Last updated:
2006-12-07 14:33

TASKING C166 ELF/DWARF APPLICATION BINARY INTERFACE

Document ID	119-EDABI
Status	Released
Version	1.2
Date	2006-12-07

1 Revision History

- v1.0: Initial version
- v1.1: Version made available with the last v1.0 beta of the TASKING VX-toolset for C166. Switched formally to DWARF 3.0
- v1.2: Added SHF_TASKING_PROTECTED. Changed values of EF_C166_DATA_*. Added return_address_register in DWARF information. Updated call stack frame section.

2 Introduction

This document describes the implementation of the ELF object format and the DWARF 3 debug information for the TASKING VX-toolset for C166. The implementation is based on:

- System V Application Binary Interface - DRAFT - 17 December 2003
  see http://www.caldera.com/developers/gabi/2003-12-17/contents.html
- DWARF Debugging Information Format, Version 3, December 20, 2005
  see http://dwarf.freestandards.org

3 ELF Implementation

3.1 ELF Header

The following paragraphs define C166 specific items in the ELF header.

3.1.1 e_ident field

The e_ident field values are defined as follows:

Field	Value	Description
e_ident[EI_CLASS]	ELFCLASS32	Identifies 32 bit architecture.
e_ident[EI_DATA]	ELFDATA2LSB	Identifies 2's complement little endian data encoding.

3.1.2 E_MACHINE

The E_MACHINE is defined as follows:

E_MACHINE	Value	Description
EM_C166	116	Infineon C16x/XC16x processor

3.1.3 E_FLAGS

The E_FLAGS field will be used to distinguish between memory models and extended architectures:

Bit	Type	Values	Meaning
0-3	EF_C166_CORE_UNDEFINED	0	Architecture not defined
	EF_C166_CORE_8X166	1	Classic 8xC166
	EF_C166_CORE_C16X	2	Infineon C16x
	EF_C166_CORE_ST10	3	STMicroelectronics ST10
	EF_C166_CORE_ST10MAC	4	STMicroelectronics ST10 with MAC unit (e.g., ST10x272)
	EF_C166_CORE_XC16X	5	Infineon XC16X
	EF_C166_CORE_SUPER10	6	STMicroelectronics Super10
	EF_C166_CORE_SUPER10M345	7	STMicroelectronics Super10M345 and derivatives
	EF_C166_CORE_C166SV1	8	Infineon C166S V1 core
		9-15	reserved for future use
4-7	EF_C166_DATA_UNDEFINED	0	Data model not defined
	EF_C166_DATA_NEAR	1	Near data model
	EF_C166_DATA_FAR	2	Far data model
	EF_C166_DATA_SHUGE	3	Segmented huge data model
	EF_C166_DATA_HUGE	4	Huge data model
		5-15	reserved for future use
8-10	EF_C166_CODE_UNDEFINED	0	Code model not defined
	EF_C166_CODE_HUGE	1	Code model with huge functions
	EF_C166_CODE_NEAR	2	Code model with near functions
		3-7	reserved for future use
11	EF_C166_SYSTEM_STACK	0	System stack is used as default for return values
	EF_C166_USER_STACK	1	User stack is used as default for return values
12	EF_C166_FLOAT_DOUBLE	0	Double precission floating point is treated as double precission
	EF_C166_FLOAT_NODOUBLE	1	Double precission floating point is treated as single precission
13-31		0	Reserved for future use

3.2 ELF Section Attribute Flags

Section attribute flags are defined in the sh_flags field of the section header record. The TASKING defined flags are in the SHF_MASKOS or the SFR_MASKPROC range:

Name	Value
SHF_MASKOS	0x0FF00000
SHF_MASKPROC	0xF0000000
SHF_TASKING_PROTECTED	0×08000000
SHF_TASKING_ABSOLUTE	0×10000000
SHF_TASKING_SEPARATE	0×20000000
SHF_TASKING_NOCLEAR	0×40000000
SHF_TASKING_PAGED	0×80000000

SHF_TASKING_PROTECTED
Sections with this flag set are protected. Sections with the SHF_TASKING_PROTECTED flag set are excluded from unreferenced section removal and duplicate section removal.

SHF_TASKING_ABSOLUTE
Sections with this flag set are absolute. The sh_addr field in the section header contains the absolute address.

SHF_TASKING_SEPARATE
Sections with the same type, attributes and name are concatenated by the linker. Sections with the SHF_TASKING_SEPARATE flag set will not be concatenated.

SHF_TASKING_NOCLEAR
These sections must have type SHT_NOBITS. Normally, sections of this type must be cleared on startup, but sections with the flag SHF_TASKING_NOCLEAR set should not be cleared.

SHF_TASKING_PAGED
Sections with this flag set are relocatable, the sh_addr field in the section header is interpreted as a page size by the linker. The section must be located within a page of this size. Pages start at a multiple of the page size. If the section name is of the form “name@group”, the linker must place all sections with the same group postfix in the same page. The size of the page depends on the section type and address space.

‘Max sections’
When the SHF_MERGE flag is set in combination with the SHF_TASKING_NOCLEAR flag, all sections with the same name type and flags are combined into a single section, with size equal to the largest input section. This are so-called ‘max sections’.
Note that this only applies to scratch sections.

3.3 Address Spaces

Address space information for sections and symbols that is to be used by the linker is encoded in an additional field that is added to the ELF section headers and symbol table entries. If present, the value for this field must be non-zero for sections that have the SHF_ALLOC flag set. The addional address space fields are only present in relocatable ELF object files. The fields are not present in the absolute ELF file as generated by the linker.

The Section Header definition for relocatable object files:

```
 typedef struct {
        Elf32_Word      sh_name;
        Elf32_Word      sh_type;
        Elf32_Word      sh_flags;
        Elf32_Addr      sh_addr;
        Elf32_Off       sh_offset;
        Elf32_Word      sh_size;
        Elf32_Word      sh_link;
        Elf32_Word      sh_info;
        Elf32_Word      sh_addralign;
        Elf32_Word      sh_entsize;
        unsigned char   sh_addrspace;   // additional address space field
        unsigned char   sh_reserved[3]; // reserved for future use
} Elf32_Shdr;
```

The Symbol Table Entry definition for relocatable object files:

```
typedef struct {
        Elf32_Word      st_name;
```

```
        Elf32_Addr      st_value;
        Elf32_Word      st_size;
        unsigned char   st_info;
        unsigned char   st_other;
        Elf32_Half      st_shndx;
        unsigned char   st_addrspace;   // additional address space field
        unsigned char   st_reserved[3]; // reserved for future use
} Elf32_Sym;
```

The sh_reserved and st_reserved fields are required to pad to a 32 bit boundary.

The following address space values are defined:

Space	Value
bit	1
bita	2
iram	3
near	4
far	5
shuge	6
huge	7
code	8

3.4 Relocation Expression Stack

For those situations in which the relocation value cannot be expressed as a simple symbol value plus an addend, there are three special relocation types (ELF32_R_TYPE) used to evaluate an arbitrary expression on a relocation stack. These relocation types are referred to as extended relocations. Other relocation types are ordinary relocations.

A relocation stack is a standard last-in-first-out data structure containing 32-bit values. A hosted environment must not place any arbitrary limit on the depth of the stack. An embedded environment may impose any limit on stack depth or omit the relocation stack entirely (effectively, a maximum stack depth of zero).

A target supporting the relocation expression stack must define the following relocation types in addition to the target specific relocation types:

Relocation type	Value
R_TASKING_PUSH	253
R_TASKING_OPER	254
R_TASKING_POP	255

R_TASKING_PUSH
This relocation type indicates that the sum of the symbol value (the value of symbol number zero is zero) plus the signed r_addend value should be pushed onto the relocation stack.

R_TASKING_OPER
This relocation type defines an operation to be performed on one or more stack values. The operation is specified by the sum of the symbol value (the value of symbol number zero is zero) plus the signed r_addend value. Operations are shown in Table 8. In the table, Stack 0 indicates the value on the top of the stack, and Stack 1 indicates the value one level beneath the top of the stack.

R_TASKING_POP
Indicates the end of a relocation expression. When the R_TASKING_POP operation is encountered, there should be exactly one value on the stack. This value, which is consumed by this operation, becomes the new relocation value for the ordinary relocation type specified in the R_TASKING_POP relocation. The relocation type is specified by the sum of the symbol value (the value of symbol number zero is zero) plus the signed r_addend value It is the responsibility of the relocation engine to ensure that the stack is empty after a R_TASKING_POP, before an ordinary relocation, and after linking is complete. A sequence of relocations which causes a stack underflow does not conform to this specification.

The following Relocation Stack Operations are defined:

Relocation Value	Stack 0 Before Operation	Stack 1 Before Operation	Stack 0 After Operation	Operation

0	X		X	No operation
1	X		-X	Negation (2s complement)
2	X		~X	Bitwise NOT (1s complement)
3	X		!X	Boolean NOT (zero ->1, nonzero -> 0)
4	Y	X	X * Y	Multiplication
5	Y	X	X / Y	Division
6	Y	X	X % Y	Remainder
7	Y	X	X + Y	Addition
8	Y	X	X - Y	Subtraction
9	Y	X	X <<< Y	Logical shift left
10	Y	X	X >>> Y	Logical shift right
11	Y	X	X << Y	Arithmetic shift left
12	Y	X	X >> Y	Arithmetic shift right
13	Y	X	X < Y	1 if X < Y, otherwise 0
14	Y	X	X <= Y	1 if X <= Y, otherwise 0
15	Y	X	X > Y	1 if X > Y, otherwise 0
16	Y	X	X >= Y	1 if X >= Y, otherwise 0
17	Y	X	X == Y	1 if X equals Y, otherwise 0
18	Y	X	X != Y	1 if X does not equal Y, otherwise 0
19	Y	X	X & Y	Bitwise AND
20	Y	X	X | Y	Bitwise OR
21	Y	X	X ^ Y	Bitwise XOR
22	Y	X	X && Y	1 if X and Y both nonzero, otherwise 0
23	Y	X	X | | Y	1 if X or Y or both nonzero, otherwise 0

Note that in most cases, the stack values are treated as unsigned. However, arithmetic shifts and logical shifts are treated differently.

Logical shift left:
Zeroes are shifted in on the right.

Logical shift right:
Zeroes are shifted in on the left.

Arithmetic shift left:
Zeroes are shifted in on the right, and the most significant bit is always unaffected. Arithmetic shift right: Copies of the most significant bit are shifted in on the left.

4 DWARF Debug Information

The C166 tool chain uses DWARF for passing HLL debug information from the compiler to the debugger.

4.1 DWARF register mapping

DWARF represents register names effectively as small integers. These numbers are used in the OP_REG and OP_BASEREG atoms to locate values. The mapping of DWARF register numbers to the C166 register set is as follows.

Register	Atom	Ranges
Rn	a = n	0 < = n < = 15; 0 < = a < = 15
RLn	a = 16 + n*2	0 < = n < = 7; 16 < = a < = 30 (even)
RHn	a = 17 + n*2	0 < = n < = 7; 17 < = a < = 31 (odd)
Rn.m	a = 32 + n*16 + m	0 < = n < = 15; 0 < = m < = 15; 32 < = a < = 287
USR0	a	a = 288
SP	a	a = 289
MAC	a	a = 290
MAH	a	a = 291
MAL	a	a = 292
MAE	a	a = 293
MRW	a	a = 294
IDX0	a	a = 295
IDX1	a	a = 296
QX0	a	a = 297
QX1	a	a = 298
QR0	a	a = 299
QR1	a	a = 300
CF Info return_address_register	a	a = 301
IP	a	a = 302
CSP	a	a = 303
SPSEG	a	a = 304
DPP0	a	a = 305
DPP1	a	a = 306
DPP2	a	a = 307
DPP3	a	a = 308

Note: the “CF Info return_address_register” register has been specified to prevent the number from being used for a regular register in the future, which could potentially confuse debuggers when reading older objects where the number would have been used for the return_address_register instead of the regular register in the call frame information. The return_address_register is virtual and it is not intended to show up in any DWARF expression.

4.2 Function Attributes

Function attributes describing the combination of memory model, stack model and other calling convention details, are conveyed with additional tool-chain specific values using the DWARF calling convention attribute DW_AT_calling_convention.

4.2.1 DWARF Function Calling Convention

Encoding	Symbolic Value	Meaning
0×01	DW_CC_normal	Huge function model, return address on system stack (default)
0×02	DW_CC_program	Not used (see DWARF 3 specification)
0×03	DW_CC_nocall	Not used (see DWARF 3 specification)

0×65	DW_CC_interrupt	Function is an interrupt handler, return address on system stack
0×66	DW_CC_near_system_stack	Near function model, return address on system stack
0×67	DW_CC_near_user_stack	Near function model, return address on user stack
0×68	DW_CC_huge_user_stack	Huge function model, return address on user stack

4.3 TASKING Type Qualifier Extensions

TASKING Type Qualifiers Extension Encoding

The additional C type qualifiers are specified using the DW_AT_address_class attribute.

Qualifier	Value	Remark
__bita	1	
__near	2	
__far	3	
__shuge	4	
__huge	5	
__code	6	not really used; is implicit for functions

4.4 Call frame information

The following information should be read in conjunction with the definitions in Section 6.4 of the DWARF standard document.

4.4.1 Call Stack and Memory Models

The size and the save area of the return address differ across the various memory models. This has to be reflected by the debug info for the debugger to be able to walk up the stack.

Basic Facts

- Each stack word is 16 bits in size.
- The return address consists basically of CSP:IP. Yet in some memory models only IP will be pushed on the stack. Likewise, the stack where the return address is pushed is also memory model dependent.
- Some derivatives use SPSEG to determine which segment is used for the system stack, while other derivatives omit an SPSEG register altogether.
- R15 is combined with one of four DPP registers. The top 2 bits of R15 select the DPP register, and DPP*i* is shifted 2 bits to the right before combining.
- CSP does not change for the duration of one function.

Pending Issues

- Functions where variable length arrays (VLA) are used, switch to using R8 as the frame pointer in order to access automatic variables, while R15 still acts as SP. However, R15 is changed based on run-time data, when resizing VLAs, which cannot be determined at compile time. Therefore in VLA situations R8 should be used in the CFA calculations.
- Infrequently the C compiler needs to save the PSW register to the system stack for a very short period of time, causing the SP register to change in value. These so-called stack deltas also need to be reflected in the call frame information.

Known Limitations

- When single-stepping individual instructions into a function call in a user-stack model application, the return address is pushed onto the user-stack using multiple instructions. For these instructions no call frame information is

issued, causing call frame information to be insufficient for stack walking or saved register retrieval when halting anywhere in such a push sequence.

Near Functions, Return Address on System-Stack

Saved value	Stack
Return address	SP stack
Local automatic variables	R15 stack
CPU registers	R15 stack

Stack Layout	
+0	IP

Huge Functions, Return Address on System-Stack

Saved value	Stack
Return address	SP stack
Local automatic variables	R15 stack
CPU registers	R15 stack

Stack Layout	
+2	CSP
+0	IP

Near Functions, Return Address on User-Stack

Saved value	Stack
Return address	R15 stack
Local automatic variables	R15 stack
CPU registers	R15 stack

Stack Layout	
+0	IP

Huge functions, Return Address on User-Stack

Saved value	Stack
Return address	R15 stack
Local automatic variables	R15 stack
CPU registers	R15 stack

Stack Layout	
+2	CSP
+0	IP

Interrupt Functions

Saved value	Stack
Return address	SP stack
Local automatic variables	R15 stack
CPU registers	SP stack

Stack Layout	
+4	PSW

+2	CSP
+0	IP

4.4.2 Self-containedness

The compiler generates the call frame information in such a way that no information from sections other than .debug_frame should be required to produce a stack trace. For example, it should not be necessary to look up DW_AT_calling_convention attributes.

4.4.3 Definition of CFA

The canonical frame address (CFA) for an address A belonging to function foo is defined as follows:

- If foo has the __interrupt attribute, the CFA expression associated with A evaluates to the value that the system stack pointer (SP or SP combined with SPSEG) had just before the interrupt occurred. If no system stack manipulations happen in foo itself, the CFA expression will therefore come down to “SP + 6” or “SP + 4” (if no SPSEG), depending on the SGTDIS bit.
- If foo is the _cstart function, the CFA expression evaluates to the initial value of the system stack pointer, i.e. “top of system stack”.
- For all other kinds of functions, the CFA expression at A evaluates to the value that the user stack pointer (R15 combined with the appropriate DPP register) had just before the call (or jump) that led to the invocation of foo. This is always the value *before*
  the return address was pushed on the stack, so in the below example for the user stack model, the CFA expression at _foo is “R15 + 4” (ignoring the DPP registers), i.e. the CFA is equal to the value that R15 had at _71, not at _73.

```
_71: movw r8,#@seg(_74)
movw [-r15],r8
movw r8,#@sof(_74)
movw [-r15],r8
_73: jmp _foo
_74:

_foo: .proc far
mov r4, r2
mov r12, [r15+]
mov r11, [r15+]
atomic #0x3
push r11
push r12
ret
```

It should be emphasized that in general debuggers do not actually have to know with which of the two stacks the CFA is associated in a given function, because it is an abstract concept. However, it may be referenced from a location expression via DW_OP_call_frame_cfa. A related point is that the CFA and the DW_AT_frame_base are often related, but they should not be equated.

4.4.4 Determining stack pointer values

The values of the system (SP) and user (R15) stack pointer registers in higher frames can be determined in exactly the same way as those of other registers. For example, an empty huge function will have a rule

```
DW_CFA_val_expression: reg=289, expr=bregx 289 offset=4
```

which states that the value of SP (289) in this frame’s caller is this frame’s SP value plus 4, i.e. the “stack delta” is 4. The same applies to R15.