



链滴

【读书笔记】《Computer System A Programmer's Perspective》 Writing TMin32 in C

作者: [armstrong](#)

原文链接: <https://ld246.com/article/1357370795290>

来源网站: [链滴](#)

许可协议: [署名-相同方式共享 4.0 国际 \(CC BY-SA 4.0\)](#)

Writing TMin32 in C (P105, second edition)

```
<div>
</div>
<div>
<br />
</div>
<div>
    &nbsp; &nbsp;&nbsp;C语言头文件<limits.h>中定义了TMin和TMax的取值:
</div>
<div>
<br />
</div>
<div> &nbsp; &nbsp; &nbsp; &nbsp;
<pre>    /*
 * Maximum and minimum values for ints.
 */
#define INT_MAX 2147483647
#define INT_MIN (-INT_MAX-1)
#define UINT_MAX 0xffffffff
/* </pre>
<br />
</div>
<div>
<br />
</div>
<div>
    &nbsp;可以看到32位有符号整数取值范围为[-2147483648,2147483647]。书中提到了一
有趣的问题：为什么最小值要写成-2147483647-1 (-INT_MAX-1)？为什么不直接写成-21474836
8？
</div>
<div>
<br />
</div>
<div> &nbsp;答：由于Tow's Complement正负数的不对称性，直接写成-2147483648可
会导致溢出。在某些编译器中，对于-X的计算，先取数值X再取负数。2147483648超过了32位有符
整数的表示范围，造成了溢出，因此必须写成-2147483647-1。
</div>
<div>
<br />
</div>
<div>
<br />
</div>
<div>
<br />
</div>
<div>
    <h2> Errata for<br /> Computer Systems: A Programmer's Perspective (CS:APP) </h2>
</div>
<div>
<br />
</div>
<div>
<br />
</div>
<div>
    <a href="http://csapp.cs.cmu.edu/public/1e/public/errata.html">http://csapp.cs.cmu.edu/pu
lic/1e/public/errata.html</a>&nbsp;
<br />
```


numeric constant -2147483648 is handled in a peculiar way on a 32-bit, two's complement machine.

<p> The problem can be corrected by writing -2147483647-1, rather than -2147483648 in any C code. </p>

<h1> Description of Problem </h1> The ANSI C standard requires that an integer constant too large to be represented as a signed integer be ``promoted'' to an unsigned value. When GC encounters the value 2147483648, it gives a warning message: ``warning: decimal constant is too large that it is unsigned.'' The result is the same as if the value had been written 2147483648U.

<p> The compiler processes an expression of the form -X by first reading the expression X and then negating it. Thus, when the C compiler encounters the constant -2147483648, it first processes 2147483648, yielding 2147483648U, and then negates it. The unsigned negation of this value is also 2147483648U. The bit pattern is correct, but the type is wrong! </p>

<h1> Writing TMin in Code </h1> The ANSI C standard states that the maximum and minimum integers should be declared as constants INT_MAX and INT_MIN in the file limits.h. Looking at this file on an IA32 Linux machine (in the directory /usr/include), we find the following declarations:

<p>
 </p>

```
/* Minimum and maximum values a 'signed int' can hold. */#define INT_MAX 2147483647#define INT_MIN (-INT_MAX - 1)
```

<p> This method of declaring INT_MIN avoids accidental promotion and also avoids any warning messages by the C compiler about integer overflow. </p>

<p> The following are ways to write TMin_32 for a 32-bit machine that give the correct value and type, and don't cause any error messages: </p>

- -2147483647-1
- (int) 2147483648U
- 1<<31

 The first method is preferred, since it indicates that the result will be a negative number.

</div>