



Hexagon Application Binary Interface

Specification

80-N2040-23 Rev. A

August 28, 2013

Submit technical questions at:
<https://support.cdmatech.com>

Qualcomm and Hexagon are trademarks of QUALCOMM Incorporated, registered in the United States and other countries. All QUALCOMM Incorporated trademarks are used with permission. Other product and brand names may be trademarks or registered trademarks of their respective owners.

This technical data may be subject to U.S. and international export, re-export, or transfer (“export”) laws. Diversion contrary to U.S. and international law is strictly prohibited.

**Qualcomm Technologies, Inc.
5775 Morehouse Drive
San Diego, CA 92121
U.S.A.**

© 2006, 2008, 2011-2013 Qualcomm Technologies, Inc.

Contents

1 Introduction.....	8
1.1 Overview	8
1.2 Using the document.....	9
2 Data Types.....	10
2.1 Overview	10
2.2 Basic data types	10
2.3 Memory alignment	11
2.3.1 Static allocation.....	12
2.3.2 Stack allocation.....	13
2.3.3 Heap allocation	13
3 Software Stack	14
3.1 Overview	14
3.2 Stack sections	14
4 Parameter Passing	16
4.1 Overview	16
4.2 Fixed argument list function calls	16
4.2.1 Examples.....	17
4.3 Variable argument list function calls	18
4.3.1 Examples.....	18
5 Return Values.....	19
5.1 Overview	19
6 Register Usage Across Calls.....	20
6.1 Overview	20
6.2 Outgoing memory arguments	20
7 C++ Exception Handling	21
7.1 Overview	21

8 Operating System Interface	22
8.1 Overview	22
9 Process Initialization	23
9.1 Overview	23
9.2 Special registers	23
9.3 General registers	23
10 Program Loading	24
10.1 Overview	24
10.2 Position-independent objects	24
11 Object Files	25
11.1 Overview	25
11.2 Compatibility	25
11.3 ELF header	26
11.4 Sections	27
11.5 Relocation	27
11.5.1 Relocation fields	28
11.5.2 Relocation symbols	29
11.5.3 Relocation types	30
11.5.4 Special relocation types	33
12 Program Headers	34
12.1 Overview	34
13 Dynamic Linking	35
13.1 Overview	35
13.2 Dynamic section	36
13.3 Global offset table	37
13.4 Procedure linkage table	37
14 Thread-Local Storage	40
14.1 Overview	40
14.2 Thread-local storage section	40
14.3 Runtime allocation	40
14.4 Load-time allocation	41
14.5 Interface	41
14.6 Thread-local storage access	41
14.6.1 General dynamic (GD)	42
14.6.2 Local dynamic (LD)	42

14.6.3 Initial executable (IE)	42
14.6.4 Local executable (LE).....	43
15 Coding Examples	44
15.1 Overview	44
15.2 Examples	45
15.2.1 Register use.....	45
15.2.2 Assembler symbols.....	45
15.2.3 Addressing constraints.....	47
15.3 Function prolog and epilog.....	47
15.3.1 Absolute	48
15.3.2 Position-independent.....	49
15.4 Direct function call.....	50
15.4.1 Absolute	50
15.4.2 Position-independent.....	50
15.5 Indirect function call.....	51
15.5.1 Absolute	51
15.5.2 Position-independent.....	52
15.6 Direct branch	53
15.6.1 Absolute & position-independent	53
15.7 Indirect branch.....	54
15.7.1 Absolute	54
15.7.2 Position-independent.....	55
15.8 Data access	56
15.8.1 Absolute	56
15.8.2 Position-independent.....	57
15.9 Thread-local storage	59
15.9.1 Absolute	59
15.9.2 Position-independent.....	62
A Version-Specific Coding Examples	63
A.1 Overview	63
A.2 Function prolog and epilog.....	64
A.2.1 Position-independent.....	64
A.3 Indirect function call.....	65
A.3.1 Position-independent.....	65
A.4 Data access	66
A.4.1 Position-independent.....	66
A.5 Thread-local storage	67
A.5.1 Absolute	67
A.5.2 Position-independent.....	70

Tables

Table 2-1	Basic data type representation	10
Table 2-2	Data type alignment	11
Table 6-1	Register usage across calls	20
Table 11-1	Object processor version flags	26
Table 11-2	Highest ISA version flags	26
Table 11-3	Section indexes for common small data	27
Table 11-4	Relocation fields	28
Table 11-5	Relocation symbols	29
Table 11-6	Relocation types	30
Table 11-7	Special relocation types	33
Table 13-1	Dynamic array tags	36
Table 15-1	Assembler symbols	45
Table 15-2	Function prolog and epilog (absolute)	48
Table 15-3	Function prolog and epilog (absolute – GD or LD access)	48
Table 15-4	Function prolog and epilog (absolute – IE or LE access)	48
Table 15-5	Direct function call (absolute)	50
Table 15-6	Direct function call (position-independent)	50
Table 15-7	Indirect function call (absolute)	51
Table 15-8	Indirect function call (absolute – no small data)	51
Table 15-9	Indirect function call (position-independent)	52
Table 15-10	Indirect function call (position-independent – no small data)	52
Table 15-11	Ind. function call (position-independent – no small data, GOT < 64KB).. 53	53
Table 15-12	Direct branch (absolute & position-independent)	53
Table 15-13	Data access (absolute)	56
Table 15-14	Data access (absolute – no small data)	56
Table 15-15	Data access (position-independent)	57
Table 15-16	Data access (position-independent – no small data)	57
Table 15-17	Data access (position-independent – no small data, GOT < 64KB)	58
Table 15-18	Thread-local storage (V4 absolute – GD access)	59
Table 15-19	Thread-local storage (V4 absolute – LD access)	60
Table 15-20	Thread-local storage (V4 absolute – IE access)	61
Table 15-22	Thread-local storage (V4 position-independent – IE access)	62
Table A-1	Indirect function call (V4 position-independent)	65
Table A-2	Data access (V4 position-independent)	66
Table A-3	Thread-local storage (V4 absolute – GD access)	67
Table A-4	Thread-local storage (V4 absolute – LD access)	68
Table A-5	Thread-local storage (V4 absolute – IE access)	69
Table A-7	Thread-local storage (V4 position-independent – IE access)	70

Revision History

Revision	Date	Description
A	May 2006	Initial release
B	Oct 2006	<ul style="list-style-type: none"> ■ By default “char” is now unsigned. ■ Bit-fields within structures are also unsigned by default. ■ Any code that relies on char variables or bit-fields being signed need an explicit “signed” keyword in the declaration. ■ Arrays are no longer aligned to 8-byte boundaries by default. Arrays are aligned based on the alignment of a single element. ■ Any specific alignment must be handled explicitly using the <code>aligned</code> attribute. For example: <pre>int buffer[N] __attribute__((aligned(8))) = {0};</pre> ■ A bug exists in the linker that uninitialized arrays with an <code>aligned</code> attribute are not properly aligned. Initialize all <code>aligned</code> arrays. ■ Variables of type <code>enum</code> are assigned the smallest type that fits all of the values, and alignment is based on the alignment of the underlying type.
C	Oct 2006	<ul style="list-style-type: none"> ■ Marked document as [Draft]. No changes to text.
D	Mar 2008	<ul style="list-style-type: none"> ■ Updated legal statements and usage of Hexagon. No changes to text.
E	Apr 2008	<ul style="list-style-type: none"> ■ Numerous corrections; more info on alignment; improved doc organization.
F	Dec 2008	<ul style="list-style-type: none"> ■ Hexagon V3 additionally defines R16-R23 as callee saved.
G	Dec 2008	<ul style="list-style-type: none"> ■ Hexagon V3 uses R0-R3 for exception handling.
H	May 2011	<ul style="list-style-type: none"> ■ Added Hexagon V4. Added following chapters: OS interface, process initialization, program loading, object files, program headers, dynamic linking, thread-local storage, coding examples.
J	March 2012	<ul style="list-style-type: none"> ■ Added Hexagon V5. ■ Corrected bitmaps for reloc fields <code>Word32_GP</code> and <code>Word32_R16</code>. ■ Globally replaced <code>Word32_R16</code> with <code>Word32_U16</code>. ■ Changed field for reloc type <code>R_HEX_9_X</code> to <code>Word32_U6</code>.
K	July 2012	<ul style="list-style-type: none"> ■ Removed nonexistent reloc type <code>R_HEX_IE_16</code>. ■ Added new reloc types <code>R_HEX_LD_PLT_B22_PCREL</code>, <code>R_HEX_LD_GOT_LO16</code>, <code>R_HEX_LD_GOT_HI16</code>, <code>R_HEX_LD_GOT_32</code>, <code>R_HEX_LD_GOT_16</code>, <code>R_HEX_LD_GOT_32_6_X</code>, <code>R_HEX_LD_GOT_16_X</code>, <code>R_HEX_LD_GOT_11_X</code> ■ Added asm symbols <code>name@LDPLT</code>, <code>name@LDGOT</code> ■ Rewrote Examples section in Code Examples chapter to clarify register use conventions ■ Code corrections in TLS and position-independent prolog/epilog code examples

Revision	Date	Description
L	January 2013	<ul style="list-style-type: none">■ Removed stub section (Symbol Table) from Object Files chapter■ Updated section 1.2 (text and URLs)
M	January 2013	<ul style="list-style-type: none">■ Corrected relocation formula for reloc types R_HEX_B22_PCREL_X, R_HEX_B15_PCREL_X, R_HEX_B13_PCREL_X, R_HEX_B9_PCREL_X, R_HEX_B7_PCREL_X (removed ">>2")
N	May 2013	<ul style="list-style-type: none">■ Updated processor versions (removed V2/V3-specific details; added V55)■ Merged coding examples appendix into chapter
A'	August 2013	<ul style="list-style-type: none">■ External doc version

1 Introduction

1.1 Overview

This document describes the application binary interface (ABI) for the Hexagon™ V4, V5, and V55 processors.

The ABI defines a set of conventions which enables the inter-operation of code written in different languages and at different times. These conventions include:

- Hardware representation of C data types
- Software stack configuration
- Parameter passing
- Return values
- Register usage across function calls
- C++ exception handling
- Operating system interface
- Process initialization
- Program loading
- Object files
- Program headers
- Dynamic linking
- Thread-local storage

Coding examples are provided for some of the basic operations.

NOTE Hexagon V2 and V3 object files are not compatible with object files for the other Hexagon processor versions ([Section 11.2](#)).

1.2 Using the document

The Hexagon ABI is based on the System V ABI.

This document describes only the Hexagon-specific aspects of the ABI. For information on the complete ABI specification, see the following documents:

- *System V Application Binary Interface, Edition 4.1*
(<http://www.sco.com/developers/devspecs/gabi41.pdf>)
- *Solaris Linker and Libraries Guide*, chapter 8 (*Thread-Local Storage*)
(<http://docs.oracle.com/cd/E19253-01/817-1984/>)

2 Data Types

2.1 Overview

This chapter describes how C data types are represented on the Hexagon processor. It covers the following topics:

- Hardware representation of C data types
- Memory alignment requirements

2.2 Basic data types

Table 2-1 lists the basic C data types and how they are represented in hardware on the Hexagon processor.

Table 2-1 Basic data type representation

C Type	Size (in bytes)	Hardware Representation
Char	1	Byte
Short	2	Halfword
Int	4	Word
Long	4	Word
Long long	8	Doubleword
Float	4	Word
Double	8	Doubleword
Long double	8	Doubleword
Enum	variable	variable
Pointer	4	Word

By default the `char` data type is unsigned. If a `char` variable must be signed, declare the variable with the type `signed char`.

By default enumeration types are assigned the smallest integer type that can store the enumeration values. For example, an `enum` type containing constants in the range 0-255 is stored in memory as a single byte.

NOTE The command option `-fno-short-enums` causes the C compilers to allocate all enumeration types in 4 bytes.

2.3 Memory alignment

The Hexagon processor requires data types to be properly aligned in memory when they are accessed. An item of size N bytes is *aligned* when its memory address mod N yields zero. For example, a 4-byte memory access is aligned when the address is an integral multiple of 4.

Table 2-2 lists various C data types and their corresponding alignment requirements.

Table 2-2 Data type alignment

Data Type	Alignment Requirement
Scalar	Align to size of item (Section 2.2).
Enumeration	Align according to underlying data type.
Array	<ul style="list-style-type: none"> ■ Align array according to data type of its array element. ■ Align array elements according to their data type.
Structure	<ul style="list-style-type: none"> ■ Align structure to alignment of structure member with largest alignment. ■ Align members according to their data type. ■ If necessary, insert pad bytes between members. ■ If necessary, insert pad bytes after last member to make the structure size an integral multiple of the alignment requirement. (This ensures correct pointer arithmetic for pointers to structures.) <p>NOTE Members are allocated in memory in the order they appear in a <code>struct</code> definition.</p>
Bit field	<ul style="list-style-type: none"> ■ Same size and alignment rules as non-bit field structure members. ■ Allocated from right to left (i.e., least to most significant). ■ Must reside entirely within a storage unit appropriate for their data type. ■ Unsigned by default. To create signed bit fields use the <code>signed</code> type specifier.
Union	Align to alignment of member with largest alignment.

Data types need to be aligned only when they are accessed. Unaligned types are allowed and should be handled properly by the compilers. Problems arise only when dereferencing pointers are cast to data types with stricter alignment requirements than the original type. For example:

```
int g(int *p){
    return *p;
}

int f(void){
    unsigned char a[4];
    a[0] = 0xbeU;
    a[1] = 0xbaU;
    a[2] = 0xfeU;
    a[3] = 0xcaU;
    return g(a);
}
```

In this example array `a` is guaranteed only to be 1-byte aligned, but function `g()` accesses `a` as an `int`, which (unless it is declared with the `packed` attribute) must be 4-byte aligned. Thus the compiler will generate a `memw` instruction in `g()`.

Alignment attributes

Nonstandard memory alignment requirements must be explicitly specified using the `aligned` attribute. For example:

```
int buffer[N] __attribute__((aligned(8)));
```

NOTE The programmer can specify any alignment value larger than the one required.

The `packed` attribute can be used to declare that a data type should not be aligned. `packed` and `aligned` can be used together to align a data type to more than 1 byte but less than the type's "natural" alignment. For example:

```
int i __attribute__((packed, aligned(2)));
```

In this example `i` will be 2-byte aligned.

2.3.1 Static allocation

Data types statically allocated in memory are subject to the alignment requirements listed in [Section 2.3](#). No additional requirements exist.

2.3.2 Stack allocation

Data types allocated on the software stack ([Chapter 3](#)) are subject to the alignment requirements listed in [Section 2.3](#).

Data allocated on the stack is only guaranteed to be up to 8-byte aligned. To allocate stack data with more than 8-byte alignment, the data address must be manually aligned. This can be done with the following procedure:

1. Declare a local `char` array with the following number of elements:
`(size(data) + alignment(data) - 1)`
2. Cast the address of the array to an integer.
3. Clear the $\log_2(\text{alignment})$ low bits of the integer.
4. Cast the resulting value to a pointer to the type to be allocated.

NOTE The compilers do not currently support stack realignment.

2.3.3 Heap allocation

Data types allocated on the heap are subject to the alignment requirements listed in [Section 2.3](#).

The standard allocator functions (`malloc`, `realloc`, `calloc`) return a pointer that is aligned to the largest alignment required by any type in the absence of the aligned attribute. For the Hexagon processor this value is currently 8 bytes.

To allocate an object on the heap with alignment greater than 8 bytes, use the `memalign` function.

NOTE `memalign` automatically performs the procedure described in [Section 2.3.2](#).

3 Software Stack

3.1 Overview

The software stack is divided into *frames*. In general, there is one frame for each active function. The exception to this rule is described later. First, we describe the layout of each frame.

The current stack frame is marked by the *frame pointer register (R30)* and the *stack pointer register (R29)*. The stack grows toward smaller addresses in memory. Stack frames are always aligned on 8-byte boundaries. This enables the alignment of data within each frame.

3.2 Stack sections

Each stack frame contains the following sections, in decreasing memory order:

1. **Saved R31, saved R30**

We save the values of R31 (the function return address) and R30 (frame pointer) during the function prologue. This is done with the *allocframe* instruction.

2. **Locals/spill area**

This optional area contains enough space for any local variables that require stack space and any register spills (including callee saved registers) needed. The size of this area is determined at compile time on a per function basis. It is recommended, but not required, that the compiler or programmer order the data in this section to minimize the amount of padding between variables.

3. **Alloca area**

This optional area contains space allocated by *alloca*. The size of this area is determined at run time based on the *alloca* calls that are executed. Initially, its size is zero.

4. **Outgoing memory arguments**

This optional area is used when the function calls one or more functions whose arguments cannot all be passed in registers. The size of this area is determined at compile time on a per function basis. Within the function, the compiler examines all function calls and determines the maximum space required for any call. Note that this scheme makes function calls fast and simple at the expense of some wasted space.

Each section should be aligned to an 8-byte boundary. If necessary, padding should be added in the high memory address of each section.

Within each section, the normal data alignment rules apply ([Section 2.3](#)).

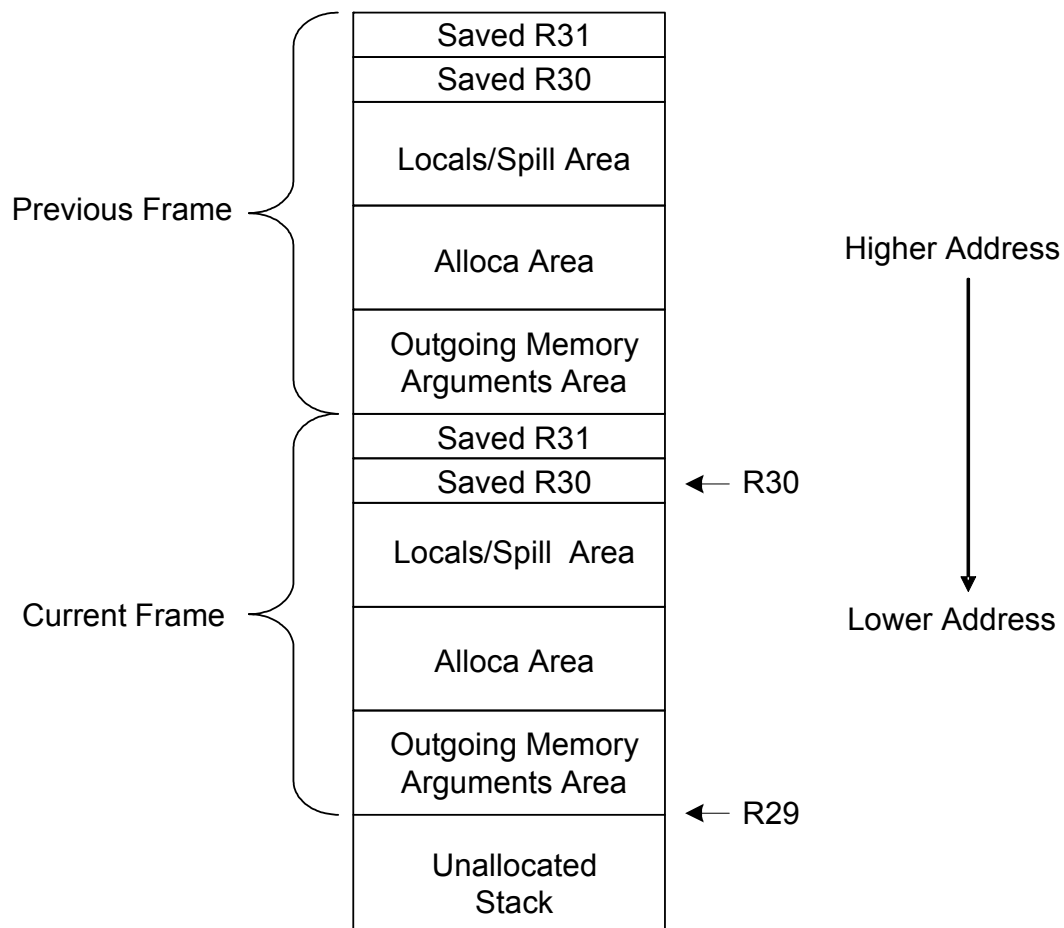


Figure 3-1 Stack in memory

NOTE For performance reasons certain functions are allowed to avoid the overhead of allocating and deallocating a frame. A function that does not make any calls is called a *leaf function*. A leaf function that does not require any stack space (i.e., the size of each optional area is zero) may not allocate a frame.

4 Parameter Passing

4.1 Overview

Parameters are passed to functions in up to six general purpose registers: R0–R5. Any parameters that cannot fit in these registers are passed on the software stack.

The calling function allocates space in the lower portion of its frame for these. Each function allocates the maximum amount of space required by any of its calls in its outgoing memory arguments area.

If the function does not make any calls, or if all of its calls can pass all parameters in registers, then the outgoing memory arguments area of the stack frame is not allocated.

4.2 Fixed argument list function calls

The most common case for function calls is a fixed argument length function. In C, these are functions with:

- A function prototype that contains a fixed number of arguments.
- No function prototype, also known as Kernighan and Ritchie (K&R) style.

Note that a variable argument length function whose prototype is missing at the call site can lead to subtle program errors. This is due to the mismatch between the parameter passing of the caller and callee.

For each fixed argument length call, the parameters are processed from left to right as follows:

- Any parameter (including aggregates) with size up to 64 bits is a *candidate* to be passed in a register. Parameters up to 32 bits in size are passed in a single register. Larger parameters are passed in a register pair.
- Candidates are assigned to registers R0 through R5 in order.
- Candidates larger than 32 bits are passed in a register pair. If the next available register is an odd-numbered register, that register is skipped (left unused), and the candidate is placed in the next even-odd pair. If the next available register is R5, then the candidate is passed on the stack.
- Candidates larger than 64 bits are passed on the stack. The alignment rules are the same as for static data ([Section 2.3](#)).

- Note that we do not attempt to recover skipped parameter registers by filling them with subsequent parameters. However, we will continue using parameter registers after a candidate larger than 64 bits has been placed on the stack.
- Once all the parameter registers have been filled or skipped, all remaining parameters are passed on the stack. The normal alignment rules apply ([Section 2.3](#)).

4.2.1 Examples

Example 1: scalars

```
extern int foo(short, float, int, double);
foo(i, a, b, c);
```

The parameters are passed as follows:

```
i R0
a R1
b R2
c R5:R4
```

Example 2: arrays and structures

```
extern int bar();
struct {int length, width;} st1;
struct {int a; int bvec[8];} st2;
bar(i, st1, st2);
```

The parameters are passed as follows:

```
i R0
st1 R3:R2
st2.a In memory, at R29+0
st2.bvec In memory, at R29+4
```

NOTE R29 is the stack pointer (sp).

4.3 Variable argument list function calls

Variable argument list functions are a special type of function in C where the number and types of arguments can vary. The best-known example of this is the `printf` function in the C library, where the function's number and type of arguments change, depending on the first argument given: a format string. In general, variable argument functions determine the number and type of their arguments by programming conventions.

For these functions, we pass the named (and typed) parameters in the same manner as for fixed argument list functions. The remainder of the arguments are passed on the stack. The normal alignment rules apply ([Section 2.3](#)).

4.3.1 Examples

Example 1

```
extern int vfoo(int, long long, short, ...);
vfoo(a1, a2, a3, a4, (double)a5, a6);
```

The parameters are passed as follows:

```
a1 R0
a2 R3:R2
a3 R4
a4 In memory, at R29+0
a5 In memory, at R29+8, due to 8 byte alignment
a6 In memory, at R29+16
```

NOTE R29 is the stack pointer (`sp`).

5 Return Values

5.1 Overview

Functions can return values up to 32 bits in size (including structures) in register R0 and up to 64 bits in size (including structures) in the R1:0 register pair. For return values larger than 64 bits, the caller allocates space in its frame large enough to hold the return value.

The space is aligned according to the normal alignment rules for the returned type ([Section 2.3](#)). The caller passes an extra argument at the beginning of the parameter list (i.e., in R0) containing an address of the allocated space.

Example:

```
typedef struct {
    int a;
    int b;
    int c;
} S;

S foo(int x)
{
    S s;
    ...
    return s;
}

void bar(int y)
{
    S s = foo(y);
}
```

The above code will be automatically converted by the compiler to the following:

```
void foo(S *hidden, int x)
{
    S s;
    ...
    *hidden = s;
}

void bar(int y)
{
    S s;
    foo(&s, y);
}
```

6 Register Usage Across Calls

6.1 Overview

The registers are divided into two groups, according to how they are used across function calls:

- **Caller saved** – At a call site, the caller must assume that the callee changes the values of these registers. Therefore, the *caller* must save the values of these registers before making the call (if necessary), and restore them after the call.
- **Callee saved** – At a call site, the caller can assume that the callee does *not* change the values of these registers. Therefore, the *callee* must save and restore the values of these registers, if necessary.

Table 6-1 lists the register usage across calls.

Table 6-1 Register usage across calls

Register	Usage	Saved by
R0 - R5	Parameters ^a	–
R6 - R15	Scratch ^b	Caller
R16 - R27	Scratch	Callee
R28	Scratch ^b	Caller
R29 - R31	Stack frames	Callee ^c
P3:0	Processor state	Caller

^a The callee can change parameter values (Chapter 4).

^b R14-R15 and R28 are used by the procedure linkage table (Section 13.4).

^c R29-R31 are saved and restored by allocframe and deallocframe (Chapter 3).

6.2 Outgoing memory arguments

The outgoing-memory-arguments section of the caller's frame is caller-saved. In other words, the caller should assume that the callee will overwrite the values of parameters passed. Note that this only applies to the outgoing memory arguments space used by the current call, and not to the total space allocated by the caller.

7 C++ Exception Handling

7.1 Overview

The exception handling model for the Hexagon processor is based on the DWARF2 stack unwinding mechanism.

Registers R0 to R3 are reserved to communicate between exception handling library routines and exception handlers.

The stack location for saved LR (see [Figure 3-1](#)) is used to store the address of an exception handler to which the procedure returns.

8 Operating System Interface

8.1 Overview

Processes run in a 32-bit virtual address space. The Hexagon memory management hardware translates virtual addresses into physical addresses.

The memory management system supports the following page sizes (in bytes): 4K, 16K, 64K, 256K, 1M, 4M, 16M.

9 Process Initialization

9.1 Overview

The initial state of a process is established by setting certain Hexagon processor registers to well-defined initial values.

9.2 Special registers

The GP register is set to the starting address of the process's *small data area*, as referenced by the program symbol `_SDA_BASE_`.

The UGP register is set to the highest-memory-address-plus-one of the process's thread-local storage area ([Chapter 14](#)).

The R28 register is set to the address of a function which the process must call when it terminates. The process is responsible for saving this address so it can be called later on. After saving the address, the process can freely use the R28 register.

NOTE All other special registers contain unspecified values at process initialization, and operating systems are free to specify their contents.

9.3 General registers

Except for R28 and R29 (also known as SP, which points to the top of the stack), the Hexagon general registers contain unspecified values at process initialization.

NOTE The contents of stack memory below the top of stack is always undefined.

10 Program Loading

10.1 Overview

When loading an object into memory, the relative positions of the object's several segments must be preserved.

Additionally, absolute objects must reside at the same virtual address that was used to build them.

10.2 Position-independent objects

A position-independent object must be loaded at a virtual address with address alignment equal to the alignment of its largest segment. This ensures the alignment of all the object's segments.

The small data area must be aligned to a 64-byte address boundary. Therefore, position-independent objects that contain a small data area must be loaded at a virtual address aligned to 64 bytes.

NOTE The alignment requirement for position-independent objects may change in future versions of the Hexagon processor.

11 Object Files

11.1 Overview

Hexagon object files are stored in the ELF file format (short for *Executable and Linkable Format*). This chapter describes object file compatibility and how various ELF elements are defined in Hexagon. It covers the following topics:

- Compatibility
- ELF header
- ELF sections
- Relocation

11.2 Compatibility

Object files for the Hexagon V4, V5, and V55 processors are backward compatible when executed in user mode.

Kernel-mode object files are not backward compatible due to changes in the system-level instruction set.

NOTE Hexagon V2 and V3 object files are not compatible with object files for the other Hexagon processor versions.

11.3 ELF header

The ELF header member `e_machine` is set to the symbolic value `EM_HEXAGON` (which has the decimal numeric value 164).

The member `e_type` can additionally be set to the symbolic value `ET_HEXAGON_IR` (numeric value `0xff00`) to indicate that an object file contains compiler-generated intermediary representation language.

Bits [3:0] of the member `e_flags` indicate the version of the Hexagon processor that the object file was created for. [Table 11-1](#) lists the possible values for this bit field.

Table 11-1 Object processor version flags

Name	Value	Processor Version
<code>EF_HEXAGON_MACH_V2</code>	<code>0x1</code>	Hexagon V2
<code>EF_HEXAGON_MACH_V3</code>	<code>0x2</code>	Hexagon V3
<code>EF_HEXAGON_MACH_V4</code>	<code>0x3</code>	Hexagon V4
<code>EF_HEXAGON_MACH_V5</code>	<code>0x4</code>	Hexagon V5
<code>EF_HEXAGON_MACH_V55</code>	<code>0x5</code>	Hexagon V55

Bits [7:4] of `e_flags` indicate the *highest* version of the Hexagon processor that the object file contains code for.

In certain cases, object files created for different Hexagon versions can be linked together into a single object file containing code from the different versions. As long as this code is execution-compatible on the target processor version, the resulting object file is considered a valid executable, and its mixed-code status is indicated in this bit field.

[Table 11-2](#) lists the possible values for the highest-version bit field.

Table 11-2 Highest ISA version flags

Name	Value	Description
<code>EF_HEXAGON_ISA_MACH</code>	<code>0x00</code>	Indicates that processor version is same as specified in bits [3:0] of <code>e_flags</code>
<code>EF_HEXAGON_ISA_V2</code>	<code>0x10</code>	Hexagon V2 ISA
<code>EF_HEXAGON_ISA_V3</code>	<code>0x20</code>	Hexagon V3 ISA
<code>EF_HEXAGON_ISA_V4</code>	<code>0x30</code>	Hexagon V4 ISA
<code>EF_HEXAGON_ISA_V5</code>	<code>0x40</code>	Hexagon V5 ISA
<code>EF_HEXAGON_ISA_V55</code>	<code>0x50</code>	Hexagon V55 ISA

Code from the V4, V5, and V55 processor versions can be mixed in an object file; in this case, the highest-version field is always set to the highest processor version in the mix.

NOTE “ISA” is short for “Instruction Set Architecture”, which in this context can be considered equivalent to “processor version”.

11.4 Sections

In addition to supporting the standard ELF section indexes, Hexagon object files may also include processor-specific section indexes for common symbols stored in the small data area ([Section 9.2](#)). These indexes are listed in [Table 11-3](#).

The suffix on the processor-specific section index names indicates the access size (in bytes) of the data in the corresponding section. The one name lacking a suffix is used to specify sections that have an access size other than one of the explicitly-specified values.

Table 11-3 Section indexes for common small data

Name	Value	Description
SHN_HEXAGON_SCOMMON	0xff00	Other access sizes
SHN_HEXAGON_SCOMMON_1	0xff01	Byte-sized access
SHN_HEXAGON_SCOMMON_2	0xff02	Half-word-sized access
SHN_HEXAGON_SCOMMON_4	0xff03	Word-sized access
SHN_HEXAGON_SCOMMON_8	0xff04	Double-word-size access

The member `sh_flags` can be set to the symbolic value `SHF_HEXAGON_GPREL` (numeric value 0x10000000) to mark a section as residing in the small data area.

11.5 Relocation

Many different types of relocations are performed on the instructions and data stored in a Hexagon object file. The relocation types differ according to the instruction and data fields that are affected, and the kinds of calculations that are applied to the fields to perform the relocation.

11.5.1 Relocation fields

The relocation types are expressed using the data fields defined in [Table 11-4](#):

- **Name** specifies the name of the relocatable field (referenced in [Table 11-6](#)).
- **Width** specifies the width (in bits) of the relocatable bit field. Depending on the relocation type, either all or some of these bits are modified by the relocation.
- **Effective Bits** specifies how many of the bits in the relocation field are modified.
- **Bitmap** specifies a bit pattern which indicates the bits in the relocation field that are actually modified by the relocation.
- **Byte Address Alignment** specifies the byte address alignment (1, 2, or 4) that is required for the relocation field.

Table 11-4 Relocation fields

Name	Width	Effective Bits	Bitmap	Byte Address Alignment
Word8	8	8	0xff	1
Word16	16	16	0xffff	2
Word32	32	32	0xffffffff	4
Word32_LO	32	16	0x00c03fff	4
Word32_HL	64 ^a	16 and 16	0x00c03fff 0x00c03fff	4
Word32_GP	32	16	n/a ^b	4
Word32_B7	32	7	0x00001f18	4
Word32_B9	32	9	0x003000fe	4
Word32_B13	32	13	0x00202ffe	4
Word32_B15	32	15	0x00df20fe	4
Word32_B22	32	22	0x01ff3ffe	4
Word32_R6	32	6	0x000007e0	4
Word32_U6	32	6	n/a ^b	4
Word32_U16	32	16	n/a ^b	4
Word32_X26	32	26	0x0fff3fff	4

^a The relocation calculations are applied to two 32-bit words in sequential addresses: first “(S + A) >> 16”, then “(S + A)”.

^b The bitmap varies according to the instruction opcode.

NOTE Bitmaps are used to specify operands in instruction op-codes that must be relocated. When applying a relocation in these cases, the bits that are not specified in the bitmap must be protected from accidental modification.

11.5.2 Relocation symbols

Each relocation type is defined as a formula which is expressed in terms of the symbols listed in [Table 11-5](#).

Table 11-5 Relocation symbols

Symbol	Description
A	Addend used to compute value of relocatable field.
B	Base address of object loaded into memory.
G	Offset into global offset table (Section 13.3) for the entry of a symbol.
GOT	Address of entry zero in global offset table (Section 13.3).
GP	Value of GP register, typically the base address of the small data area (<code>_SDA_BASE_</code>).
L	Offset into procedure linkage table (Section 13.4) for the entry of a symbol.
P	Place address of the field being relocated. NOTE - This value is computed using <code>r_offset</code> in the relocation entry.
S	Value of the symbol whose index resides in the relocation entry (unless the object is a shared object and the symbol index is <code>SHN_UNDEF</code> , in which case this is equivalent to symbol B).
TLS	Thread-pointer-relative offset to a thread-local symbol.
T	Base address of the static thread-local template that contains a thread-local symbol.

11.5.3 Relocation types

The relocation types are applied in a similar manner to relocatable, executable and shared object files, except where noted otherwise.

Only `Elf32_Rela` relocation entries are used. Therefore, the original content of the instruction and data fields is irrelevant when calculating the relocation.

The relocation types are defined in [Table 11-6](#):

- **Name** specifies the symbolic name of the relocation type.
- **Value** specifies the corresponding numeric identifier for the relocation type.
- **Field** specifies the data field affected by the relocation type ([Table 11-4](#)).
- **Relocation** specifies the algebraic formula used to perform the relocation (with the formula symbols defined in [Table 11-5](#)).
- **Result** specifies the numeric format (signed or unsigned) of the relocated value.
- **Action** specifies an additional operation performed on the relocated value as part of the relocation:
 - **Truncate** – Truncate the relocated value to fit the relocation field.
 - **Verify** – Check that the relocated value can fit in the relocation field, and generate an error if the value is too big to fit in its relocation field.
 - **None** – Perform no operation on the relocated value.

Table 11-6 Relocation types

Name	Value	Field	Relocation	Result	Action
R_HEX_NONE	0	None	None	None	None
R_HEX_B22_PCREL	1	Word32_B22	$(S + A - P) \gg 2$	Signed	Verify
R_HEX_B15_PCREL	2	Word32_B15	$(S + A - P) \gg 2$	Signed	Verify
R_HEX_B7_PCREL	3	Word32_B7	$(S + A - P) \gg 2$	Signed	Verify
R_HEX_LO16	4	Word32_LO	$(S + A)$	Unsigned	Truncate
R_HEX_HI16	5	Word32_LO	$(S + A) \gg 16$	Unsigned	Truncate
R_HEX_32	6	Word32	$(S + A)$	Unsigned	Truncate
R_HEX_16	7	Word16	$(S + A)$	Unsigned	Truncate
R_HEX_8	8	Word8	$(S + A)$	Unsigned	Truncate
R_HEX_GPREL16_0	9	Word32_GP	$(S + A - GP)$	Unsigned	Verify
R_HEX_GPREL16_1	10	Word32_GP	$(S + A - GP) \gg 1$	Unsigned	Verify
R_HEX_GPREL16_2	11	Word32_GP	$(S + A - GP) \gg 2$	Unsigned	Verify
R_HEX_GPREL16_3	12	Word32_GP	$(S + A - GP) \gg 3$	Unsigned	Verify
R_HEX_HL16	13	Word32_HL	$(S + A) \gg 16$ and $(S + A)$	Unsigned	Truncate
R_HEX_B13_PCREL	14	Word32_B13	$(S + A - P) \gg 2$	Signed	Verify

Table 11-6 Relocation types (continued)

Name	Value	Field	Relocation	Result	Action
R_HEX_B9_PCREL	15	Word32_B9	(S + A - P) >> 2	Signed	Verify
R_HEX_B32_PCREL_X	16	Word32_X26	(S + A - P) >> 6	Signed	Truncate
R_HEX_32_6_X	17	Word32_X26	(S + A) >> 6	Unsigned	Verify
R_HEX_B22_PCREL_X	18	Word32_B22	(S + A - P) & 0x3f	Signed	Verify
R_HEX_B15_PCREL_X	19	Word32_B15	(S + A - P) & 0x3f	Signed	Verify
R_HEX_B13_PCREL_X	20	Word32_B13	(S + A - P) & 0x3f	Signed	Verify
R_HEX_B9_PCREL_X	21	Word32_B9	(S + A - P) & 0x3f	Signed	Verify
R_HEX_B7_PCREL_X	22	Word32_B7	(S + A - P) & 0x3f	Signed	Verify
R_HEX_16_X	23	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_12_X	24	Word32_R6	(S + A)	Unsigned	Truncate
R_HEX_11_X	25	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_10_X	26	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_9_X	27	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_8_X	28	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_7_X	29	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_6_X	30	Word32_U6	(S + A)	Unsigned	Truncate
R_HEX_32_PCREL	31	Word32	(S + A - P)	Signed	Verify
R_HEX_COPY	32	Word32	(see below)		
R_HEX_GLOB_DAT	33	Word32	(S + A) (see below)	Unsigned	Truncate
R_HEX_JMP_SLOT	34	Word32	(S + A) (see below)	Unsigned	Truncate
R_HEX_RELATIVE	35	Word32	(B + A) (see below)	Unsigned	Truncate
R_HEX_PLT_B22_PCREL	36	Word32_B22	(L + A - P) >> 2	Signed	Verify
R_HEX_GOTREL_LO16	37	Word32_LO	(S + A - GOT)	Signed	Truncate
R_HEX_GOTREL_HI16	38	Word32_LO	(S + A - GOT) >> 16	Signed	Truncate
R_HEX_GOTREL_32	39	Word32	(S + A - GOT)	Signed	Truncate
R_HEX_GOT_LO16	40	Word32_LO	(G)	Signed	Truncate
R_HEX_GOT_HI16	41	Word32_LO	(G) >> 16	Signed	Truncate
R_HEX_GOT_32	42	Word32	(G)	Signed	Truncate
R_HEX_GOT_16	43	Word32_U16	(G)	Signed	Verify
R_HEX_DTPMOD_32	44	Word32			
R_HEX_DTPREL_LO16	45	Word32_LO	(S + A - T)	Signed	Truncate
R_HEX_DTPREL_HI16	46	Word32_LO	(S + A - T) >> 16	Signed	Truncate
R_HEX_DTPREL_32	47	Word32	(S + A - T)	Signed	Truncate
R_HEX_DTPREL_16	48	Word32_U16	(S + A - T)	Signed	Verify
R_HEX_GD_PLT_B22_PCREL	49	Word32_B22	(L + A - P) >> 2	Signed	Verify

Table 11-6 Relocation types (continued)

Name	Value	Field	Relocation	Result	Action
R_HEX_GD_GOT_LO16	50	Word32_LO	(G)	Signed	Truncate
R_HEX_GD_GOT_HI16	51	Word32_LO	(G) >> 16	Signed	Truncate
R_HEX_GD_GOT_32	52	Word32	(G)	Signed	Truncate
R_HEX_GD_GOT_16	53	Word32_U16	(G)	Signed	Verify
R_HEX_IE_LO16	54	Word32_LO	(G + GOT)	Signed	Truncate
R_HEX_IE_HI16	55	Word32_LO	(G + GOT) >> 16	Signed	Truncate
R_HEX_IE_32	56	Word32	(G + GOT)	Signed	Truncate
R_HEX_IE_GOT_LO16	57	Word32_LO	(G)	Signed	Truncate
R_HEX_IE_GOT_HI16	58	Word32_LO	(G) >> 16	Signed	Truncate
R_HEX_IE_GOT_32	59	Word32	(G)	Signed	Truncate
R_HEX_IE_GOT_16	60	Word32_U16	(G)	Signed	Verify
R_HEX_TPREL_LO16	61	Word32_LO	(TLS - S - A)	Signed	Truncate
R_HEX_TPREL_HI16	62	Word32_LO	(TLS - S - A) >> 16	Signed	Truncate
R_HEX_TPREL_32	63	Word32	(TLS - S - A)	Signed	Truncate
R_HEX_TPREL_16	64	Word32_U16	(TLS - S - A)	Signed	Verify
R_HEX_6_PCREL_X	65	Word32_U6	(S + A - P)	Unsigned	Truncate
R_HEX_GOTREL_32_6_X	66	Word32_X26	(S + A - GOT) >> 6	Signed	Truncate
R_HEX_GOTREL_16_X	67	Word32_U6	(S + A - GOT)	Unsigned	Truncate
R_HEX_GOTREL_11_X	68	Word32_U6	(S + A - GOT)	Unsigned	Truncate
R_HEX_GOT_32_6_X	69	Word32_X26	(G) >> 6	Signed	Truncate
R_HEX_GOT_16_X	70	Word32_U6	(G)	Signed	Verify
R_HEX_GOT_11_X	71	Word32_U6	(G)	Unsigned	Truncate
R_HEX_DTPREL_32_6_X	72	Word32_X26	(S + A - T) >> 6	Signed	Truncate
R_HEX_DTPREL_16_X	73	Word32_U6	(S + A - T)	Unsigned	Truncate
R_HEX_DTPREL_11_X	74	Word32_U6	(S + A - T)	Unsigned	Truncate
R_HEX_GD_GOT_32_6_X	75	Word32_X26	(G) >> 6	Signed	Truncate
R_HEX_GD_GOT_16_X	76	Word32_U6	(G)	Unsigned	Truncate
R_HEX_GD_GOT_11_X	77	Word32_U6	(G)	Unsigned	Truncate
R_HEX_IE_32_6_X	78	Word32_X26	(G + GOT) >> 6	Signed	Truncate
R_HEX_IE_16_X	79	Word32_U6	(G + GOT)	Unsigned	Truncate
R_HEX_IE_GOT_32_6_X	80	Word32_X26	(G) >> 6	Signed	Truncate
R_HEX_IE_GOT_16_X	81	Word32_U6	(G)	Unsigned	Truncate
R_HEX_IE_GOT_11_X	82	Word32_U6	(G)	Unsigned	Truncate
R_HEX_TPREL_32_6_X	83	Word32_X26	(TLS - S - A) >> 6	Signed	Truncate
R_HEX_TPREL_16_X	84	Word32_U6	(TLS - S - A)	Unsigned	Truncate

Table 11-6 Relocation types (continued)

Name	Value	Field	Relocation	Result	Action
R_HEX_TPREL_11_X	85	Word32_U6	(TLS - S - A)	Unsigned	Truncate
R_HEX_LD_PLT_B22_PCREL	86	Word32_B22	(L + A - P) >> 2	Signed	Verify
R_HEX_LD_GOT_LO16	87	Word32_LO	(G)	Signed	Truncate
R_HEX_LD_GOT_HI16	88	Word32_LO	(G) >> 16	Signed	Truncate
R_HEX_LD_GOT_32	89	Word32	(G)	Signed	Truncate
R_HEX_LD_GOT_16	90	Word32_R16	(G)	Signed	Verify
R_HEX_LD_GOT_32_6_X	91	Word32_X26	(G) >> 6	Signed	Truncate
R_HEX_LD_GOT_16_X	92	Word32_U6	(G)	Unsigned	Truncate
R_HEX_LD_GOT_11_X	93	Word32_U6	(G)	Unsigned	Truncate

NOTE If a relocation formula contains a reference to either G or GOT, an entry for the indicated symbol is implicitly created in the global offset table (“GOT”).

If a relocation formula contains a reference to L, an entry for the indicated symbol is implicitly created in the procedure linkage table (“PLT”).

11.5.4 Special relocation types

Four of the relocation types listed in [Table 11-6](#) have special semantics – they are the only ones the loader should support, while generating an error for the other relocation types.

[Table 11-7](#) lists the special relocation types.

Table 11-7 Special relocation types

Name	Description
R_HEX_COPY	Dynamic linking support – the offset refers to a location in a writable segment, and the symbol table index specifies a symbol that should exist both in the current object file and in a shared object. During execution, the dynamic linker copies the field associated with the shared object's symbol into the location in the current object specified by the offset.
R_HEX_GLOB_DAT	Similar to relocation type R_HEX_32, except that it sets a GOT entry to the address of the specified symbol.
R_HEX_JMP_SLOT	Dynamic linking support – the offset refers to the location of a GOT entry, which the dynamic linker modifies with the address of the specified symbol.
R_HEX_RELATIVE	Dynamic linking support – the offset refers to a location in a shared object which contains a value representing a relative address. The loader computes the corresponding virtual address by adding the relative address to the virtual address that the shared object was loaded at. NOTE The symbol table index must be set to 0.

12 Program Headers

12.1 Overview

When a loadable segment is loaded into memory, the contents of element `p_flags` specify the corresponding segment permissions.

13 Dynamic Linking

13.1 Overview

Dynamic linking involves loading objects into an application program (and linking them) at runtime rather than compile time.

The mechanics of shared object loading are specific to the target operating system. Hexagon object files use the same data structures for dynamic linking that are used in the System V ABI ([Section 1.2](#)).

Shared objects are often intended to be referenced by more than one client object, which requires that their code be executable in a reentrant manner. Reentrant objects are commonly implemented by maintaining separate instances of the object data section for each reference to the shared object.

The mechanism used for creating multiple data section instances is implementation-defined. Here are two examples of how the data sections can be implemented:

1. The object loader marks all pages for a data section as read-only, and relies on the virtual memory manager to create writeable copies of the pages, which are written to privately from the calling-process context (a policy known as *copy-on-write*).
2. The object loader creates a new instance of the shared object for each reference.

13.2 Dynamic section

In addition to supporting the standard ELF dynamic array tags, Hexagon object files may also include processor-specific dynamic array tags to support dynamic linking.

[Table 13-1](#) lists the processor-specific dynamic array tags (along with the standard ELF dynamic array tag `DT_PLTGOT`):

- **Name** specifies the name of the tag.
- **Value** specifies the numeric identifier of the tag.
- **d_un** specifies whether the tag represents an integer or program virtual address. (`d_un` is the name of a union that is used to store elements in a dynamic section.)
- **Executable** specifies whether or not the dynamic linking array for an executable object file must contain a tag entry of the specified type. `Optional` indicates that an entry can appear in the array, but is not required.
- **Shared Object** specifies whether or not the dynamic linking array for a shared object file must contain a tag entry of the specified type.

Table 13-1 Dynamic array tags

Name	Value	d_un	Executable	Shared Object
<code>DT_PLTGOT</code>	<code>0x00000003</code>	<code>d_ptr</code>	Optional	Optional
<code>DT_HEXAGON_SYMSZ</code>	<code>0x70000000</code>	<code>d_val</code>	Optional	Optional
<code>DT_HEXAGON_VER</code>	<code>0x70000001</code>	<code>d_val</code>	Mandatory	Mandatory
<code>DT_HEXAGON_PLT</code>	<code>0x70000002</code>	<code>d_val</code>	Optional	Optional

These tags have the following descriptions:

- `DT_PLTGOT` – Image offset of the GOT ([Section 13.3](#)).
- `DT_HEXAGON_SYMSZ` – Size (in bytes) of the symbol table pointed by `DT_SYMTAB`. This value is equivalent to the value of `DT_SYMENT` multiplied by the value of field `nchain` in the hash table pointed by `DT_HASH`.
- `DT_HEXAGON_VER` – Version of interface with dynamic linker ([Section 13.4](#)). Currently it can be set to the integer value 2 or 3. For the object to be compatible with the Hexagon ABI, it must be set to 3. The default is 2.
- `DT_HEXAGON_PLT` – Image offset of the PLT ([Section 13.4](#)).

13.3 Global offset table

The global offset table (or “GOT”) is an array of absolute addresses which enables an object to use position-independent code.

At load-time, the relocations related to GOT-relative dynamic symbols are resolved, and the GOT is populated with their absolute addresses.

The zero'th GOT entry (index = 0) is reserved for holding the address of the dynamic structure, which is referenced through the symbol `_DYNAMIC`. The three subsequent entries are reserved for use by the dynamic linker.

The GOT is referenced through the symbol `_GLOBAL_OFFSET_TABLE_`. It can reside anywhere in the section `.got`, and can accept positive and negative indexes given its declaration as:

```
extern Elf32_Addr _GLOBAL_OFFSET_TABLE_ [];
```

NOTE Any change to the format of the GOT will be reflected in `DT_HEXAGON_VER` ([Section 13.2](#)).

13.4 Procedure linkage table

The procedure linkage table (or “PLT”) enables execution transfers from one object to another, with the resolution of function symbols performed at runtime.

Each object that references external functions has a PLT, which is composed of an array of stubs for each of the external functions. The dynamic linker determines the absolute addresses of the destinations and stores them in the GOT, from which they are loaded by the stub code in the PLT entry.

A relocation table is associated with the PLT. The `DT_JMPREL` tag in the `_DYNAMIC` array indicates the location of the first relocation entry. The relocation table entries, after any reserved entry, match the PLT entries in a one-to-one correspondence – for example, relocation table entry 16 applies to the 16th PLT entry after the reserved PLT entries.

The relocation type associated with each non-reserved PLT entry must be `R_HEX_JMP_SLOT`. The relocation offset must specify the address of a GOT entry containing the address of the function, and the symbol table index must reference the appropriate symbol.

The resolution of external function symbols is beyond the scope of this specification; however, the PLT is designed to allow on-demand resolution of such symbols. In other words, the external function symbols need not be resolved before execution begins: each symbol is resolved only when it is called (a policy known as *lazy binding*, which avoids the overhead of resolving and relocating symbols that are never called).

The interface between the PLT and dynamic linker is performed via registers R14 and R15, which respectively contain the PLT relocation table entry index and the object identification. Additionally, R28 is available for use as a scratch register by the PLT stub code.

The details of the PLT are implementation-defined. For example, the following code example presents a possible PLT design, with the first four entries reserved for interfacing with the dynamic linker to perform lazy binding.

```
.PLT0:                                // entry reserved for dyn linker
r15 = pc                             // address of .PLT0
r28.h = #hi (.PLT0@GOT)              // offset of .PLT0 from GOT

r28.l = #lo (.PLT0@GOT)
r28 = sub (r15, r28)                 // address of GOT
r15 = memw (r28 + #8)                // object ID at GOT [2]
r14 = sub (r14, r28)                 // offset of @GOT from GOT

.PLT1:                                // entry reserved for dyn linker
r14 = add (r14, #-16)                // subtract reserved PLT entries
r14 = asr (r14, #2)                  // index of PLT relocation
r28 = memw (r28 + #4)                // dynamic linker at GOT [1]
jumpr r28                            // call dynamic linker
nop
nop

.PLT2:                                // entry reserved for dyn linker
...

.PLT3:                                // entry reserved for dyn linker
...

.PLT4:                                // entry for name1
r15 = pc                             // address of .PLT4
r14.h = #hi (name1's GOT - .PLT4)    // offset of name1 GOT from entry
r14.l = #lo (name1's GOT - .PLT4)
r14 = add (r15, r14)                 // address of name1 GOT
r28 = memw (r14)                     // contents of name1 GOT
jumpr r28                            // call it

.PLT5:                                // entry for name2
r15 = pc
r14.h = #hi (name2's GOT - .PLT5)
r14.l = #lo (name2's GOT - .PLT5)
r14 = add (r15, r14)
r28 = memw (r14)
jumpr r28

.PLT6:
...
```

NOTE For details on the @GOT suffix see [Section 15.2.2](#).

The following steps are assumed in resolving external function symbols:

1. The program loader sets the second entry of the GOT to the address of the dynamic linker (in order to resolve external function symbols), and the third entry of the GOT to some identifying information unique to the object.
2. The linker initializes the first few PLT entries reserved to it with code that will marshal arguments to the dynamic linker. In this example, the code at the beginning of the first PLT entry finds the offset from it to the GOT (as set by the linker); the offset is then used to calculate the absolute address of the GOT. Register R14 is then set to index into the PLT relocation table for function `name1`, and R15 is set to the object ID. The address of the dynamic linker is loaded from a GOT entry and called.
3. When the function `name1` is called, execution is transferred to its assigned PLT entry (in this example, `.PLT4`).
4. The PLT entry contains a stub which sets R15 to its address and then calculates the address of the GOT entry for the symbol `name1`. Control is then transferred indirectly to the contents of this GOT entry, which initially contains the address of the initial PLT entry (`.PLT0`).
5. The preceding PLT code example is executed, starting at label `.PLT0`. With the PLT entry index for `name1` and its corresponding GOT entry offset, the dynamic linker may find the associated relocation entry for the symbol, whose offset points to the GOT entry for the symbol, and whose symbol index points to the function symbol (`name1` in this example).
6. When the dynamic linker resolves the symbol `name1`, it modifies the GOT entry associated with its PLT entry with its actual address, and then transfers control to it. Subsequent calls to `name1` land immediately at its address (after a stop at its PLT entry), but without calling the dynamic linker again.

NOTE The example used in the above procedure uses 24-byte long PLT entries.

Any change to the interface with the dynamic linker will be reflected in `DT_HEXAGON_VER` ([Section 13.2](#)).

14 Thread-Local Storage

14.1 Overview

The Hexagon development tools support thread-local storage (TLS), which is data that is allocated at runtime and accessible only by the thread that allocates it.

This chapter describes how TLS features can be used.

14.2 Thread-local storage section

Both initialized and uninitialized TLS data can be specified. At runtime, when a thread is created the initialized data must be provided. This data is provided in the sections identified with `SHF_TLS` flag:

- `.tdata` – Initialized TLS data
- `.tbss` – Uninitialized data defined as `COMMON` symbols

These combined sections form the TLS template that is used when a thread is created.

Symbols associated with the TLS data are assigned type `STT_TLS`. Their offset is relative to the beginning of the TLS template – this offset is used in the `st_value` field of defined TLS symbols.

The TLS template resides in a segment marked with `PT_TLS`.

14.3 Runtime allocation

The TLS area is allocated in the following situations:

- Whenever a thread is created (including the main executable thread)
- Whenever the TLS area of a shared object is referenced for the first time after an executable was started

When creating the TLS area for a thread (including the main executable thread), the runtime linker combines the TLS templates for all loaded shared objects – including the executable and shared libraries – into a single TLS template. This single TLS template is then used to initialize the TLS area allocated for each new thread.

14.4 Load-time allocation

When a shared object is loaded after process startup, and the object contains a TLS template, then a TLS area may be allocated when the TLS data is first referenced.

When a shared object is unloaded, the memory associated with its TLS template is freed (depending on the access method used – see [Section 14.6](#) for details).

14.5 Interface

The TLS area is accessed at the processor level through the special register UGP ([Section 9.2](#)). This register is set to the address one location above the TLS area, which grows downwards from UGP.

Depending on the TLS access method in use ([Section 14.6](#)), different methods are used to locate the TLS area.

The following (implementation-defined) function is assumed to be available:

```
typedef struct
{
    size_t ti_moduleid;
    ptrdiff_t ti_tloffset;
} TLS_index;

extern void *__tls_get_addr (TLS_index *);
```

This function returns a pointer to the address of the specified TLS data.

14.6 Thread-local storage access

The location of specific data items in an object's TLS area is known at link time, and this information can be used to directly access the data. However, shared objects containing a TLS template may be loaded after process startup; in this case, more general methods are needed for an executable object to access its TLS data.

The following general methods can be used to access TLS data – they are listed here in order from the most general to the most restrictive (with the most restrictive methods also being the most efficient):

- General dynamic (GD)
- Local dynamic (LD)
- Initial executable (IE)
- Local executable (LE)

NOTE Subject to the specified constraints, any of these methods can be used in a given shared object – the choice depends on which method is the most convenient to use in a specific instance.

14.6.1 General dynamic (GD)

In the GD method, TLS data can be referenced from either a shared object or an executable object.

To obtain the address of a variable in the TLS area, the code must call the function `__tls_get_addr ()` with the address of a `TLS_index` structure ([Section 14.5](#)).

NOTE In a shared object, the TLS template can be loaded only after process startup. Otherwise, the (standard ELF) dynamic tag `DT_FLAGS` is set to the value `DF_STATIC_TLS`, which prevents the TLS data from being loaded after process startup.

14.6.2 Local dynamic (LD)

In the LD method, if a shared object contains TLS data that is protected or has local visibility, the location of the data in the TLS template is known at link time, and this location information can be used directly.

To obtain the address of a variable in the TLS area, the code must first obtain the TLS base address by calling the function `__tls_get_addr ()` with the address of a `TLS_index` structure ([Section 14.5](#)) that has a `ti_tlsoffset` member with a value of zero. The TLS variables can be accessed using the LD method through offsets from the obtained base address.

NOTE In a shared object, the TLS template can be loaded only after process startup. Otherwise, the (standard ELF) dynamic tag `DT_FLAGS` is set to the value `DF_STATIC_TLS`, which prevents the TLS data from being loaded after process startup.

14.6.3 Initial executable (IE)

In the IE method, an object can access TLS data only if it is part of the initial TLS template established at link time. When a shared object is loaded after process startup, its TLS data cannot be accessed using this method.

To obtain the address of a variable in the TLS area, the dynamic linker creates a GOT entry to store the variable's offset in the initial TLS area.

NOTE When a shared object is loaded after process startup, its TLS data is not part of the initial TLS template. Therefore, if a shared object tries to use this method to access its TLS data, the access must be rejected.

14.6.4 Local executable (LE)

In the LE method, an object can access the TLS data of executable objects only.

To obtain the address of a variable in the initial TLS area, simply use a static offset from the top of the initial TLS area (since the variable location is known at link time).

NOTE If an object tries to use this method to access the TLS data of a non-executable object, the access must be rejected.

15 Coding Examples

15.1 Overview

This chapter presents examples showing how the following operations can be coded in Hexagon assembly code:

- Function prologs and epilogs
- Function calls (direct and indirect)
- Branches (direct and indirect)
- Data access
- Thread-local storage

15.2 Examples

Most of the examples in this chapter include separate coding solutions for two cases:

- Absolute code (which must be loaded at a specific address)
- Position-independent code (which can be loaded at an arbitrary address)

Position-independent code is used in shared objects, which typically execute in the context of another object. Special attention should be paid to the differences between the coding solutions for absolute and position-independent code.

NOTE The examples in this chapter are not intended as definitive coding solutions – they merely demonstrate how certain basic operations can be implemented. Optimal coding solutions can be derived from the examples.

15.2.1 Register use

The examples in this chapter use the following register conventions:

- The absolute address of the GOT ([Section 13.3](#)) is stored in R24
- The absolute address of the application TLS area ([Section 14.6](#)) is stored in R25

Both R24 and R25 are initialized in the function prolog ([Section 15.3](#)).

NOTE These register choices are arbitrary, and do not imply any standard coding convention for Hexagon assembly code.

15.2.2 Assembler symbols

[Table 15-1](#) lists special assembler symbols which are used in the examples.

Table 15-1 Assembler symbols

Name	Description
<code>_GLOBAL_OFFSET_TABLE_</code>	Reference symbol in the global offset table (Section 13.3). Implicit address of the zero'th (index = 0) GOT entry.
<code>name@PCREL</code>	Offset to position of the symbol <code>name</code> from the current location.
<code>name@GOT</code>	Offset of GOT entry for the symbol <code>name</code> from <code>_GLOBAL_OFFSET_TABLE_</code> . Instructs linker to create GOT entry for the symbol <code>name</code> .
<code>name@GOTREL</code>	Offset to location of the symbol <code>name</code> from <code>_GLOBAL_OFFSET_TABLE_</code> .
<code>name@PLT</code>	Offset to PLT entry for the symbol <code>name</code> from the current code location. Instructs linker to create a PLT entry for the symbol <code>name</code> .

Table 15-1 Assembler symbols (continued)

Name	Description
name@GDPLT	Offset to PLT entry for the function that returns the address of a TLS symbol <code>name</code> from the current code location, using the GD method (Section 14.6.1). This function takes the address of a <code>TLS_index</code> structure as its argument (Section 14.5).
name@GDGOT	Offset of first GOT entry for the <code>TLS_index</code> structure for the symbol <code>name</code> from <code>_GLOBAL_OFFSET_TABLE_</code> , using the GD method (Section 14.6.1). Instructs linker to create GOT entries for <code>TLS_index</code> structure members <code>ti_moduleid</code> and <code>ti_tlsoffset</code> .
name@LDPLT	Offset to PLT entry for the function that returns the address of a TLS symbol <code>name</code> from the current code location, using the LD method (Section 14.6.2). This function takes the address of a <code>TLS_index</code> structure as its argument (Section 14.5).
name@LDGOT	Offset of first GOT entry for the <code>TLS_index</code> structure for the symbol <code>name</code> from <code>_GLOBAL_OFFSET_TABLE_</code> , using the LD method (Section 14.6.2). Instructs linker to create GOT entries for <code>TLS_index</code> structure members <code>ti_moduleid</code> and <code>ti_tlsoffset</code> , with the latter zeroed.
name@DTPREL	Offset to position of the symbol <code>name</code> from base of TLS template, using the LD method (Section 14.6.2).
name@IE	Address of the GOT entry for the offset into the TLS area for the symbol <code>name</code> , using the IE method (Section 14.6.2). Instructs linker to create a GOT entry for the offset into the TLS area for the symbol <code>name</code> .
name@IEGOT	Offset of the GOT entry for the offset into the TLS area for the symbol <code>name</code> from <code>_GLOBAL_OFFSET_TABLE_</code> , using the IE method (Section 14.6.3). Instructs linker to create a GOT entry for the offset into the TLS area for the symbol <code>name</code> .
name@TPREL	Offset to position of the symbol <code>name</code> from base of the TLS area, using the LE method (Section 14.6.4).

15.2.3 Addressing constraints

The Hexagon processor architecture has certain addressing constraints which affect how data accesses are coded in Hexagon assembly language. The constraints depend on the instruction and data size:

- Some branching instructions use a 7- to 22-bit operand as a signed offset from the current program-counter (PC).
- When accessing data in the small data area (using GP-relative addressing), a 16-bit value is used as an unsigned offset from the GP register. This value is shifted left before it is added to the GP register – the shift amount is determined by the size of the accessed data:
 - 0 for bytes
 - 1 for half-words
 - 2 for words
 - 3 for double-words

This limits the range that GP-relative addressing can access: 64K for bytes, 128K for half-words, 256K for words and 512K for double-words.

- When indexing an entry in the GOT ([Section 13.3](#)), instead of using a larger and slower 32-bit index, a program can restrict itself to just 16 bits for indexing a GOT entry. If the GOT size grows larger than 64KB, then 32-bit indexes must be used.

15.3 Function prolog and epilog

Function *prologs* and *epilogs* are standard code sections which appear at the beginning and end of a function body. Their purpose is to manage the use of certain Hexagon processor registers within the function ([Section 15.2.1](#)).

This section presents examples showing how function prologs and epilogs can be coded in Hexagon assembly language.

15.3.1 Absolute

If a function calls another function, it is defined as a *non-leaf* function. In this case, in order to save the link register, a stack frame must be allocated even if no local variables are allocated from the stack.

If a function does not call another function, no stack frame is necessary (though it may be useful for traversing previous stack frames during debugging).

Table 15-2 Function prolog and epilog (absolute)

C	Hexagon
<pre>void foo (void) { ... }</pre>	<pre>foo: allocframe (#...) ... dealloc_return</pre>

When using the GD or LD method for accessing TLS ([Section 14.6](#)), the absolute address of the GOT ([Section 13.3](#)) can be stored in a register for convenience.

Table 15-3 Function prolog and epilog (absolute – GD or LD access)

C	Hexagon
<pre>void foo (void) { ... }</pre>	<pre>foo: allocframe (#...) r24.h = #HI (_GLOBAL_OFFSET_TABLE_) r24.l = #LO (_GLOBAL_OFFSET_TABLE_) ... dealloc_return</pre>

When using the IE or LE method for accessing TLS, the UGP register ([Section 9.2](#)) can be copied to a general-purpose register for convenience.

Table 15-4 Function prolog and epilog (absolute – IE or LE access)

C	Hexagon
<pre>void foo (void) { ... }</pre>	<pre>foo: allocframe (#...) r25 = ugp ... dealloc_return</pre>

15.3.2 Position-independent

Because each object has a dedicated small data area ([Section 9.2](#)), a global function in a shared object must reset the GP register to point to its own small data area before the function can access the area.

Position-independent code is required to save the previous GP register before resetting it.

```
.typefoo, @function

foo:
allocframe (#...)                // prolog
r14 = gp
memw (sp) = r14                  // save previous GP
l:
r24 = add (pc, ##_GLOBAL_OFFSET_TABLE_@PCREL) // calc GOT addr
...
deallocframe
jumpr lr                        // return

.size foo, . - foo
```

If a global function does not have access to a small data area, the GP register does not need to be reset.

```
.typefoo, @function

foo:
allocframe (#...)                // prolog
l:
r24 = add (pc, ##_GLOBAL_OFFSET_TABLE_@PCREL) // calc GOT addr
r25 = ugp                        // get TLS area addr
...
dealloc_return                  // return

.sizefoo, . - foo
```

NOTE This code example assumes that the absolute address of the GOT is in R24, and the absolute address of the TLS area in R25. This register choice is arbitrary – it does not imply any standard coding convention.

15.4 Direct function call

This section presents examples showing how direct function calls can be coded in Hexagon assembly language.

15.4.1 Absolute

Direct function calls are performed with the CALL instruction, which uses the PC-relative addressing mode to specify the function address.

Table 15-5 Direct function call (absolute)

C	Hexagon	Relocation
extern void foo (void *); foo (NULL);	.extern foo r0 = #0 call foo	R_HEX_B22_PCREL

15.4.2 Position-independent

In external function calls, the function address is not known at build time. In this case the address must be stored in a specific data structure at runtime ([Chapter 13](#)).

Assuming that the relevant data structures have been initialized, external functions are normally called through stubs contained in the procedure linkage table.

Table 15-6 Direct function call (position-independent)

C	Hexagon	Relocation
extern void foo (void *); foo (NULL);	.extern foo r0 = #0 call foo@PLT ^a	R_HEX_PLT_B22_PCREL

^a The “@PLT” suffix is optional in this symbol reference.

NOTE Local function calls use the same code sequence for both absolute and position-independent code.

15.5 Indirect function call

This section presents examples showing how indirect function calls can be coded in Hexagon assembly language.

15.5.1 Absolute

Indirect function calls are performed with the CALLR instruction.

Table 15-7 Indirect function call (absolute)

C	Hexagon	Relocation
extern void foo (void *); static void bar (void); auto void (*fun) (void *); fun = foo; fun (bar);	.extern foo .local bar r20 = const32 (#foo) r0 = const32 (#bar) callr r20	 R_HEX_GPREL16_2 + R_HEX_32 ^a R_HEX_GPREL16_2 + R_HEX_32 ^a

^a Relocation created in small data area.

If the calling function does not have access to a small data area ([Section 9.2](#)), the function addresses must be loaded from regular memory, which requires a less-efficient code sequence.

Table 15-8 Indirect function call (absolute – no small data)

C	Hexagon	Relocation
extern void foo (void *); static void bar (void); auto void (*fun) (void *); fun = foo; fun (bar);	.extern foo .local bar r20 = ##foo r0 = ##bar callr r20	 R_HEX_HI16 R_HEX_LO16 R_HEX_HI16 R_HEX_LO16

15.5.2 Position-independent

In position-independent code, indirect function calls are also performed with the CALLR instruction. However, it is necessary to calculate at runtime the absolute address of the function.

Table 15-9 Indirect function call (position-independent)

C	Hexagon	Relocation
<pre>extern void foo (void *); static void bar (void); auto void (*fun) (void *); fun = foo; fun (bar);</pre>	<pre>.extern foo .local bar: r20 = memw (r24 + ##foo@GOT) ^a r0 = add (pc, ##bar@PCREL) callr r20</pre>	<pre>R_HEX_GOT_32_6_X ^b + R_HEX_GOT_11_X R_HEX_B32_PCREL_X ^b + R_HEX_6_PCREL_X</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

If the caller function does not have access to a small data area ([Section 9.2](#)), the callee function address must be loaded from regular memory, which requires a less-efficient code sequence.

Table 15-10 Indirect function call (position-independent – no small data)

C	Hexagon	Relocation
<pre>extern void foo (void *); static void bar (void); auto void (*fun) (void *); fun = foo; fun (bar);</pre>	<pre>.extern foo .local bar r20.h = #hi (foo@GOT) r20.l = #lo (foo@GOT) r20 = add (r20, r24) r20 = memw (r20) r0.h = #hi (bar@GOTREL) r0.l = #lo (bar@GOTREL) r0 = add (r0, r24) callr r20</pre>	<pre>R_HEX_GOT_HI16 R_HEX_GOT_LO16 R_HEX_GOTREL_HI16 R_HEX_GOTREL_LO16</pre>

However, if the GOT ([Section 13.3](#)) is known to be smaller than 64 KB, then a missing small data area requires fewer loads from memory because the callee function address value fits in 16 bits.

Table 15-11 Ind. function call (position-independent – no small data, GOT < 64KB)

C	Hexagon	Relocation
extern void foo (void *); static void bar (void); auto void (*fun) (void *); fun = foo; fun (bar);	.extern foo .local bar r20 = add (r24, #foo@GOT) r20 = memw (r20) r0.h = #hi (bar@GOTREL) r0.l = #lo (bar@GOTREL) r0 = add (r0, r24) callr r20	R_HEX_GOT_16 R_HEX_GOTREL_HI16 R_HEX_GOTREL_LO16

NOTE The absolute address of the GOT is assumed to be stored in R24 ([Section 15.2.1](#)).

Because the address of the function is referenced as a data object, it is resolved along with the other data object when an object file is loaded, effectively bypassing the lazy-binding mechanism ([Section 13.4](#)).

15.6 Direct branch

This section presents an example showing how direct branches can be coded in Hexagon assembly language.

15.6.1 Absolute & position-independent

Direct branches are performed with the JUMP instruction, which uses the PC-relative addressing mode to specify the branch address.

Table 15-12 Direct branch (absolute & position-independent)

C	Hexagon	Relocation
bar: goto bar;	.lbar: jump .lbar	R_HEX_B22_PCREL

15.7 Indirect branch

This section presents examples showing how indirect branches can be coded in Hexagon assembly language.

15.7.1 Absolute

Indirect branches are performed with the JUMPR instruction. For instance, the following example uses an indirect branch to implement a jump table.

```

r22 = const32 (#.Lbar)
r23 = addasl (r22, r21, #2)
r23 = memw (r23)
jumpr r23
...
.L1:
...
.L2:
...

.section .rodata, "a", @progbits

.Lbar:
.word .L1
.word .L2
...

```

If the code performing the branch does not have access to a small data area ([Section 9.2](#)), the branch address must be loaded from regular memory, which requires a less-efficient code sequence.

```

r22.h = #hi (.Lbar)
r22.l = #lo (.Lbar)
r23 = addasl (r22, r21, #2)
r23 = memw (r23)
jumpr r23
...
.L1:
...
.L2:
...

.section .rodata, "a", @progbits

.Lbar:
.word .L1
.word .L2
...

```

15.7.2 Position-independent

In position-independent code, indirect branches are also performed with the JUMPR instruction. However, it is necessary to calculate at runtime the absolute address of the branch destination.

For instance, the following example implements a position-independent jump table.

```

r22 = const32 (#.Lbar@GOTREL)
r22 = add (r22, r24)
r23 = addasl (r22, r21, #2)
r23 = memw (r23)
r23 = add (r23, r24)
jumpr r23
...
.L1:
...
.L2:
...

.section .rodata, "a", @progbits

.Lbar:
.word .L1@GOTREL
.word .L2@GOTREL
...

```

If the code performing the branch does not have access to a small data area ([Section 9.2](#)), the branch address must be loaded from regular memory, which requires a less-efficient code sequence.

```

r22.h = #hi (.Lbar@GOTREL)
r22.l = #lo (.Lbar@GOTREL)
r22 = add (r22, r24)
r23 = addasl (r22, r21, #2)
r23 = memw (r23)
r23 = add (r23, r24)
jumpr r23
...
.L1:
...
.L2:
...

.section .rodata, "a", @progbits

.Lbar:
.word .L1@GOTREL
.word .L2@GOTREL
...

```

NOTE The absolute address of the GOT is assumed to be stored in R24 ([Section 15.2.1](#)).

15.8 Data access

This section presents examples showing how data accesses can be coded in Hexagon assembly language.

15.8.1 Absolute

Data accesses are performed using the Hexagon load and store instructions.

Table 15-13 Data access (absolute)

C	Hexagon	Relocation
extern int src; static int dst; auto int *ptr; ptr = &dst; *ptr = src;	.extern src .lcomm dst, 4 r23 = const32 (#dst) r22 = memw (#src) memw (r23) = r22	R_HEX_GPREL16_2 + R_HEX_32 ^a R_HEX_GPREL16_2

^a Relocation created in small data area.

If the code performing the data access does not have access to a small data area ([Section 9.2](#)), the data address must be loaded from regular memory, which requires a less-efficient code sequence.

Table 15-14 Data access (absolute – no small data)

C	Hexagon	Relocation
extern int src; static int dst; auto int *ptr; ptr = &dst; *ptr = src;	.extern src .lcomm dst, 4 r23.h = #hi (dst) r23.l = #lo (dst) r21.h = #hi (src) r21.l = #lo (src) r22 = memw (r21) memw (r23) = r22	R_HEX_HI16 R_HEX_LO16 R_HEX_HI16 R_HEX_LO16

However, if the GOT ([Section 13.3](#)) is known to be smaller than 64 KB, then a missing small data area requires fewer loads from memory because the data address value fits in 16 bits.

Table 15-17 Data access (position-independent – no small data, GOT < 64KB)

C	Hexagon	Relocation
extern int src; static int dst; auto int *ptr; ptr = &dst; *ptr = src;	.extern src .lcomm dst, 4 r23.h = #hi (dst@GOTREL) r23.l = #lo (dst@GOTREL) r23 = add (r23, r24) r21 = add (r24, #src@GOT) r22 = memw (r21) r22 = memw (r22) memw (r23) = r22	R_HEX_GOTREL_HI16 R_HEX_GOTREL_LO16 R_HEX_GOT_16

NOTE The absolute address of the GOT is assumed to be stored in R24 ([Section 15.2.1](#)).

15.9 Thread-local storage

This section presents examples showing how data accesses to thread-local storage (TLS) can be coded in version-specific Hexagon assembly language.

15.9.1 Absolute

This section presents code examples using the following TLS access methods:

- GD
- LD
- IE
- LE

GD TLS data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table 15-18 Thread-local storage (V4 absolute – GD access)

C	Hexagon	Relocation
<pre>extern __thread int src; extern __thread int dst; auto int *ptr; ptr = &dst; *ptr = src;</pre>	<pre>.extern src .type src, @tls_object .section .tbss, "awT", @nobits .global dst .type dst, @tls_object dst: .skip 4 .text r0 = add (r24, ##dst@GDGOT) ^a call dst@GDPLT r23 = r0 r0 = add (r24, ##src@GDGOT) ^a call src@GDPLT r22 = memw (r0) memw (r23) = r22</pre>	<pre>R_HEX_GD_GOT_32_6_X ^b + R_HEX_GD_GOT_16_X + R_HEX_DTPMOD_32 ^c + R_HEX_DTPREL_32 ^c R_HEX_GD_PLT_B22_PCREL R_HEX_GD_GOT_32_6_X ^b + R_HEX_GD_GOT_16_X + R_HEX_DTPMOD_32 ^c + R_HEX_DTPREL_32 ^c R_HEX_GD_PLT_B22_PCREL</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

^c Relocation created in GOT.

LD TLS data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table 15-19 Thread-local storage (V4 absolute – LD access)

C	Hexagon	Relocation
<pre>static __thread int src; static __thread int dst; auto int *ptr; ptr = &dst; *ptr = src;</pre>	<pre>.section .tbss, "awT", @nobits .type src, @tls_object src: .skip 4 .type dst, @tls_object dst: .skip 4 .text r0 = add (r24, ##dst@LDGOT) ^a call dst@LDPLT r23 = add (r0, ##dst@DTPREL) ^d r22 = add (r0, ##src@DTPREL) ^d memw (r23) = r22</pre>	<pre>R_HEX_LD_GOT_32_6_X ^b + R_HEX_LD_GOT_16_X + R_HEX_DTPMOD_32 ^c R_HEX_LD_PLT_B22_PCREL R_HEX_DTPREL_32_6_X ^b + R_HEX_DTPREL_16_X R_HEX_DTPREL_32_6_X ^b + R_HEX_DTPREL_16_X</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

^c Relocation created in GOT.

^d If TLS template is known to be smaller than 64 KB, the immediate extension can be omitted with consequent changes to the appropriate relocation type.

IE TLS data accesses in the V4 processor use indirect-with-register-offset addressing mode and 32-bit constant extenders to make the code sequence more efficient.

Table 15-20 Thread-local storage (V4 absolute – IE access)

C	Hexagon	Relocation
extern __thread int src;	.extern src	
	.type src, @tls_object	
extern __thread int dst;	.section .tbss, "awT", @nobits	
	.global dst	
	.type dst, @tls_object	
	dst: .skip 4	
auto int *ptr;		
ptr = &dst;	.text r23 = memw (##dst@IE)	R_HEX_IE_32_6_X ^a + R_HEX_IE_16_X + R_HEX_TPREL_32 ^b
*ptr = src;	r22 = memw (##src@IE)	R_HEX_IE_32_6_X ^a + R_HEX_IE_16_X + R_HEX_TPREL_32 ^b
	r22 = memw (r25 + r22 << #1) memw (r25 + r23 << #1) = r22	

^a Relocation created for immediate constant extender.

^b Relocation created in GOT.

LE TLS data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table 15-21 Thread-local storage (V4 absolute – LE access)

C	Hexagon	Relocation
static __thread int src;	.extern src	
	.type src, @tls_object	
static __thread int dst;	.section .tbss, "awT", @nobits	
	.type dst, @tls_object	
	dst: .skip 4	
auto int *ptr;		
ptr = &dst;	.text r23 = add (r25, ##dst@TPREL) ^a	R_HEX_TPREL_32_6_X ^b + R_HEX_TPREL_16_X
*ptr = src;	r22 = memw (r25 + ##src@TPREL)	R_HEX_TPREL_32_6_X ^b + R_HEX_TPREL_11_X
	memw (r23) = r22	

^a If TLS area is known to be smaller than 64 KB, the immediate extension can be omitted with consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

15.9.2 Position-independent

This section presents code examples using the following TLS access methods:

■ IE

Position-independent IE TLS data accesses in the V4 processor use indirect-with-register-offset addressing mode and 32-bit constant extenders to make the code sequence more efficient.

Table 15-22 Thread-local storage (V4 position-independent – IE access)

C	Hexagon	Relocation
extern __thread int src;	.extern src	
	.type src, @tls_object	
extern __thread int dst;	.section .tbss, "awT", @nobits	
	.global dst	
	.type dst, @tls_object	
	dst: .skip 4	
auto int *ptr;		
ptr = &dst;	.text	
	r23 = memw (r24 + ##dst@IEGOT) ^a	R_HEX_IE_GOT_32_6_X ^b + R_HEX_IE_GOT_11_X + R_HEX_TPREL_32 ^c
*ptr = src;	r22 = memw (r24 + ##src@IEGOT) ^a	R_HEX_IE_GOT_32_6_X ^b + R_HEX_IE_GOT_11_X + R_HEX_TPREL_32 ^b
	r22 = memw (r25 + r22 << #1)	
	memw (r25 + r23 << #1) = r22	

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

^c Relocation created in GOT.

A Version-Specific Coding Examples

A.1 Overview

This appendix presents examples showing how the following operations can be coded in Hexagon assembly language that is processor-version-specific:

- Function prologs and epilogs
- Function calls (indirect only)
- Data access
- Thread-local storage

The version-specific examples offer more efficient coding solutions than the ones presented in [Chapter 15](#), by making use of architectural features that are available in more recent processor versions.

Examples are selectively provided for absolute code (which is loaded at a specific address) and position-independent code (which can be loaded at an arbitrary address).

NOTE The examples in this chapter are not intended as definitive coding solutions – they merely demonstrate how certain basic operations can be implemented. Optimal coding solutions can be derived from the examples.

The examples in this appendix are specific to the Hexagon V4 processor. However, they can be used without change on V5 or V55.

A.2 Function prolog and epilog

This section presents an example showing how function prologs and epilogs ([Section 15.3](#)) can be coded in version-specific Hexagon assembly language.

A.2.1 Position-independent

Because the Hexagon V4 processor supports more powerful addressing modes than the preceding processor versions, it is not always necessary to set the GP register as part of a function call ([Section 15.3.2](#)). The following code example sets up pointers for both the GOT and TLS area.

```
.typefoo, @function

foo:
allocframe (#...)                               // prolog
1:
r24 = add (pc, ##_GLOBAL_OFFSET_TABLE_@PCREL) // calc GOT addr
r25 = ugp                                       // get TLS area addr
...
dealloc_return                                // return

.sizefoo, . - foo
```

NOTE The code examples in this appendix assume that the absolute address of the GOT is in R24, and the absolute address of the TLS area in R25. This register choice is arbitrary – it does not imply any standard coding convention.

A.3 Indirect function call

This section presents an example showing how indirect function calls can be coded in version-specific Hexagon assembly language.

A.3.1 Position-independent

Indirect function calls in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table A-1 Indirect function call (V4 position-independent)

C	Hexagon	Relocation
<pre>extern void foo (void *); static void bar (void); auto void (*fun) (void *); fun = foo; fun (bar);</pre>	<pre>.extern foo .local bar r20 = memw (r24 + ##foo@GOT) ^a r0 = add (pc, ##bar@PCREL) callr r20</pre>	<pre>R_HEX_GOT_32_6_X ^b + R_HEX_GOT_11_X R_HEX_B32_PCREL_X ^b + R_HEX_6_PCREL_X</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

A.4 Data access

This section presents an example showing how data accesses can be coded in version-specific Hexagon assembly language.

A.4.1 Position-independent

Data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table A-2 Data access (V4 position-independent)

C	Hexagon	Relocation
extern int src; static int dst; auto int *ptr; ptr = &dst; *ptr = src;	.extern src .lcomm dst, 4 r23 = add (pc, ##dst@PCREL) r22 = memw (r24 + ##src@GOT) ^b r22 = memw (r22) memw (r23) = r22	R_HEX_B32_PCREL_X ^a + R_HEX_6_PCREL_X R_HEX_GOT_32_6_X ^a + R_HEX_GOT_11_X

^a Relocation created for immediate constant extender.

^b If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

A.5 Thread-local storage

This section presents examples showing how data accesses to thread-local storage (TLS) can be coded in version-specific Hexagon assembly language.

A.5.1 Absolute

This section presents code examples using the following TLS access methods:

- GD
- LD
- IE
- LE

GD TLS data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table A-3 Thread-local storage (V4 absolute – GD access)

C	Hexagon	Relocation
<pre>extern __thread int src; extern __thread int dst; auto int *ptr; ptr = &dst; *ptr = src;</pre>	<pre>.extern src .type src, @tls_object .section .tbss, "awT", @nobits .global dst .type dst, @tls_object dst: .skip 4 .text r0 = add (r24, ##dst@GDGOT) ^a call dst@GDPLT r23 = r0 r0 = add (r24, ##src@GDGOT) ^a call src@GDPLT r22 = memw (r0) memw (r23) = r22</pre>	<pre>R_HEX_GD_GOT_32_6_X ^b + R_HEX_GD_GOT_16_X + R_HEX_DTPMOD_32 ^c + R_HEX_DTPREL_32 ^c R_HEX_GD_PLT_B22_PCREL R_HEX_GD_GOT_32_6_X ^b + R_HEX_GD_GOT_16_X + R_HEX_DTPMOD_32 ^c + R_HEX_DTPREL_32 ^c R_HEX_GD_PLT_B22_PCREL</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

^c Relocation created in GOT.

LD TLS data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table A-4 Thread-local storage (V4 absolute – LD access)

C	Hexagon	Relocation
<pre>static __thread int src; static __thread int dst; auto int *ptr; ptr = &dst; *ptr = src;</pre>	<pre>.section .tbss, "awT", @nobits .type src, @tls_object src: .skip 4 .type dst, @tls_object dst: .skip 4 .text r0 = add (r24, ##dst@LDGOT) ^a call dst@LDPLT r23 = add (r0, ##dst@DTPREL) ^d r22 = add (r0, ##src@DTPREL) ^d memw (r23) = r22</pre>	<pre>R_HEX_LD_GOT_32_6_X ^b + R_HEX_LD_GOT_16_X + R_HEX_DTPMOD_32 ^c R_HEX_LD_PLT_B22_PCREL R_HEX_DTPREL_32_6_X ^b + R_HEX_DTPREL_16_X R_HEX_DTPREL_32_6_X ^b + R_HEX_DTPREL_16_X</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with the consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

^c Relocation created in GOT.

^d If TLS template is known to be smaller than 64 KB, the immediate extension can be omitted with consequent changes to the appropriate relocation type.

IE TLS data accesses in the V4 processor use indirect-with-register-offset addressing mode and 32-bit constant extenders to make the code sequence more efficient.

Table A-5 Thread-local storage (V4 absolute – IE access)

C	Hexagon	Relocation
extern __thread int src;	.extern src	
	.type src, @tls_object	
extern __thread int dst;	.section .tbss, "awT", @nobits	
	.global dst	
	.type dst, @tls_object	
	dst: .skip 4	
auto int *ptr;		
ptr = &dst;	.text r23 = memw (##dst@IE)	R_HEX_IE_32_6_X ^a + R_HEX_IE_16_X + R_HEX_TPREL_32 ^b
*ptr = src;	r22 = memw (##src@IE)	R_HEX_IE_32_6_X ^a + R_HEX_IE_16_X + R_HEX_TPREL_32 ^b
	r22 = memw (r25 + r22 << #1) memw (r25 + r23 << #1) = r22	

^a Relocation created for immediate constant extender.

^b Relocation created in GOT.

LE TLS data accesses in the V4 processor use 32-bit constant extenders to make the code sequence more efficient.

Table A-6 Thread-local storage (V4 absolute – LE access)

C	Hexagon	Relocation
static __thread int src;	.extern src	
	.type src, @tls_object	
static __thread int dst;	.section .tbss, "awT", @nobits	
	.type dst, @tls_object	
	dst: .skip 4	
auto int *ptr;		
ptr = &dst;	.text r23 = add (r25, ##dst@TPREL) ^a	R_HEX_TPREL_32_6_X ^b + R_HEX_TPREL_16_X
*ptr = src;	r22 = memw (r25 + ##src@TPREL)	R_HEX_TPREL_32_6_X ^b + R_HEX_TPREL_11_X
	memw (r23) = r22	

^a If TLS area is known to be smaller than 64 KB, the immediate extension can be omitted with consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

A.5.2 Position-independent

This section presents code examples using the following TLS access methods:

- IE

Position-independent IE TLS data accesses in the V4 processor use indirect-with-register-offset addressing mode and 32-bit constant extenders to make the code sequence more efficient.

Table A-7 Thread-local storage (V4 position-independent – IE access)

C	Hexagon	Relocation
<pre>extern __thread int src; extern __thread int dst; auto int *ptr; ptr = &dst; *ptr = src;</pre>	<pre>.extern src .type src, @tls_object .section .tbss, "awT", @nobits .global dst .type dst, @tls_object dst: .skip 4 .text r23 = memw (r24 + ##dst@IEGOT) ^a r22 = memw (r24 + ##src@IEGOT) ^a r22 = memw (r25 + r22 << #1) memw (r25 + r23 << #1) = r22</pre>	<pre>R_HEX_IE_GOT_32_6_X ^b + R_HEX_IE_GOT_11_X + R_HEX_TPREL_32 ^c R_HEX_IE_GOT_32_6_X ^b + R_HEX_IE_GOT_11_X + R_HEX_TPREL_32 ^b</pre>

^a If GOT is known to be smaller than 64 KB, the immediate extension can be omitted with consequent changes to the appropriate relocation type.

^b Relocation created for immediate constant extender.

^c Relocation created in GOT.