
Subject: Re: FOR loops removal

Posted by [Wox](#) on Tue, 19 Aug 2008 13:43:40 GMT

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

On Tue, 19 Aug 2008 05:38:50 -0700 (PDT), loebasboy
<stijn.vdl@gmail.com> wrote:

```
>   FOR l = 0, n*2 DO BEGIN
>     temp = 0
>     FOR i = 0, max_y-1 DO BEGIN
>       FOR j = 0, max_x-1 DO BEGIN
>         jtemp = j + l
>         jtemp2 = j + n
>         temp = temp + (arr[i,jtemp] * arr [i,jtemp2])
>       ENDFOR
>     ENDFOR
>     output[l] = temp/(max_x*max_y)
>   ENDFOR
```

The code below is a start. Does this processing have a name? It feels familiar somehow. Btw, in IDL the first index of an array is the column and the second is the row. So in your case y are the columns and x are the rows. No problem with that off course, just check whether this is how you intended it.

```
n = 8
max_x = 5
max_y = 5
output = fltarr(2*n+1)
arr = findgen(max_y, 2*n+max_x) + 1

arr2=arr[0:max_y-1,n:max_x-1+n]
FOR l = 0, 2*n DO $
  output[l] = total(arr[0:max_y-1,l:max_x-1+l]*arr2)
output/=max_x*max_y
```

Subject: Re: FOR loops removal

Posted by [David Fanning](#) on Tue, 19 Aug 2008 13:51:26 GMT

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

Wox writes:

```
> The code below is a start. Does this processing have a name? It feels
> familiar somehow. Btw, in IDL the first index of an array is the
> column and the second is the row. So in your case y are the columns
> and x are the rows. No problem with that off course, just check
```

> whether this is how you intended it.

And, of course, it *will* matter (especially in loops) how you access the data. Always keep in mind that internally data is stored in memory in row order (column indices vary faster than row indices). Loops that keep this in mind are fast. Loops that don't keep this in mind are what IDL users complain about. :-)

Cheers,

David

--

David Fanning, Ph.D.

Fanning Software Consulting, Inc.

Coyote's Guide to IDL Programming: <http://www.dfanning.com/>

Sepore ma de ni thui. ("Perhaps thou speakest truth.")

Subject: Re: FOR loops removal

Posted by [Chris\[6\]](#) on Tue, 19 Aug 2008 20:04:04 GMT

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

> Can anybody tell me why removing one loop doesn't help in this case or
> what i'm doing wrong?

I think the main reason why your second code snippet isn't much faster than the first is because it's not a very good vectorization (the term to describe the elimination of loops in favor of array based operations).

Not being a computer scientist, I don't actually 'understand' what IDL does when it compiles and runs code. But the image in my mind is akin this old man I saw in a post office one time. He couldn't really hear that well, and kept (loudly) asking the post office clerk 'when the hell those Rat-a-ville stamps are coming in' (after eavesdropping for a while, I realized that he was actually sent by his wife to buy stamps from the movie 'Ratatouille'). After the clerk (repeatedly) told him that a) it was pronounced 'rat-uh-too-eee' and b) they would get them next week, the old man was on his way. The impressive thing was that this 90 year old man FLEW out of the post office when he was done. He was fast - like Lolo Jones fast.

How does this connect? IDL for loops are slow because the part of IDL that interprets your file is a fast but crotchety old man who can't hear you very well and may not even really be listening. Any time you tell him to do something, it takes him a while to interpret what you just said - much longer than other, less crotchety men. Once he figures out

what's going on, however, he's plenty fast (especially if you tell him to do something that he was already designed to do, for which he has been well optimized). Good vectorization, then, minimizes the number of instructions (e.g. iterations in a loop) while maximizing the amount of work to do with each instruction.

Your second loop doesn't have any fewer iterations than the first loop - it just gets rid of one nested for loop and increases the size of the previous loop. Un-nesting the loops helps a bit (looping the loop is two layers of interpretation. IDL has no patience for such tasks. He lived through the depression and fought the Germans), but you really aren't following the principle of 'loop less with bigger processing chunks in each step.'

Wox's code is the right way to vectorize your loop. It truly iterates fewer times, and gives IDL more to chew with each line of instruction. I wouldn't bother eliminating the L loop. As soon as you do some hefty processing in each iteration, the looping penalty goes away, and you don't need to worry about your vectorization creating huge temporary arrays and paying penalties in memory allocation.

As long as you don't have any loops where, at each iteration, you are simply accessing an element of an array, IDL should be pretty fast.

chris

Subject: Re: FOR loops removal
Posted by [loebasboy](#) on Wed, 20 Aug 2008 07:24:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Aug 19, 3:43 pm, Wox <nom...@hotmail.com> wrote:

> On Tue, 19 Aug 2008 05:38:50 -0700 (PDT), loebasboy
>

> <stijn....@gmail.com> wrote:

```
>>   FOR l = 0, n*2 DO BEGIN
>>     temp = 0
>>     FOR i = 0, max_y-1 DO BEGIN
>>       FOR j = 0, max_x-1 DO BEGIN
>>         jtemp = j + l
>>         jtemp2 = j + n
>>         temp = temp + (arr[i,jtemp] * arr [i,jtemp2])
>>       ENDFOR
>>     ENDFOR
>>     output[l] = temp/(max_x*max_y)
>>   ENDFOR
```

>

> The code below is a start. Does this processing have a name? It feels

```

> familiar somehow. Btw, in IDL the first index of an array is the
> column and the second is the row. So in your case y are the columns
> and x are the rows. No problem with that off course, just check
> whether this is how you intended it.
>
> n = 8
> max_x = 5
> max_y = 5
> output = fltarr(2*n+1)
> arr = findgen(max_y, 2*n+max_x) +1
>
> arr2=arr[0:max_y-1,n:max_x-1+n]
> FOR l = 0, 2*n DO $
>     output[l] = total(arr[0:max_y-1,l:max_x-1+l]*arr2)
> output/=max_x*max_y

```

Thank you for your code, it works rather well, maybe it seems familiar because it's a kind of autocorrelation that I'm calculating... .

I think I still need some vectorisation training to get IDL much faster, I've calculated a time profit of 14 h (that makes 8.5 h instead of 22.5 h), so I still have some FOR loops I can train on ;). Thanks for helping finding my way and the fast answers, I think I will definitely post again when I'm really stuck again ;).

Subject: Re: FOR loops removal
 Posted by [Jeremy Bailin](#) on Wed, 20 Aug 2008 11:50:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Aug 19, 9:43 am, Wox <nom...@hotmail.com> wrote:
 > On Tue, 19 Aug 2008 05:38:50 -0700 (PDT), loebasboy
 >

```

> <stijn....@gmail.com> wrote:
>>     FOR l = 0, n*2 DO BEGIN
>>         temp = 0
>>         FOR i =0,max_y-1 DO BEGIN
>>             FOR j=0,max_x-1 DO BEGIN
>>                 jtemp = j + l
>>                 jtemp2 = j + n
>>                 temp = temp + (arr[i,jtemp] * arr [i,jtemp2])
>>             ENDFOR
>>         ENDFOR
>>         output[l] = temp/(max_x*max_y)
>>     ENDFOR
>

```

```

> The code below is a start. Does this processing have a name? It feels
> familiar somehow. Btw, in IDL the first index of an array is the

```

```

> column and the second is the row. So in your case y are the columns
> and x are the rows. No problem with that off course, just check
> whether this is how you intended it.
>
> n = 8
> max_x = 5
> max_y = 5
> output = fltarr(2*n+1)
> arr = findgen(max_y, 2*n+max_x) +1
>
> arr2=arr[0:max_y-1,n:max_x-1+n]
> FOR I = 0, 2*n DO $
>     output[I] = total(arr[0:max_y-1,I:max_x-1+I]*arr2)
> output/=max_x*max_y

```

Following on that last version, I think we can *completely* get rid of the loop... though at the expense (as usual) of memory:

```

n = 8
max_x = 5
max_y = 5
arr = findgen(max_y, 2*n+max_x) +1

max_area = max_x*max_y
output = total( arr[rebin(lindgen(max_area),max_area,2*n+1) +
    max_y*rebin(reform(lindgen(2*n+1),1,2*n+1),max_area,2*n+1)] *
    rebin( (arr[:,n:max_x-1+n])[*, max_area,2*n+1), 1) / max_area

```

Whether that's actually faster will depend on how big max_x, max_y and n are, of course... it ends up internally storing a couple of max_x*max_y*(2*n+1) arrays, so if that is going to take you into swap then you're best off sticking with Wox's version. If that stays in physical memory, though, I bet this will win.

-Jeremy.