Subject: Re: FOR loops and efficiency
Posted by Craig Markwardt on Sat, 23 May 2009 20:24:41 GMT
View Forum Message <> Reply to Message

On May 22, 2:24 pm, Christopher Thom <ct...@oddjob.uchicago.edu>
wrote:
> Quoth Craig Markwardt:
>
>> A FOR loop will only be slow(er) when the time spent executing the
>> loop overhead is much more than the time spent doing the computations
>> in one loop iteration. A simple test would be to execute a dummy
>> loop:
>> NMAX = 100000L
>> FOR I = 0L, NMAX do begin & dummy = 1
>> Keep raising the value of NMAX until the execute time of the loop is
>> perceptible. Don't bother trying to optimize loops smaller than this.
>
>> In your case, you are only doing ten iterations, and each iteration
>> does a lot of work, so you won't gain by removing the loop.
>
> I've heard this description about FOR loops a lot, but one general
> question I've never been able to answer is, "how do i know when my loops
> are doing enough work?". How do I know when my loop overhead is a large
> fraction of the time spent on an iteration?
>
> I guess the real underlying question here is recognising when to optimise,
> and when to simply move on to more important things. Does anyone have any
> rules of thumb to help guide this recognition?

I still stand by my rule of thumb. The problem with FOR loops is the
amount of time spent doing loop overhead stuff. If you run your loop
but *take all the calculations out*, and the total execution time is
not perceptible, then you probably won't gain by optimizing/
vectorizing.

Craig