On Aug 14, 3:05 pm, David Fanning <n...@dfanning.com> wrote:
> Folks,
>
> A couple of weeks ago I was give some IDL code to,
> essentially, find blobs in a large image, do some
> processing on them, and return a result image. The
> images we want to work on are on the order of
> 8800 by 6800, pretty big.
>
> Naturally, the code contained FOR loops coming out
> the wazzoo and it was VERY slow. I don't know *how*
> slow because I turned the computer off when I noticed
> smoke coming out the back and I couldn't immediately
> locate the damn fire extinguisher.
>
> Because lousy code doesn't necessarily prevent you
> from writing scientific papers about your algorithm,
> I read the paper to get the gist and decided their
> algorithm was overly complicated and that I could do
> it a lot more simply (and in WAY fewer loops).
>
> So I code it up, using a much smaller image subset,
> and it was quicker than snot. When I thought I had
> all the kinks worked out of it, I decided to give'er
> a try on the larger images. It took about 48 minutes
> to run. Humm. Odd.
>
> But I couldn't immediately see where the problem was.
> I decided to play tennis instead of work on it anymore. :-)
>
> This morning I ran the PROFILER on the code, and saw it
> was spending most of its time in the WHERE function.
> I was using the WHERE to locate the indices of blobs I
> wanted to work with from the LABEL_REGION image. A couple
> of quick tests showed that as that as that image gets larger,
> the search time of WHERE goes up exponentially.
>
> "Humm, I wish LABEL_REGION just returned the indices like the
> HISTOGRAM function does," I thought.
>
> What!? After 20 years of working with IDL it finally sunk
> in. What I need is a HISTOGRAM!!!!
>
> Here is the bottom line. My program, which took 48 minutes

> to run yesterday, takes 10 *seconds* to run today.
>
> Sometimes you just gotta love IDL! :-)
>
> Cheers,
>
> David
> --
> David Fanning, Ph.D.
> Fanning Software Consulting, Inc.
> Coyote's Guide to IDL Programming (www.dfanning.com)
> Sepore ma de ni thui. ("Perhaps thou speakest truth.")

Yeah - I learned a little bit ago that anytime that you are looping
over a where function to pull out individual regions, a histogram with
reverse_indices is going to going to get the job done much faster