

---

Subject: Re: counting bits

Posted by [thompson](#) on Wed, 26 Feb 2003 20:15:28 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

JD Smith <jdsmith@as.arizona.edu> writes:

```
> IDL> r=ulong(randomu(sd,100)*2.^31) & for i=0,31 do print,FORMAT='(I2," ": ",I2,A)',i,total((r AND  
ulong(2.D^i)) ne 0UL),'% set'
```

```
> 0: 0% set  
> 1: 0% set  
> 2: 1% set  
> 3: 1% set  
> 4: 9% set  
> 5: 17% set  
> 6: 27% set  
> 7: 59% set  
> 8: 44% set  
> 9: 50% set  
> 10: 46% set  
> 11: 57% set  
> 12: 50% set  
> 13: 55% set  
> 14: 51% set  
> 15: 48% set  
> 16: 56% set  
> 17: 51% set  
> 18: 52% set  
> 19: 43% set  
> 20: 46% set  
> 21: 44% set  
> 22: 35% set  
> 23: 52% set  
> 24: 47% set  
> 25: 51% set  
> 26: 44% set  
> 27: 51% set  
> 28: 46% set  
> 29: 53% set  
> 30: 45% set  
> 31: 0% set
```

That's pretty simple to explain. Floating point numbers are stored with a mantissa and an exponent, both stored within the same 4 byte word. A few of the bits are devoted to the exponent, and the rest are devoted to the mantissa. When you call

```
randomu(sd,100)
```

you generate a bunch of numbers which mostly have the same exponent bits, because all the numbers are of the same order of magnitude, while the mantissa bits are generally 50% on or 50% off. When you then multiply this by  $2.^{31}$  and convert it into a long integer, you're primarily sampling the mantissa bits. In fact, the only reason why the last few bits are sometimes set at all is that some of the random numbers are close to zero, and thus end up with different exponents.

You can see this by looking directly at the bits of the original floating point numbers.

```
IDL> r=ulong(randomn(sd,100),0,100) & for i=0,31 do print,FORMAT='(I2," ",I2,A)',i,total((r AND  
ulong(2.D^i)) ne 0UL),'% set'
```

```
0: 58% set  
1: 49% set  
2: 48% set  
3: 48% set  
4: 49% set  
5: 49% set  
6: 47% set  
7: 42% set  
8: 53% set  
9: 51% set  
10: 55% set  
11: 46% set  
12: 44% set  
13: 48% set  
14: 50% set  
15: 59% set  
16: 40% set  
17: 52% set  
18: 53% set  
19: 57% set  
20: 48% set  
21: 41% set  
22: 43% set  
23: 43% set  
24: 75% set  
25: 86% set  
26: 96% set  
27: 96% set  
28: 96% set  
29: 96% set  
30: 4% set  
31: 53% set
```

See how the mantissa is stored in the lower bits, and there's very little variation in the uppermost bits where the mantissa is stored?

