
Subject: Re: Getting rid of large nested FOR loops
Posted by [Craig Markwardt](#) on Tue, 17 Aug 2004 22:51:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

paulo@craam.mackenzie.br (PJ Sim?es) writes:

> Hi,
>
> My code has three nested FOR loops (ops!)... well, the original code
> in
> FORTRAN has three nested FOR loops, but I'm re-writing (and improving)
> it to IDL. I've read the J.D. Smith's tutorial about vectorizing but
> the problem is that my inner loop is variable, it runs until the data
> converges... and I'm dealing with lots of data, something like this:
> for i=0,120...
> for j=0,100...
> for k=0,(variable: 2 to 1000) where I use BESELJ function (it
> seems to be the weak point) and maybe the memory usage will slow me
> down if I vectorize it... I know it sounds a little vague, but the
> code is rather simple, but I need
> it to be faster (it takes about 30 sec on a 1.2GHZ P4 / 1GB RAM). Is
> it better to vectorize the whole thing (hard stuff!), the inner loop
> or the outter loops? Or maybe another solution?

The place where vectorization will pay the most is where IDL is optimized for vectorization. For example, this code

```
for i = 0, 99 do a[i] = b[i] + c[i]
```

will always be slower than this code

```
a = b + c
```

because IDL optimizes vector operations. Also, the first code will be slower because it involves repeated indexing (the [i]'s), while the second code does not.

Since BESELJ can accept vector arguments, if you can reorder your loops so that you can call BESELJ with 100 or 120 values at once, then you may get a speedup.

However, you will then need to adjust your convergence criteria, i.e., you will have to vectorize that as well.

Good luck,
Craig

--

Craig B. Markwardt, Ph.D. EMAIL: craigmnet@REMOVEcow.physics.wisc.edu
Astrophysics, IDL, Finance, Derivatives | Remove "net" for better response
