■ Correct <code>alt_erase_flash_block()</code> example 	—
May 2007 v7.1.0	<ul style="list-style-type: none"> ■ Added table of contents to “Introduction” section. ■ Added Referenced Documents section. 	—
March 2007 v7.0.0	No change from previous release.	
November 2006 v6.1.0	<ul style="list-style-type: none"> ■ Program never exits system library option ■ Support C++ system library option ■ Lightweight device driver API system library option ■ Minimal character-mode API 	
May 2006 v6.0.0	<ul style="list-style-type: none"> ■ Revised text on instruction emulation. ■ Added section on global pointers. 	
October 2005 v5.1.0	<ul style="list-style-type: none"> ■ Added <code>alt_64</code> and <code>alt_u64</code> types to Table 6-1 on page 6-5. ■ Made changes to section “Placement of the Heap and Stack”. 	
May 2005 v5.0.0	Added <code>alt_load_section()</code> function information.	
December 2004 v1.2	<ul style="list-style-type: none"> ■ Added boot modes information. ■ Amended compiler optimizations. ■ Updated Reducing Code Footprint section. 	
September 2004 v1.1	Corrected DMA receive channels example code.	
May 2004 v1.0	Initial Release.	