

Introduction to binary formats

It is all about the bits

At its most fundamental, a data structure is simply how we interpret a sequence of bits.

The world is full of bits. Consider this photograph: ¹ :

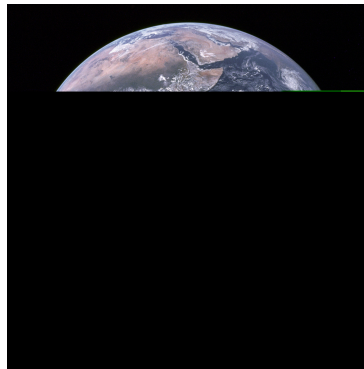


Figure 1.1: The Earth as seen from Apollo 17

Now reconsider it as a monochrome picture:

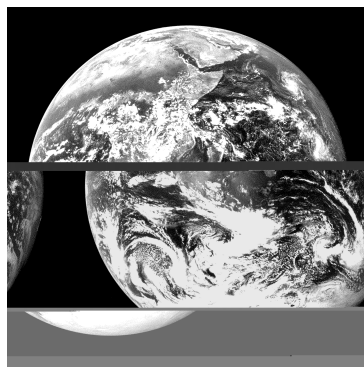


Figure 1.2: The Earth in monochrome

¹Via Wikipedia's article on Earth. This photograph was taken by the Apollo 17 crew. It is thought that either Harrison Schmitt or Ron Evans took this photograph.

Now map these black and white dots to bits; here are a selection from middle of the image, broken as a consecutive sequence of 7 bits:

```

1 0000001 0011010 0010010 0100111 1010110 1111010 1011110
2 1111111 1111000 0111101 0101111 0100100 1010100 0100000
3 1000000 0001100 0000000 0111101 0111011 1111100 0000000
4 0000110 1010011 0000000 0000000 0000001 0000000 1110101
5 1110101 1111011 1110110 1010110 1011101 0110110 1011011

```

Or, converting our groups of 7 bits as base 16

```

1 01 1A 12 27 56 7A 5E
2 7F 78 3D 2F 24 54 20
3 40 0C 00 3D 3B 7C 00
4 06 53 00 00 01 00 75
5 75 7B 76 56 5D 36 5B

```

Or rendered as (somewhat) more familiar ASCII:

```

1 SOH SUB DC2 ' V z ^
2 DEL x = / $ T SPC
3 @ FF NUL = ; | NUL
4 ACK S NUL NUL SOH NUL u
5 u { v V ] 6 [

```

But you can as easily (and more naturally) view these self-same bits also as a series of 32bit unsigned integers. Here are same bits broken into groups of 8 bits:

```

1 00000010 01101000 10010010 01111010 11011110 10101111
2 01111111 11110000 11110101 01111010 01001010 10001000
3 00100000 00001100 00000000 11110101 11011111 11000000
4 00000001 10101001 10000000 00000000 00000100 00000111
5 01011110 10111110 11111011 01010110 [...]

```

While this is a more natural view from the viewpoint of a computer, we might want to view it as base 10:

1	40,407,674	3,736,043,504	4,118,432,392
2	537,657,589	3,753,902,505	2,147,484,679
3	1,589,574,486		

From a “bottom-up” perspective, how we understand bits is the heart of structuring data. Data structures are fundamentally an understanding of bits.

Take the previous item; we can regard that series of seven 32 bit unsigned integers as a traditional C array named “v0”:

```
1  #include <stdint.h>
2
3  uint32_t v0[7];
4  v0[0] = 40407674;
5  v0[1] = 3736043504;
6  v0[2] = 4118432392;
7  v0[3] = 537657589;
8  v0[4] = 3753902505;
9  v0[5] = 2147484679;
10 v0[6] = 1589574486;
```

And, finally, completing the circle, we can write a small C program to display the bits:

```
1  #include <stdint.h>
2  #include <stdlib.h>
3  #include <stdio.h>
4
5  //
6  // The following code is courtesy of https://stackoverflow.com/
7  // questions/111928,
8  // from the posting by "ideasman42".
9  //
10 /* --- PRINTF_BYTE_TO_BINARY macro's --- */
11 #define PRINTF_BINARY_PATTERN_INT8 "%c%c%c%c%c%c%c%c "
12 #define PRINTF_BYTE_TO_BINARY_INT8(i) \
13     (((i) & 0x80ll) ? '1' : '0'), \
14     (((i) & 0x40ll) ? '1' : '0'), \
15     (((i) & 0x20ll) ? '1' : '0'), \
```

```

16     (((i) & 0x10ll) ? '1' : '0'), \
17     (((i) & 0x08ll) ? '1' : '0'), \
18     (((i) & 0x04ll) ? '1' : '0'), \
19     (((i) & 0x02ll) ? '1' : '0'), \
20     (((i) & 0x01ll) ? '1' : '0')
21
22 #define PRINTF_BINARY_PATTERN_INT16 \
23     PRINTF_BINARY_PATTERN_INT8          PRINTF_BINARY_PATTERN_INT8
24 #define PRINTF_BYTE_TO_BINARY_INT16(i) \
25     PRINTF_BYTE_TO_BINARY_INT8((i) >> 8), PRINTF_BYTE_TO_BINARY_INT8(
26     i)
27 #define PRINTF_BINARY_PATTERN_INT32 \
28     PRINTF_BINARY_PATTERN_INT16          PRINTF_BINARY_PATTERN_INT16
29 #define PRINTF_BYTE_TO_BINARY_INT32(i) \
30     PRINTF_BYTE_TO_BINARY_INT16((i) >> 16), PRINTF_BYTE_TO_BINARY_INT16
31     (i)
32 #define PRINTF_BINARY_PATTERN_INT64 \
33     PRINTF_BINARY_PATTERN_INT32          PRINTF_BINARY_PATTERN_INT32
34 #define PRINTF_BYTE_TO_BINARY_INT64(i) \
35     PRINTF_BYTE_TO_BINARY_INT32((i) >> 32), PRINTF_BYTE_TO_BINARY_INT32
36     (i)
37 /* --- end macros --- */
38
39 // End of code from stackoverflow.
40
41 int main()
42 {
43     uint32_t v0[7];
44     v0[0] = 40407674;
45     v0[1] = 3736043504;
46     v0[2] = 4118432392;
47     v0[3] = 537657589;
48     v0[4] = 3753902505;
49     v0[5] = 2147484679;
50     v0[6] = 1589574486;
51
52     printf("The base address for v0 is %p\n",v0);
53     for(int i=0; i<7; i++)
54     {
55         printf("at address %p, we have v0[%d] = "
56             PRINTF_BINARY_PATTERN_INT32 "\n",&v0[i],i,
57             PRINTF_BYTE_TO_BINARY_INT32(v0[i]));
58     }
59     uint64_t *v1 = (uint64_t *)v0;
60
61     printf("\n");
62
63     printf("The base address for v1 is %p\n",v1);
64     for(int i=0; i<3; i++)
65     {
66         printf("at address %p, we have v1[%d] = "

```



```
        PRINTF_BINARY_PATTERN_INT64 "\\n",&v1[i],i,  
        PRINTF_BYTE_TO_BINARY_INT64(v1[i]));  
62     }  
63  
64 }
```

And here's what we see when we run our program:

```
1 $ ./example1  
2 The base address for v0 is 0x7fff140be7f0  
3 at address 0x7fff140be7f0, we have v0[0] = 00000010 01101000 10010010  
   01111010  
4 at address 0x7fff140be7f4, we have v0[1] = 11011110 10101111 01111111  
   11110000  
5 at address 0x7fff140be7f8, we have v0[2] = 11110101 01111010 01001010  
   10001000  
6 at address 0x7fff140be7fc, we have v0[3] = 00100000 00001100 00000000  
   11110101  
7 at address 0x7fff140be800, we have v0[4] = 11011111 11000000 00000001  
   10101001  
8 at address 0x7fff140be804, we have v0[5] = 10000000 00000000 00000100  
   00000111  
9 at address 0x7fff140be808, we have v0[6] = 01011110 10111110 11111011  
   01010110  
10  
11 The base address for v1 is 0x7fff140be7f0  
12 at address 0x7fff140be7f0, we have v1[0] = 11011110 10101111 01111111  
   11110000 00000010 01101000 10010010 01111010  
13 at address 0x7fff140be7f8, we have v1[1] = 00100000 00001100 00000000  
   11110101 11110101 01111010 01001010 10001000  
14 at address 0x7fff140be800, we have v1[2] = 10000000 00000000 00000100  
   00000111 11011111 11000000 00000001 10101001
```

We can regard those individual bits as a sequence of 7 bit ASCII, or 8 bit bytes, or 32 bit unsigned integers, or 64 bit unsigned integers: it's not the bits that make it information, it's our interpretation of these bits.

Likewise, a data structure is our agreement as to the meaning of a given arrangement of bits.

A data structure can be a contiguous, like all of our data bits from the monochromatic rendition of Earth:

XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX
0000001	0011010	0010010	0100111	1010110	1111010	1011110
1111111	1111000	0111101	0101111	0100100	1010100	0100000
1000000	0001100	0000000	0111101	0111011	1111100	0000000
0000110	1010011	0000000	0000000	0000001	0000000	1110101
1110101	1111011	1110110	1010110	1011101	0110110	1011011
XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX	XXXXXXX

How do we understand common data structures?

Arrays of native types are generally implemented in C and C++ in this fashion, as you can see from running example1. C strings are also done in the fashion, but a new issue starts to creep in, alignment. ASCII is a 7 bit standard, but we store it 8 bit bytes. ²

Structs and classes are also laid in this fashion, but, like with ASCII, you start to run into issues such as alignment; for instance, look at this code:

```

1  #include <stdint.h>
2  #include <stdlib.h>
3  #include <stdio.h>
4
5  struct s0
6  {
7      int8_t f0;
8      int32_t f1;
9      int8_t f2;
10     int16_t f3;
11     int8_t f4;
12 };
13
14 int main()
15 {
16     struct s0 s;
17
18
19     printf("field f0 is at address %p\n",&s.f0);
20     printf("field f1 is at address %p\n",&s.f1);
21     printf("field f2 is at address %p\n",&s.f2);

```

²This actually turns out to be a benefit; UTF-8 encoding of Unicode actually turns out to embed 7bit ASCII by the convention that if the high bit is zero, then it uses the 7-bit ASCII value, and if the top bit is instead set, then it maps this byte and succeeding bytes to a Unicode character.

```
22     printf("field f3 is at address %p\n",&s.f3);
23     printf("field f4 is at address %p\n",&s.f4);
24
25 }
```

When you compile this code with clang and run it, you get:

```
1 clang -o example2 example2.c
2 langley@localhost ~/mounts/www/public_html/COP4530/Lectures $ ./
  example2
3 field f0 is at address 0x7ffee5d8b280
4 field f1 is at address 0x7ffee5d8b284
5 field f2 is at address 0x7ffee5d8b288
6 field f3 is at address 0x7ffee5d8b28a
7 field f4 is at address 0x7ffee5d8b28c
```

But f0 is only 8bits, so you might have expected f1 to start at 0x7ffee5d8b281 rather than at 0x7ffee5d8b284. But compilers tend to optimize for speed rather than memory efficiency. While the x86_64 family of processors can read f1 at either address (not a given with all processors), such unaligned access does have a significant runtime penalty for the x86_64 architecture.

If you add a *pragma*, you can ask the compiler to do such packing for you:

```
1 #include <stdint.h>
2 #include <stdlib.h>
3 #include <stdio.h>
4
5 struct s0
6 {
7     int8_t f0;
8     int32_t f1;
9     int8_t f2;
10    int16_t f3;
11    int8_t f4;
12 } __attribute__((packed));
13
14 int main()
15 {
16     struct s0 s;
17
18
19     printf("field f0 is at address %p\n",&s.f0);
20     printf("field f1 is at address %p\n",&s.f1);
```

```
21     printf("field f2 is at address %p\n",&s.f2);
22     printf("field f3 is at address %p\n",&s.f3);
23     printf("field f4 is at address %p\n",&s.f4);
24
25 }
```

Now when you run the code, you get these in truly consecutive order:

```
1 ./example2
2 field f0 is at address 0x7ffc3756f540
3 field f1 is at address 0x7ffc3756f541
4 field f2 is at address 0x7ffc3756f545
5 field f3 is at address 0x7ffc3756f546
6 field f4 is at address 0x7ffc3756f548
```

Data structures are composed of elements, such as integers, and the relationships among those elements.

There are two ways of thinking of data structures: an “abstract” data structure, and a “realized” data structure. An abstract data structure only specifies that some sort of relationship exists between elements; a realized data structure specifies the actual relationships and the actual elements.

A realized data structure extends this idea from merely adjacent bits as a basic type, such as an integer, to multiple elements. The elements follow an agreed pattern; the agreement can be based on simple adjacency (i.e., adjacent elements have (effectively) a zero distance between them (subject to alignment issues)), or it can be based on internal components that specify the location of other elements, such as pointers.

Arrays are usually implemented by agreement; the most common agreement is that 1) each element of the array is of uniform type, and 2) that all elements are laid consecutively. For example, let’s declare an array `arr` of type `uint32_t` with 20 elements, and look at how it’s laid out in memory:

```
1 #include <stdint.h>
2 #include <stdlib.h>
3 #include <stdio.h>
4
5
6 int main()
7 {
8     uint32_t arr[20];
9
10    printf("Array 'arr' begins at address %p and is %lu bytes in size.\n",
11           arr, sizeof(arr));
12    for(int i = 0; i < 20; i++)
13    {
```

```
13
14     printf("\t element %02d starts at %p and is %lu bytes in size.\n"
15           ,i,&arr[i],sizeof(arr[i]));
16 }
```

When you run this code, you see that the array is made of 20 uniformly 4 byte integers packed side-by-side:

```
1 ./example3
2 Array 'arr' begins at address 0x7ffde225d760 and is 80 bytes in size.
3     element 00 starts at 0x7ffde225d760 and is 4 bytes in size.
4     element 01 starts at 0x7ffde225d764 and is 4 bytes in size.
5     element 02 starts at 0x7ffde225d768 and is 4 bytes in size.
6     element 03 starts at 0x7ffde225d76c and is 4 bytes in size.
7     element 04 starts at 0x7ffde225d770 and is 4 bytes in size.
8     element 05 starts at 0x7ffde225d774 and is 4 bytes in size.
9     element 06 starts at 0x7ffde225d778 and is 4 bytes in size.
10    element 07 starts at 0x7ffde225d77c and is 4 bytes in size.
11    element 08 starts at 0x7ffde225d780 and is 4 bytes in size.
12    element 09 starts at 0x7ffde225d784 and is 4 bytes in size.
13    element 10 starts at 0x7ffde225d788 and is 4 bytes in size.
14    element 11 starts at 0x7ffde225d78c and is 4 bytes in size.
15    element 12 starts at 0x7ffde225d790 and is 4 bytes in size.
16    element 13 starts at 0x7ffde225d794 and is 4 bytes in size.
17    element 14 starts at 0x7ffde225d798 and is 4 bytes in size.
18    element 15 starts at 0x7ffde225d79c and is 4 bytes in size.
19    element 16 starts at 0x7ffde225d7a0 and is 4 bytes in size.
20    element 17 starts at 0x7ffde225d7a4 and is 4 bytes in size.
21    element 18 starts at 0x7ffde225d7a8 and is 4 bytes in size.
22    element 19 starts at 0x7ffde225d7ac and is 4 bytes in size.
```

Now we consider the case where we have an indicator from one element to another element; this particular indicator is a “pointer”, which is just a variable that has memory address in it. Here is some code that implements a simple list of integers using pointers:

```
1
2 #include <stdint.h>
3 #include <stdlib.h>
4 #include <stdio.h>
```

```
5
6 struct st
7 {
8     int val;
9     struct st *next;
10 };
11
12 int main()
13 {
14     struct st *struct0;
15     struct st *struct1;
16     struct st *struct2;
17
18     struct0 = malloc(sizeof(struct st));
19     printf("For struct0, we allocated %lu bytes at memory location %p.\n"
20           ,sizeof(struct st),struct0);
21     struct1 = malloc(sizeof(struct st));
22     printf("For struct1, we allocated %lu bytes at memory location %p (
23           distance struct1-struct0 is %ld bytes).\n",sizeof(struct st),
24           struct1,(void *)struct1-(void *)struct0);
25     struct2 = malloc(sizeof(struct st));
26     printf("For struct2, we allocated %lu bytes at memory location %p (
27           distance struct2-struct1 is %ld bytes).\n",sizeof(struct st),
28           struct2,(void*)struct2-(void *)struct1);
29
30     struct0->val = 1;
31     struct0->next = struct1;
32
33     struct1->val = 2;
34     struct1->next = struct2;
35
36     struct2->val = 3;
37     struct2->next = NULL;
38
39     struct st *s = struct0;
40     while(s)
41     {
42         printf("This element is at memory location %p; it has value %d,
43               and a pointer %p to a next element.\n",
44               s,
45               s->val,
46               s->next);
47         s=s->next;
48     }
49 }
```

If we were to run this code, we can see now that while the pointers are linearly increasing (the heap

grows up), these are not contiguous:

```
1 $ clang -g -o example4 example4.c
2 fsucs@localhost ~/mounts/www/public_html/COP4530/Lectures $ ./example4
3 For struct0, we allocated 16 bytes at memory location 0x2349260.
4 For struct1, we allocated 16 bytes at memory location 0x2349690 (
   distance struct1-struct0 is 1072 bytes).
5 For struct2, we allocated 16 bytes at memory location 0x23496b0 (
   distance struct2-struct1 is 32 bytes).
6 This element is at memory location 0x2349260; it has value 1, and a
   pointer 0x2349690 to a next element.
7 This element is at memory location 0x2349690; it has value 2, and a
   pointer 0x23496b0 to a next element.
8 This element is at memory location 0x23496b0; it has value 3, and a
   pointer (nil) to a next element.
```

Finally, we will look at an actual data structure being created in x86_64 assembly language. This data structure is called an “ELF header”. It’s what is used, for instance, as the header for every binary executable on an x86_64 Linux computer.

Here’s the NASM code for a trivial program that lays out its own ELF header and then has a trivial body that only exits (with 42, naturally):

```
1
2     ;; inspired by http://www.muppetlabs.com/~breadbox/software/tiny/teensy.html
3     ;; https://blog.stalkr.net/2014/10/tiny-elf-3264-with-nasm.html
4     ;; ... and others of similar ilk
5
6
7 BITS 64
8
9 ORG 0x400000
10
11 ;;;
12 ;;; Definitions from "ELF-64 Object File Format" (aka "EOFF document")
13 :
14 ;;;
15 ;;; Elf64_Addr    8 bytes, aligned on 8 bytes ; program address
16 ;;; Elf64_Off    8 bytes, aligned on 8 bytes ; file offset
```

```

17 ;;;      Elf64_Half      2 bytes, aligned on 2 bytes ; medium integer
18 ;;;      Elf64_Word      4 bytes, aligned on 4 bytes ; integer
19 ;;;      Elf64_Sword      4 bytes, aligned on 4 bytes ; signed integer
20 ;;;      Elf64_Xword      8 bytes, aligned on 8 bytes ; long integer
21 ;;;      Elf64_Sxword      8 bytes, aligned on 8 bytes ; signed long
    integer
22 ;;;      unsigned char 1 byte, aligned on 1 byte ; small integer
23
24
25 elf64_file_header:      ; This is often just called the elf64 header
26
27 ;;; at 0: unsigned char e_ident[16]
28 db 127,"ELF"           ; e_ident[0-3]: EI_MAG{0,1,2,3} (aka "magic
    number")
29 db 2                   ; e_ident[4]: EI_CLASS; ELFCLASS32=1, ELFCLASS64=2
    (aka "File class")
30 db 1                   ; e_ident[5]: EI_DATA; ELFDATALSB=1, ELFDATAMSB=2 (
    aka "Data encoding")
31 db 1                   ; e_ident[6]: EI_VERSION; EV_CURRENT=1 (aka "File
    version")
32 db 0                   ; e_ident[7]: EI_OSABI; ELFOSABI_SYSV=0 (aka "OS/
    ABI identification")
33 db 0                   ; e_ident[8]: EI_ABI (always zero)
34 times 7 db 0           ; e_ident[9-15]: EI_PAD
35
36 ;;; at 16: Elf64_Half
37 dw 2                   ; e_type: 2 = "executable file" (aka "object file
    type")
38 ;;; at 18: Elf64_Half
39 dw 62                   ; e_machine: EM_X86_64 = 62 (aka "machine type")
40 ;;;      that is found for Linux in "include/uapi/linux/elf-em
    .h"
41
42 ;;; at 20: Elf64_Word
43 dd 1                   ; e_version: always 1
44 ;;; at 24: Elf64_Addr
45 dq _start              ; e_entry: the address where you want to start
    running
46 ;;; at 32: Elf64_Off
47 dq elf64_program_header - $$
48 ;;;      ; e_phoff: offset to program header(s) start - required
    in
49 ;;;      ; all executables since they give the actual segments
    to be
50 ;;;      ; to be laid out in memory
51 ;;; at 40: Elf64_Off
52 dq 0                   ; e_shoff: offset to section header(s) start - not
    required in
53 ;;;      ; static executables since sections are
    only important for relocation
54 ;;; at 48: Elf64_Word

```



```

55     dd 0                ; e_flags: processor-specific flags
56     ;;                  ; (where is this documented in the kernel
    source code?)
57     ;;
58     ;; at 52: Elf64_Half
59     dw elf64_file_header_size ; e_ehsize: elf64 file header size
60     ;; at 54: Elf64_Half
61     dw elf64_program_header_entry_size
62     ;;                  ; e_phentsize: size of one program header
    entry
63     ;;
64     ;; at 56: Elf64_Half
65     dw 1                ; e_phnum: how many program header entries do we
    have?
66     ;; at 58: Elf64_Half
67     dw 0                ; e_shentsize: size of one section header entry
68     ;; at 60: Elf64_Half
69     dw 0                ; e_shnum: how many section header entries do we
    have?
70     ;; at 62: Elf64_Half
71     dw 0                ; e_shstrndx: section name string table index
72     ;;
73     ;;
74     elf64_file_header_size equ $ - elf64_file_header
75     ;;
76     ;;                  ; compute how big the header was (64 bytes
    !)
77
78
79     elf64_program_header:      ; This is our only program header since we
    only want one segment
80
81     ;; at 64: Elf64_Word
82     dd 1                ; p_type: type of segment, 1 = "loadable segment" (
    from EOFF document)
83     ;; at 68: Elf64_Word
84     dd 7                ; p_flags: segment attributes; 0x1 = execute
    permission
85     ;;                  ; 0x2 = write
    permission
86     ;;                  ; 0x4 = read
    permission
87     ;;
88     ;; at 72: Elf64_Off
89     dq 0                ; p_offset: offset in file -- where does this
    segment start in file?
90     ;; at 80: Elf64_Addr
91     dq $$               ; p_vaddr: virtual address of the segment in memory
92     ;; at 88: Elf64_Addr
93     dq $$               ; p_paddr: reserved for systems with physical
    addressing

```

```
94     ;; at 96: Elf64_Xword
95     dq total_size      ; p_filesz: size of segment in file
96     ;; at 104: Elf64_Xword
97     dq total_size      ; p_memsz: size of segment in memory
98     ;; at 112: Elf64_Xword
99     dq 0x1000          ; p_align: alignment of the segment. p_offset =
    p_vaddr MOD p_align
100     ;;
101     elf64_program_header_entry_size equ $ - elf64_program_header
102
103
104     _start:
105         mov rax, 231      ; sys_exit_group
106         mov rdi, 42      ; answer to everything
107         syscall
108
109     total_size equ $ - $$
```