

## Counting in Binary

The polynomial expansion of the decimal number 362 is  $3 \times 10^2 + 6 \times 10^1 + 2 \times 10^0$ . Binary numbers work in exactly the same way, but with powers of 2 instead of powers of 10. So the number  $00100101 = 2^5 + 2^2 + 2^0 = 37$ .

## Data and Encodings

Imagine that we write a small amount of textual data to file a file called `output.txt` using the following Python code.

```
1 with open('output.txt', 'wt', encoding='ASCII') as fout:
2     print("De La Soul is Dead", file=fout, )
```

If we cat the file in unix we see

```
$ cat output.txt
De La Soul is Dead
```

We also know that the data must really just be a sequence of zeros and ones, so we can use the `xxd` command to see this underlying representation

```
$ xxd -b output.txt
0000000: 01000100 01100101 00100000 01001100 01100001 00100000  De La
0000006: 01010011 01101111 01110101 01101100 00100000 01101001  Soul i
000000c: 01110011 00100000 01000100 01100101 01100001 01100100  s Dead
0000012: 00001010  .
```

Each line tells you the start of the next byte, in hexadecimal, the next 6 bytes, and then the *textual* interpretation of the bytes. Bytes are just 8 bits of data. Notice that the first byte is `01000100`, which is  $2^6 + 2^2 = 68$ . The first character is “D”. This is not a coincidence. When the “D” was written to disk, it was written using the ASCII encoding. ASCII stands for American Standard Code for Information Interchange. It encodes the character “D” using the integer 68. ASCII encodes 128 characters, requiring 7 bits, which you can see in the cart below.

Dec	Hx	Oct	Chr	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr
0	0	000	<b>NUL</b> (null)	32	20	040	<b>Space</b>	64	40	100	<b>@</b>	96	60	140	<b>#96;</b>			
1	1	001	<b>SOH</b> (start of heading)	33	21	041	<b>!</b>	65	41	101	<b>A</b>	97	61	141	<b>#97;</b>			
2	2	002	<b>STX</b> (start of text)	34	22	042	<b>"</b>	66	42	102	<b>B</b>	98	62	142	<b>#98;</b>			
3	3	003	<b>ETX</b> (end of text)	35	23	043	<b>#</b>	67	43	103	<b>C</b>	99	63	143	<b>#99;</b>			
4	4	004	<b>EOT</b> (end of transmission)	36	24	044	<b>\$</b>	68	44	104	<b>D</b>	100	64	144	<b>#100;</b>			
5	5	005	<b>ENQ</b> (enquiry)	37	25	045	<b>%</b>	69	45	105	<b>E</b>	101	65	145	<b>#101;</b>			
6	6	006	<b>ACK</b> (acknowledge)	38	26	046	<b>&amp;</b>	70	46	106	<b>F</b>	102	66	146	<b>#102;</b>			
7	7	007	<b>BEL</b> (bell)	39	27	047	<b>'</b>	71	47	107	<b>G</b>	103	67	147	<b>#103;</b>			
8	8	010	<b>BS</b> (backspace)	40	28	050	<b>(</b>	72	48	110	<b>H</b>	104	68	150	<b>#104;</b>			
9	9	011	<b>TAB</b> (horizontal tab)	41	29	051	<b>)</b>	73	49	111	<b>I</b>	105	69	151	<b>#105;</b>			
10	A	012	<b>LF</b> (NL line feed, new line)	42	2A	052	<b>*</b>	74	4A	112	<b>J</b>	106	6A	152	<b>#106;</b>			
11	B	013	<b>VT</b> (vertical tab)	43	2B	053	<b>+</b>	75	4B	113	<b>K</b>	107	6B	153	<b>#107;</b>			
12	C	014	<b>FF</b> (NP form feed, new page)	44	2C	054	<b>,</b>	76	4C	114	<b>L</b>	108	6C	154	<b>#108;</b>			
13	D	015	<b>CR</b> (carriage return)	45	2D	055	<b>-</b>	77	4D	115	<b>M</b>	109	6D	155	<b>#109;</b>			
14	E	016	<b>SO</b> (shift out)	46	2E	056	<b>.</b>	78	4E	116	<b>N</b>	110	6E	156	<b>#110;</b>			
15	F	017	<b>SI</b> (shift in)	47	2F	057	<b>/</b>	79	4F	117	<b>O</b>	111	6F	157	<b>#111;</b>			
16	10	020	<b>DLE</b> (data link escape)	48	30	060	<b>0</b>	80	50	120	<b>@</b>	112	70	160	<b>#112;</b>			
17	11	021	<b>DC1</b> (device control 1)	49	31	061	<b>1</b>	81	51	121	<b>A</b>	113	71	161	<b>#113;</b>			
18	12	022	<b>DC2</b> (device control 2)	50	32	062	<b>2</b>	82	52	122	<b>B</b>	114	72	162	<b>#114;</b>			
19	13	023	<b>DC3</b> (device control 3)	51	33	063	<b>3</b>	83	53	123	<b>C</b>	115	73	163	<b>#115;</b>			
20	14	024	<b>DC4</b> (device control 4)	52	34	064	<b>4</b>	84	54	124	<b>D</b>	116	74	164	<b>#116;</b>			
21	15	025	<b>NAK</b> (negative acknowledge)	53	35	065	<b>5</b>	85	55	125	<b>E</b>	117	75	165	<b>#117;</b>			
22	16	026	<b>SYN</b> (synchronous idle)	54	36	066	<b>6</b>	86	56	126	<b>F</b>	118	76	166	<b>#118;</b>			
23	17	027	<b>ETB</b> (end of trans. block)	55	37	067	<b>7</b>	87	57	127	<b>G</b>	119	77	167	<b>#119;</b>			
24	18	030	<b>CAN</b> (cancel)	56	38	070	<b>8</b>	88	58	130	<b>X</b>	120	78	170	<b>#120;</b>			
25	19	031	<b>EM</b> (end of medium)	57	39	071	<b>9</b>	89	59	131	<b>Y</b>	121	79	171	<b>#121;</b>			
26	1A	032	<b>SUB</b> (substitute)	58	3A	072	<b>:</b>	90	5A	132	<b>Z</b>	122	7A	172	<b>#122;</b>			
27	1B	033	<b>ESC</b> (escape)	59	3B	073	<b>;</b>	91	5B	133	<b>[</b>	123	7B	173	<b>#123;</b>			
28	1C	034	<b>FS</b> (file separator)	60	3C	074	<b>&lt;</b>	92	5C	134	<b>\</b>	124	7C	174	<b>#124;</b>			
29	1D	035	<b>GS</b> (group separator)	61	3D	075	<b>=</b>	93	5D	135	<b>]</b>	125	7D	175	<b>#125;</b>			
30	1E	036	<b>RS</b> (record separator)	62	3E	076	<b>&gt;</b>	94	5E	136	<b>^</b>	126	7E	176	<b>#126;</b>			
31	1F	037	<b>US</b> (unit separator)	63	3F	077	<b>?</b>	95	5F	137	<b>_</b>	127	7F	177	<b>#127;</b>			

Source: [www.LookupTables.com](http://www.LookupTables.com)

## ASCII is so 1997

There's a sociology thesis to be written about ASCII—it encodes the English alphabet, numbers, some punctuation, and a few out-of-date control characters for your dot matrix printer and terminal like a LINE FEED (ASCII 10) and BELL (ASCII 7)—but nothing more. It is like saying the world is all of Iowa or any other state in the union. But it is also pervasive, so we need to deal with it intelligently.

## Unicode

Unicode recognizes that the world doesn't start and end with the English alphabet. It associates an integer (called a code point) with a platonic character, like the symbol "@" (U+0040, 64 in decimal) or a glyph of an anchor (U+2693, 9875 in decimal). There are 1,114,112 code points overall, in the range from 0 to 10FFFF in hexadecimal.

## Encodings

One represents unicode through *encodings*. For example, one could represent strings of characters as code points directly. For example, we could just use a standard 32-bit integer to store a code point and then strings would be lists or arrays of integers. This has several disadvantages:

- Because code points are grouped by locale, most of the bits in single encoded code point will be 0. This seems wasteful.
- Each character requires 21 bits to represent it. Because the basic unit of addressable memory on a computer is the byte, which is traditionally 8-bits, we are really working with 24 bits, so even going with a fixed-length, 3-byte encoding seems wasteful.
- Different processors order bytes differently, so a fixed-length encoding is not portable.

## UTF-8

UTF-8 stands for Unicode Transformation Format. The 8 means that encodings are done in 8-bit blocks (called octets by UTF-8). UTF-8 is a variable-length encoding, meaning that depending on which unicode code point it's transforming, the corresponding encoding will either be 1,2,3 or 4 bytes long. Here is a schematic of how UTF-8 encodes data.

Bits in code pt	First code pt	Last code pt	# Bytes	Byte 1	Byte 2	Byte 3	Byte 4
7	U+0000	U+007F	1	0xxxxxxx			
11	U+0080	U+07FF	2	110xxxxx	10xxxxxx		
16	U+0800	U+FFFF	3	1110xxxx	10xxxxxx	10xxxxxx	
21	U+10000	U+1FFFFF	4	11110xxx	10xxxxxx	10xxxxxx	10xxxxxx

Here are several advantages of UTF-8.

- UTF-8 is compatible with ASCII because it encodes all ASCII characters as ASCII values;
- UTF-8 can encode all the unicode code points; and
- UTF-8 is self-synchronizing—one can use the byte signatures to both determine the number of bytes and the order of the bytes.

## Unicode in Python

The `str` type in Python is a sequence of unicode code points. Those points are stored compactly, under-the-hood, so space is not wasted. Given a character, one can grab it's code point using the `ord` built-in function.

```
1 >>> print("".join([str(ord(c)) for c in 'De La Soul is Dead']))
2 >>> 68 101 32 76 97 32 83 111 117 108 32 105 115 32 68 101 97 100
```

One can convert a unicode code point into a character using the built-in function `chr`.

```
1 >>> codepoints = [int(x) for x in "68 101 32 76 97 32 83 111 117 108 32 105 115 32 68 101 97 100".split()]
2 >>> print("".join([chr(cp) for cp in codepoints]))
3 >>> De La Soul is Dead
```

One can embed unicode code points in hexadecimal directly in string literals using the `u` preface, so

```
>>> '\u2603'
```

represents a picture of a snowman.

## Encoding Unicode

Writing strings to files mean encoding them in some format. By default this is UTF-8, but Python supports writing in many different formats including ASCII. Use the `encode` method of Python to encode unicode data in a particular format.

```
>>> "De La Soul is Dead".encode("utf-8")
b'De La Soul is Dead'
```

Notice the `b` prefix means this is binary data—the interpreter is printing those characters to the screen by actually decoding them (it will do this for any ASCII character) but if we throw on a unicode star (U+2b50) we'll see some raw hex.

```
>>> "De La Soul is Dead\u2b50".encode("utf-8")
b'De La Soul is Dead\xe2\xad\x90'
```

Notice that in UTF-8, it takes three bytes to encode a star. These bytes in binary are

```
>>> "".join(["{:b}".format(x) for x in "\u2b50".encode("utf-8")])
'11100010 10101101 10010000'
```

So if we wrote this string out to disk

```
1 with open('output2.txt', 'wt') as fout:
2     print("De La Soul is Dead\u2b50", file=fout, end="")
```

we see that the final three bytes are faithfully preserved and `xxd` does not know how to decode them, so it uses the period. Here you see firsthand that UTF-8 uses ASCII encodings for ASCII characters and is also variable length.

```
$ xxd -b output2.txt
0000000: 01000100 01100101 00100000 01001100 01100001 00100000  De La
0000006: 01010011 01101111 01110101 01101100 00100000 01101001  Soul i
000000c: 01110011 00100000 01000100 01100101 01100001 01100100  s Dead
0000012: 11100010 10101101 10010000  . . . .
```