Alpha Assembly Language Guide

Randal E. Bryant
Carnegie Mellon University
Randy.Bryant@cs.cmu.edu

February 23, 1998

1. Overview

This document provides an overview of the Alpha instruction set and assembly language programming conventions. More complete documentation is available from Digital Equipment Corporation [1, 2].

1.1. Data Types

The most notable feature of the Alpha is that it is a true 64-bit machine. All integer registers are 64 bits wide. Manipulating 64-bit addresses and 64-bit integers is fully supported. In addition, there is support for 32-bit integers.

For historical reasons, (dating back to the PDP-11, a 16-bit machine), Digital has an idiosyncratic terminology for word sizes. They consider a "word" to be 16-bits. Based on this, they refer to 32-bit quantities as "long words" (the word size of the VAX.) They refer to 64-bit quantities as "quad words." We are mostly interested in long words and quad words.

Table 1 shows the machine representations used for the primitive data types of C. Note that variables declared as int's are stored as long (4-byte) words. If you want an 8-byte number, you need to declare it

C declaration	Alpha Data Type	Size (Bytes)
char	Byte	1
short	Word	2
int	Long Word	4
unsigned	Long Word	4
long int	Quad Word	8
long unsigned	Quad Word	8
char *	Quad Word	8
float	S_Floating	4
double	T_Floating	8

Table 1: Sizes of standard data types

as `long`. All pointers (shown here as `char *`) are stored as 8-byte quad words. Don't confuse the two uses of the word "long" here—Alpha long words are 4 bytes, but C `long int`'s are 8.

Within the machine, all integer registers hold quad words. Long words are converted to quad words by sign extension. That is, when converting from a long word lw to a quad word qw, the high order bit of lw is replicated as the most significant 33 bits of qw.

1.2. Porting C Code to the Alpha

When porting C code originally developed on a 32-bit machine to an Alpha, the difference between the sizes for pointers and `int`'s is a common source of nonportability. Lots of code has been written assuming that you could store pointers in locations declared as `int`'s with no loss of information.

Another source of problems is with integer constants. By default, constants in C are assumed to be `int`'s. If you want to make them long, you need to add the suffix "L". Without that suffix, the number is truncated to 32 bits and then sign extended to 64. Here are some examples illustrating this effect:

```
long int a = 0xFFFFFFFF7FFFFFFF;  /* 0x000000007FFFFFFFL */
long int b = 0xFFFFFFFF7FFFFFFFL; /* 0xFFFFFFFF7FFFFFFFL */
long int c = 0x0000000080000000;  /* 0xFFFFFFFF80000000L */
long int d = 0x0000000080000000L; /* 0x0000000080000000L */
```

Observe that values `b` and `d` are most likely what the programmer intended them to be. Values `a` and `c`, on the other hand are not, because the "L" suffix was omitted. Their high order bits are either all 0's or all 1's depending on bit 31 of the declared constant.

When you want to print out 8-byte integers with `printf`, you need to use the directive `%ld`, rather than the standard `%d`. Similarly for printing in hexidecimal (`%lx`) and unsigned (`%lu`) formats.

2. Instructions

The Alpha instruction set is relatively simple. Arithmetic operations apply only to register data. Explicit load and store operations are needed to move data between memory and registers. Conditional branches can only test the relation between a register and the value zero.

2.1. Arithmetic Operations

Alpha supports integer operations for both 4-byte and 8-byte integers. The 8-byte versions treat the operands as full precision values. The 4-byte versions mimic the behavior one would obtain by executing the operations on a 32-bit machine. That is, they compute a value based on only the low order 4 bytes of the operands to generate a 4-byte value. They then sign extend this value to obtain the 8-byte result.

Table 2 lists the arithmetic instructions having both 4-byte and 8-byte versions (note the suffixes "l" and "q".) Arithmetic operations have three operands: two source and one destination. The destination is

2

Long Word	Quad Word	Description	Computation
addl	addq	Add	c = a + b
s4addl	s4addq	Scaled by 4 Add	c = 4*a + b
s8addl	s8addq	Scaled by 8 Add	c = 8*a + b
subl	subq	Subtract	c = a - b
s4subl	s4subq	Scaled by 4 Subtract	c = 4*a - b
s8subl	s8subq	Scaled by 8 Subtract	c = 4*a - b
mull	mulq	Multiply	c = a * b
divl	divq	Divide	c = a / b
reml	remq	Remainder	c = a % b

Table 2: Arithmetic Operations. Each instruction has a (4-byte) word and a quad (8-byte) word form.

given as the rightmost operand, [in contrast to MIPS where the destination is given as the leftmost operand.] Arithmetic operations can have one of two formats (shown for instruction `addq`):

$$\texttt{addq} \quad R_a, \quad R_b, \quad R_c$$
$$\texttt{addq} \quad R_a, \quad Lit_b, \quad R_c$$

where R_a, R_b, and R_c denote registers, and Lit_b denotes a "literal" constant between 0 and 255. The first two operands denote the operation sources: the first must be from register R_a, the second can either be from a register R_b or a literal value Lit_b. The third operand denotes the destination, which must always be a register R_c.

Table 2 shows the effect of each of these instructions using C notation, with source operands `a` and `b`, and destination operand `c`.

For operations requiring constants that don't fit within the 8 bit limit of the standard operations, it is common to use instructions `lda` (load address) and `ldah` (load address high). These are documented in Section 2.4 describing load and store operations, even though they do not reference memory. Alternatively, constants can be declared as part of the assembly program data and stored in memory. Load instructions can then put these values into registers.

The scaled operations, having prefixes "s4" and "s8" scale the first source value by a factor of 4 or 8. These are commonly used for array indexing.

The following are some examples of typical assembly code using arithmetic operations:

```
# Add $1 and $2 and store in $3
addq $1, $2, $3
# Register $8 points to integer array a
# Register $17 contains index i
# Want to set $9 to &a[i]:
s4addq $17, $8, $9
```

2.2. Comparison Operations

Instruction	Description	Computation
cmpeq	Equality	c = (a == b)
cmple	Less than or equal	c = (a <= b)
cmplt	Less than	c = (a < b)
cmpule	Unsigned less than or equal	c = (ua <= ub)
cmpult	Unsigned less than	c = (ua < ub)

Table 3: Comparison Operations

Instruction	Description	Effect
and	And	c = a & b
bic	Bit Clear	c = a & ~b
bis	Bit Set	c = a | b
eqv	Logical Equivalence	c = ~(a ^ b)
xor	Exclusive-Or	c = a ^ b
ornot	Or-Not	c = a | ~b
sra	Shift Right Arithmetic	c = a >> (b % 64)
sll	Shift Left	c = a << (b % 64)
srl	Shift Right Logical	c = ua >> (b % 64)

Table 4: Bit-Level Operations

All comparisons operations operate on quad (8-byte) words. They have the same format as arithmetic operations. They result in destination register R_c being set to 1 (true) or 0 (false). Table 3 lists the different comparison possibilities. Note that the inequality tests have both signed and unsigned versions. These are indicated with C syntax using operands a and b as signed values and ua and ub as unsigned values.

2.3. Bit-Level and Logical Operations

All bit-level and shift operations operate on quad words. They have the same format as arithmetic instructions. Table 4 lists the different possibilities.

Left shift inserts 0's into the low order bit positions. Logical right shift inserts 0's into the high order bit positions (used for unsigned operands). Arithmetic right shift copies the high order bit of operand a into the new bit positions (used for signed operands).

Note that the shift amount must be between 0 and 63. Any larger number is reduced modulo 64 (by simply masking off all but the low order 6 bits).

2.4. Loads and Stores

Load and store operations are used to transfer data between registers and memory. Separate instructions are used to perform long (4-byte) word accesses and quad (8-byte) word accesses. In addition, instructions lda and ldah have the format of a load operation, but they do not cause any memory references.

4

Instruction	Description	Bytes Accessed	Effective Address	Effect
ldl	Load Long	4	$EA = b + D$	a = *EA
ldq	Load Quad	8	$EA = b + D$	a = *EA
stl	Store Long	4	$EA = b + D$	*EA = a
stq	Store Quad	8	$EA = b + D$	*EA = a
lda	Load Address	0	$EA = b + D$	a = EA
ldah	Load Address High	0	$EA = b + D \cdot 65536$	a = EA

Table 5: Load and Store Operations

Load and store instructions have the following format, shown with instruction ldq (load quad word):

$$\texttt{ldq} \quad \texttt{R}_a, \quad Disp(\texttt{R}_b)$$

Operands $\texttt{R}_a$ and $\texttt{R}_b$ indicate registers, while $Disp$ is a constant *displacement* ranging between -32,768 and +32,767. In most cases $\texttt{R}_a$ indicates the destination (load) or source (store) of the data, while the combination of $\texttt{R}_b$ and $Disp$ indicates the memory location to access. Note that load instructions are the only Alpha instructions for which the destination is written on the left.

Table 5 describes the load and store operations. The column labeled "Effective Address" shows how the contents of register $\texttt{R}_b$, denoted b, and the value of the displacement $Disp$, denoted D are combined to generate an effective address EA. In most cases the values are simply added. For the ldah instruction, the value D is scaled by a factor of $2^{16} = 65536$.

The column labeled "Effect" in Table 5 describes the behavior of the operation in C notation, where the effective address EA is represented by a pointer variable EA, and register $\texttt{R}_a$ is represented by variable a. Load instructions read from the effective address and place the result in register $\texttt{R}_a$. The quad word version ldq reads 8 bytes. The long word version ldl reads only 4 bytes and sign extends them to 8. Conversely, the store operations write the value in register $\texttt{R}_a$ to the memory locations indicated by the effective address. The quad word version stq writes all 8 bytes, while the long word version writes only the low order 4 bytes of $\texttt{R}_a$.

Instructions lda and ldah place the effective address in register $\texttt{R}_a$ without accessing any memory locations. They are useful for setting addresses, for performing pointer arithmetic, and even for performing integer operations involving constants.

```
# Set $1 to absolute address 0x000F0FF0
# Use property that $31 is always 0
ldah $1, 15($31)   # 0xF
lda $1, 4080($1)  # 0x0FF0
# Compute p++ for integer pointer p in register $2
lda $2, 4($2)
# Compute x -= 17 for integer x in register $5
lda $5, -17($5)
```

2.5. Conditional Moves

5

Instruction	Description	Move Condition
cmoveq	Conditional Move on Equal	a == 0
cmovne	Conditional Move on Not Equal	a != 0
cmovgt	Conditional Move on Greater Than	a > 0
cmovge	Conditional Move on Greater Than or Equal	a >= 0
cmovlt	Conditional Move on Less Than	a < 0
cmovle	Conditional Move on Less Than or Equal	a <= 0
cmovlbc	Conditional Move on Lower Bit Clear	!(a & 0x1)
cmovlbs	Conditional Move on Lower Bit Set	a & 0x1

Table 6: Conditional Move Instructions

Form	Description	Actual Implementation
nop	No operation	bis $31, $31, $31
mov $1, $2	Move register	bis $31, $1, $2
mov 17, $2	Move literal	bis $31, 17, $2
sextl $1, $2	Move long word and sign-extend	addl $31, $1, $2

Table 7: Derived Assembly Language Operations.

Conditional move operations provide a means of conditionally updating a register without using any branch operations. In modern machines such as Alpha, this can yield much better performance than the traditional technique of conditionally branching around the updating code.

These instructions have the same format as arithmetic operations, e.g., for instruction cmoveq:

$$\text{cmoveq} \quad R_a, \quad R_b, \quad R_c$$
$$\text{cmoveq} \quad R_a, \quad Lit_b, \quad R_c$$

Register R_a indicates the tested value, either register R_b or Lib_b indicates the source data, and register R_c designates the move destination. Whether or not the move takes place is based on the result of comparing register R_a to 0.

Table 6 lists the different conditional move instructions and the type of comparison performed. C expression syntax is used, where variable a denotes the contents of register R_a. For example, the cmoveq instruction is equivalent to the following C code:

```
if (a == 0)
    c = b;
```

where variable b represents the source data and variable c represents the destination.

2.6. Derived Arithmetic Operations

In an attempt to make assembly language more readable, some commonly used patterns are given special names. These typically involve degenerate cases, such as copying from one register to another. They exploit

Instruction	Description	Branch Condition
beq	Branch on Equal	a == 0
bne	Branch on Not Equal	a != 0
bgt	Branch on Greater Than	a > 0
bge	Branch on Greater Than or Equal	a >= 0
blt	Branch on Less Than	a < 0
ble	Branch on Less Than or Equal	a <= 0
blbc	Branch on Lower Bit Clear	!(a & 0x1)
blbs	Branch on Lower Bit Set	a & 0x1
br	Branch	1
bsr	Branch to Subroutine	1

Table 8: Branch Instructions

the fact that integer register $31 is always 0. Table 7 lists some typical cases and their translations into actual Alpha instructions. No-op instructions nop are commonly used to pad code to meet specified alignment requirements. Move instructions mov are used to transfer from one register to another, or to set a register to a constant value. The sign extension operation sextl is the common method for converting from C int's to long int's. For example, it is common to see instructions of the form sextl $16, $16 to convert an integer argument passed in register $16 into the numerically equivalent quad word.

2.7. Transfers of Control

Transfers of control come in two flavors: branches and jumps. Most branches are conditional—whether or not they are taken depends on the result of comparing an operand register to 0. They have format (shown for beq):

$$\texttt{beq}\quad \text{R}_a\texttt{,}\ \textit{Label}$$

where R_a denotes the register being tested and *Label* is a label designating some position in the assembly code. The assembler automatically translates this label into an offset relative to the program counter. The upper part of Table 8 documents the different conditional branch types and the condition under which they are taken. C syntax is used with variable a denoting the contents of register R_a.

As shown in the lower part of Table 8, two special branch forms: br and bsr branch unconditionally. The unconditional branch br has the form:

$$\texttt{br}\quad \textit{Label}$$

That is, it simply designates the branch target. The branch to subroutine instruction has the same format as other branches, but register argument R_a is used in a totally different way. It designates where the current value of the program counter should be stored to allow the subroutine to return to the calling point. The convention is to use register $26 for this purpose.

Jump instructions provide unconditional transfers of control with the target address specified by a register. We will use three different forms:

7

S_floating	T_floating	Description	Computation
`adds`	`addt`	Add	`c = a + b`
`subs`	`subt`	Subtract	`c = a - b`
`muls`	`mult`	Multiply	`c = a * b`
`divs`	`divt`	Divide	`c = a / b`

Table 9: Floating Point Arithmetic Operations. Each instruction has a single precision (S_floating) and a double precision (T_floating) version.

Instruction	Description	Computation
`cmpteq`	Equality	`c = (a == b) ? 2.0 : 0.0`
`cmptle`	Less than or equal	`c = (a <= b) ? 2.0 : 0.0`
`cmptlt`	Less than	`c = (a < b) ? 2.0 : 0.0`

Table 10: Floating Point Comparison Operations

$$
\begin{array}{lll}
\texttt{jmp} & \texttt{(R}_b\texttt{)} & Hint \\
\texttt{jsr}\quad \texttt{R}_a\texttt{,} & \texttt{(R}_b\texttt{)} & Hint \\
\texttt{ret} & \texttt{(R}_b\texttt{)} & Hint \\
\end{array}
$$

In all cases $Hint$ is optional information inserted by the compiler to help the processor predict the jump target. The exact nature of these hints is not our concern.

The unconditional jump instruction `jmp` designates the jump target address in register R_b. The jump to subroutine instruction `jsr` gives the target in register R_b and the register to store the current program counter as argument R_a. The convention is to use register $26 for this purpose. The return instruction `ret` is functionally equivalent to a jump—it gives the target address as register R_b. By conventional this will be register $26, holding to program counter set by the preceding `bsr` or `jsr`.

2.8. Floating Point

Floating point instructions use a set of 32 floating point registers, named $f0 to $f31$. Four floating point formats are supported, but we are only interested in two: the S_floating format implementing IEEE single precision and the T_floating format implementing IEEE double precision. Each floating point register is 8 bytes, but it can hold either a single precision or a double precision value.

In general, the floating point operations mirror the behavior of a subset of the integer operations. For example, floating point arithmetic operations have just one format (shown for instruction `addt`):

$$\texttt{addt}\quad \texttt{F}_a\texttt{,}\quad \texttt{F}_b\texttt{,}\quad \texttt{F}_c$$

where F_a and F_b indicate the two source registers, and F_c indicates the destination register. Table 9 lists the common arithmetic operations, using C variables `a` and `b` to denote the source operands and `c` to denote the destination.

8

Instruction	Description	Bytes Accessed	Effective Address	Effect
`lds`	Load S_floating	4	$EA = b + D$	`a = *EA`
`ldt`	Load T_floating	8	$EA = b + D$	`a = *EA`
`sts`	Store S_floating	4	$EA = b + D$	`*EA = a`
`stt`	Store T_floating	8	$EA = b + D$	`*EA = a`

Table 11: Floating Point Load and Store Operations

Instruction	Description	Move Condition
`fcmoveq`	Conditional Move on Equal	`a == 0.0`
`fcmovne`	Conditional Move on Not Equal	`a != 0.0`
`fcmovgt`	Conditional Move on Greater Than	`a > 0.0`
`fcmovge`	Conditional Move on Greater Than or Equal	`a >= 0.0`
`fcmovlt`	Conditional Move on Less Than	`a < 0.0`
`fcmovle`	Conditional Move on Less Than or Equal	`a <= 0.0`

Table 12: Conditional Move Instructions

Table 10 lists some of the floating point comparison operations. These set the destination register F_c to 2.0 if the comparison holds and to 0.0 if it does not.

Floating point load and store instructions have the same format as their integer counterparts, except that they use floating point registers for data. For instruction `lds`, the format is:

$$\text{lds} \quad F_a, \quad Disp(R_b)$$

Operands F_a and R_b indicate registers, while $Disp$ is a constant *displacement* ranging between -32,768 and +32,767. Floating point register F_a indicates the destination (load) or source (store) of the data, while the combination of R_b and $Disp$ indicates the memory location to access. Table 11 lists the different instructions, their effective address calculations (matching the calculations for integer loads and stores), and the instruction effect.

As indicated in Table 12, there are also conditional moves for floating point values. These have the same format as arithmetic operations (shown here for `fcmoveq`):

$$\text{fcmoveq} \quad F_a, \quad F_b, \quad F_c$$

The value in register F_b is conditionally copied to register F_c based on the result of comparing F_a to 0.0.

Table 13 describes the conditional branch instructions for floating point. They have format similar to the integer branch instructions (shown for `fbeq`):

$$\text{fbeq} \quad F_a, \quad Label$$

The decision of whether or not to branch is based on the result of comparing register F_a to 0.0.

9

Instruction	Description	Branch Condition
fbeq	Branch on Equal	a == 0.0
fbne	Branch on Not Equal	a != 0.0
fbgt	Branch on Greater Than	a > 0.0
fbge	Branch on Greater Than or Equal	a >= 0.0
fblt	Branch on Less Than	a < 0.0
fble	Branch on Less Than or Equal	a <= 0.0

Table 13: Floating Point Branch Instructions

Instruction	From	To
cvtqs	Quad integer	S_floating
cvtqt	Quad integer	T_floating
cvtsq	S_floating	Quad integer
cvttq	T_floating	Quad integer
cvtts	T_floating	S_floating
cvtst	S_floating	T_floating

Table 14: Floating Point Conversion Operations

Finally, as Table 14 indicates, there is a set of operations for converting between the different numeric formats. Each of these has the same format (shown here for cvtqs):

$$\text{cvtqs} \quad F_b, \quad F_c$$

where F_b indicates the source register, and F_c indicates the destination register. Surprisingly, a floating point registers is used even when the source or destination is a quad word integer. The 8 bytes of the floating point register is simply interpreted as a two's complement number in these cases. In fact, the only way to transfer data between the floating point registers and integer registers is to store to memory and then load them into load back to the other register set.

3. Programming Conventions

The Alpha hardware provides only low-level support for handling tasks such as setting up procedure calls, maintaining the calling stack, and allocating space for data structures. Built on top of this low-level support is a set of conventions that all compiler writers and assembly code generators are supposed to follow. Having uniform conventions makes it possible to link together code generated from different sources, e.g., to have your C code be able to use routines from the standard Unix libraries. In addition, it enables tools such as debuggers, profilers, and performance monitors to work on code independent of how it was generated.

In this section we summarize some of the key features of the Alpha programming conventions. More extensive documentation can be found in [1].

Register Name	Software Name	Use
$0	v0	Returned value from integer functions
$1–$8	t0–t7	Temporaries
$9–$14	s0–s5	Callee saved
$15	s6	Callee saved
$15 or $fp	fp	Frame pointer
$16–$21	a0–a5	Integer arguments
$22–$25	t8–t11	Temporaries
$26	ra	Return address
$27	pv	Address of current procedure
$27	t12	Temporary
$28 or $at	AT	Reserved for assembler
$29 or $gp	gp	Global pointer
$30 or $sp	sp	Stack pointer
$31	zero	Always 0

Table 15: Integer Register Usage Conventions

3.1. Register Usage

Alpha has 32 integer registers, identified as $0 up to $31. As far as the hardware goes, only one of these is special—register $31 is always equal to 0. Even if it is the destination of an operation, its value never changes.

The remaining integer registers are partitioned into different groups with different uses, as shown in Table 15. Registers can be identified in several ways: by their numbers $0–$31, by special names $fp, $at, $gp, $sp, or by "software" names that are supposed to more clearly identify their usage conventions. These software names are actually just macro definitions from file regdef.h. We will generally refer to registers by number, since this is seen in the .s files generated by the C compiler.

Observe that some lines in the table refer to the same register, indicating overlapping usages. For example, register $27 generally holds the starting address of the currently executing procedure, but it can also be used just for temporary storage.

Register $0 is used to return an integer (or pointer) value to the calling procedure. Registers $1 to $8 and $22 to $25 can be used by a procedure for arbitrary temporary values. However, if procedure A calls procedure B, there is no guarantee that the values in these registers will be unchanged when B returns back to A. Registers $9 to $14 are "callee saved" registers. That means that any procedure using them must first save the old values on the stack and then restore them before it returns. Thus, if A calls B, it can be assured that these register values are unchanged when B returns back to A, either because B did not alter them, or because B saved, altered, and later restored them.

Register $15 can be used as a "frame pointer," indicating the start of the current stack frame. Most of the time this is not done, however—all stack addressing is done relative to the stack pointer.

Registers $16 to $21 are used to pass integer (or pointer) arguments to a procedure. If needed, more arguments can be passed on the program stack.

Register $26 generally holds the address to which the currently-executing procedure should return.

Register Name	Use
$f0	Returned value from floating point functions
$f1	Returned imaginary value from complex functions
$f2–$f9	Callee saved
$f10–$f15	Temporaries
$f16–$f21	Floating point arguments
$f22–$f30	Temporaries
$f31	Always 0.0

Table 16: Floating Point Register Usage Conventions

Register $27 generally points to the currently executing procedure. The typical way for procedure A to call procedure B is to load the starting address of B into register $27 and then execute the instruction:

```
jsr $26, ($27)
```

Procedure B then returns to A by executing the instruction

```
ret ($26)
```

Register $29 is used as a "global pointer," indicating a region in memory where global data and linkage information is maintained.

Register $30 is used as the "stack pointer," indicating the address of the top element of the stack. The stack grows toward lower addresses.

Table 16 indicates the usage conventions for the floating point registers. Like the integer registers, only floating point register $f31 has any special hardware-implemented features—it is always equal to 0.0. The remainder are partitioned by convention into return values ($f0–$f1), callee saved ($f2–$f9), temporaries ($f10–$f15 and $f22–$f30), and procedure arguments ($f16–$f21).

3.2. Stack Frames

The program stack is used as the working storage for procedures. Each procedure requiring local storage to hold local data or linkage information allocates a stack frame upon entry and deallocates it before returning.

Not all procedures require stack space. As long as a procedure does not call any other procedures and can fit all the data it requires in registers, it need not allocate a frame.

Figure 1 shows the general form of a stack frame. Note that the stack grows toward lower addresses, so the top of the stack is actually the lowest address. We will draw the frame with the "top" of the stack on the bottom of the figure, as is done in our textbook. The figure illustrates the most general stack frame form—not all procedures require all parts.

We will refer to the parts of the frame relative to the currently active procedure, i.e., the one who's frame is at the top of the stack. If this procedure had more arguments than could be passed in the integer or

Stack Bottom

Passed Arg. *n*

·
·
·

Frame Pointer ➡ Passed Arg. 7

Caller's
Frame

Locals
and
Temporaries

Saved Reg. *m*

·
·
·

Saved Reg. 2

Saved Reg. 1

Argument
Build
Area

Current
Frame

Increasing
Address

Stack Pointer ➡

Stack Top

Figure 1: Stack Frame Structure

floating point registers (e.g., if it had more than 6 integer or pointer arguments), the remaining arguments would be on the stack as part of the caller's frame. The "frame pointer" indicates the top of the caller's frame. Typically, this frame pointer is "virtual", meaning that it's value is defined in terms of some offset relative to the stack pointer. The first portion of the frame holds any local or temporary data that cannot be held in registers. This will typically include local arrays and any local variable to which a pointer must be generated. The next part of the frame provides storage for any registers that the procedure needs to save. Typically this includes the return address pointer (stored first) and the old values of any callee save registers to be used by the current procedure. Finally, the top part of the frame consists of temporary space to be used in building the arguments to procedures to be called by the current procedure. This area is required only if a procedure is called having more arguments than can be passed through registers.

A final requirement is that the frame size must be a multiple of 16 bytes. This requirement is satisfied by padding the region for locals and temporaries as needed.

References

[1] *Alpha Assembly Programmer's Guide*, Digital Equipment Corporation, 1996. Adobe acrobat version (`alpha-asm.pdf`) and Postscript version (`alpha-asm.ps`) in directory:

`/afs/cs.cmu.edu/academic/class/15347-s98/public/doc/`

[2] R. L. Sites, R. T. Witek, *Alpha AXP Architecture*, 2nd edition, Digital Press, 1995.