
Subject: Re: My code is way to slow.. How can I optimize!?
Posted by [David Fanning](#) on Mon, 28 Nov 2005 23:00:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

jelle.ferwerda@rmit.edu.au writes:

- > This is my first post here. I was desperately looking for info on the
- > web on optimizing IDL code, and came across this group. Looks like you
- > might be able to help me..
- >
- > For the past 3 years I have been working with IDL to get some basic
- > processing done for my phd work. Now I have been working on processing
- > some signals from a fieldspectrometer, using a bootstrapping routine. I
- > was wondering how I can optimize this code, as the current code takes
- > over a day for one run to finish.
- >
- > Is it possible to post some code on this forum? The code is quite long,
- > and I'd rather put it as a downloadable file (code is 440 lines).

Posting code implies you have money burning a hole in your pocket and you wish to spend it. Asking a question about a particular *aspect* of your code that seems particularly slow will get you a better response, probably. In any case, if we need to look at the code we are usually foolish enough to ask for it. :-)

Cheers,

David

--

David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: <http://www.dfanning.com/>

Subject: Re: My code is way to slow.. How can I optimize!?
Posted by [leatherback](#) on Mon, 28 Nov 2005 23:27:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi David,

Thanks for your reply. Nope, i do not have money burning a hole in my pocket. It is more that I know that the flowchart of the actual routine I am using will be published in a remote sensing journal soon anyway (Also already available online as chapter 6 in http://www.itc.nl/library/Papers_2005/phd/ferwerda.pdf). So I do not really care about people getting the code I used, It has served my

purpose, and I am now applying it to a dataset, where not this code, but the data it will be processing is crucial.

Anyway. It basically is a series of nested routines, within which I perform regressions on subsets, and generate statistics of these regressions. Since about 300 lines of the code fall within these loops, i am not particularly sure which parts take most of the time. It probably is part of an array concatenation part, which in a normal dataset is performed about 500,000 times. But it could also be the regression which is also performed about 500,000 times.

I am convinced that my code is sub-sub-sub optimal. Just do not know where to start optimizing it, as I am not a programmer, but a mere ecologist :-)

J.

Subject: Re: My code is way to slow.. How can I optimize!?
Posted by [David Fanning](#) on Mon, 28 Nov 2005 23:42:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

jelle.ferwerda@rmit.edu.au writes:

> Anyway. It basically is a series of nested routines, within which I
> perform regressions on subsets, and generate statistics of these
> regressions.

I assume you mean nested loops, which is the usual culprit with slow routines written in IDL. And you are right, we have (as a group) had some success speeding this kind of code up:

http://www.dfanning.com/code_tips/slowloops.html

> Since about 300 lines of the code fall within these loops,
> i am not particularly sure which parts take most of the time.

Since it takes a day to run, I can understand not wanting to run it with timing code. :-)

Still it would be worth while knowing what is slow and what is REALLY slow.

> It probably is part of an array concatenation part, which in a normal
> dataset is performed about 500,000 times. But it could also be the
> regression which is also performed about 500,000 times.

I expect you could find tips for improving the array concatenation

part up by searching through the archives of this newsgroup. If JD hasn't written about it 50 times it would be a small miracle.

> I am convinced that my code is sub-sub-sub optimal. Just do not know
> where to start optimizing it, as I am not a programmer, but a mere
> ecologist :-(

I'll take your word for its optimized nature. I'm just saying I don't care to look at 400 lines of code unless I'm also looking at your credit card details. :-)

But why speed it up? It's served its purpose, it sounds like. Maybe this is one mess that should just be buried and forgotten. Let those guys who read your paper figure it out! :-)

Cheers,

David

--

David Fanning, Ph.D.

Fanning Software Consulting, Inc.

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

Subject: Re: My code is way to slow.. How can I optimize!?
Posted by [Rick Towler](#) on Mon, 28 Nov 2005 23:43:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

jelle.ferwerda wrote:

> Anyway. It basically is a series of nested routines, within which I
> perform regressions on subsets, and generate statistics of these
> regressions. Since about 300 lines of the code fall within these loops,
> i am not particularly sure which parts take most of the time. It
> probably is part of an array concenation part, which in a normal
> dataset is performed about 500,000 times. But it could also be the
> regression which is also performed about 500,000 times.

Before you need to know *how* to optimize, you need to know *what* to optimize. Check out the docs on the PROFILER procedure. Use it to determine where your program is spending time. Then start asking the how questions.

> I am convinced that my code is sub-sub-sub optimal. Just do not know
> where to start optimizing it, as I am not a programmer, but a mere
> ecologist :-(

We'll try not to hold that against you. :) Heck, I was a wildlife

biologist once. Hummm. That probably explains a lot about my code.

-Rick

Subject: Re: My code is way to slow.. How can I optimize!?
Posted by [Mark Hadfield](#) on Mon, 28 Nov 2005 23:57:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

jelle.ferwerda@rmit.edu.au wrote:

> I am convinced that my code is sub-sub-sub optimal. Just do not know
> where to start optimizing it, as I am not a programmer, but a mere
> ecologist :-(

A mere ecologist?! How dare you intrude on this forum? This forum is frequented only by real scientists (ie. physical scientists) who discuss only serious, high-powered, real science questions like which T-shirts are the best babe magnets.

I suggest you start by reading the section in the IDL manual on profiling.

--

Mark Hadfield "Kei puwaha te tai nei, Hoesa tahi tatou"
m.hadfield@niwa.co.nz
National Institute for Water and Atmospheric Research (NIWA)

Subject: Re: My code is way to slow.. How can I optimize!?
Posted by [leatherback](#) on Tue, 29 Nov 2005 00:29:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

Profiler.. Brilliant!

So, now I know.. 76% of the time is taken up by transposing arrays, and summing results. I suppose I'll have a look and see whether I can transpose data outside the loops.

RE David:

Unfortunately: I am now using the routine in some actual data analysis. So I need the output, and this program will need to be run on 40 datasets, which would now take a month to run: Optimizing is a must. Profiling seems to help me focus on the problem area!

RE: Mark Hadfield:

I am so sorry! I'll try to maintain a low-profile on the forum, and be non-intrusive. I must admit.. My opinion on T-Shirts does not include

IDL shirts. I'd reckon they should give all users a free TShirts, considering the price they charge for simple licences. Maybe thinkgeek.com can produce a batch of them? Else..
<http://www.thinkgeek.com/tshirts/oreillyshirts/605c/>

Anyway.. You've pushed me in the right direction, I am sure,

Thanks!

Jelle.
