
Subject: SMP experiences with IDL

Posted by [Kirt Schaper](#) on Tue, 24 Feb 1998 08:00:00 GMT

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

Does anyone have first hand experience with IDL (preferably on a Linux box) running with multiple processors? Is there any speedup? (I'm talking about IDL v5.0 for Unix).

Our experiences with single processor Pentium/Linux boxes suggests that they are at least as fast, if not faster, than much more expensive HP, Dec and Sun boxes. Aside from the problem of being a little-endian architecture, I haven't been able to see the down-side yet.

More grist for the Linux performance mill...

Here are some timing results from a simple benchmark program (the program simply generated a 100x100x50 random float array and convolved it with a 10x10x10 kernel). I know that elapsed time is not a very precise benchmark, but the systems were all unloaded at the time of the test, and elapsed time is what makes a system usable or not.

```
; idl version 4.01
; SS10/51 (50MHz) ----- elapsed time = 59.1 seconds
; Dec 600 5/266 (266MHz) ----- elapsed time = 43.0 seconds
; HP 9000 C180 (180MHz) ----- elapsed time = 19.7 seconds
; Pentium Pro (200MHz), Linux -- elapsed time = 12.1 seconds
; Pentium II (300MHz), Linux --- elapsed time = 9.0 seconds
;
; idl version 5.0
; SS10/51 (50MHz) ----- elapsed time =138.6 seconds
; Pentium Pro (200MHz), Linux -- elapsed time = 45.0 seconds
; HP 9000 C180 (180MHz) ----- elapsed time = 33.7 seconds
; Pentium II (300MHz), Linux --- elapsed time = 31.3 seconds
```

I find several things interesting about the above experience.

(1) A 200MHz Pentium Pro box is running as fast as a (much more expensive, even with 50% academic discount) HP box. This is totally contrary to the published SPECfp95_base numbers (17.2 for the HP and 5.54 for the Pentium)

(2) RSI did something quite bad to the convolution function from v4 to v5.

kt
