Subject: Re: nonuniform FFT
Posted by MKatz843 on Mon, 07 Apr 2003 07:09:58 GMT
View Forum Message <> Reply to Message

b_gom@hotmail.com (Brad Gom) wrote in message
> has anyone out there implemented a FFT routine that handles
> nonuniformly gridded samples?

If your samples aren't uniformly gridded, I'm not sure how you're going to get an "FFT" to work. However, there's no reason why you can't implement a simple (discreet) Fourier Transform minus the "fast" part.

If you know your non-uniform x values, then for arbitrary k values you could always Fourier transform your input array, A, like this

```
i = complex(1,0)
f_k = total( A * exp( i*k*x) )
```

That's just one Fourier component, and one dimension for x, but scaling that is trivial.
Of course, computing the FT in this way, one component at a time, should probably be called a SFT (Slow Fourier Transform).

Depending on your input/output needs, you may be able to implement a DFT (Discreet Fourier Transform) using IDL's vector/matrix math and then it can be Much faster. You have a vector of x values where you've sampled A. You have a vector of k values where you want to know the Fourier components. If memory serves, the outer product of those two vectors gives a matrix which is the k*x part of exp(i * k*x). The proper matrix multiplication by the input array A can yield your FT in a jiffy--much faster than using a for-do loop.

MKatz843

Subject: Re: nonuniform FFT
Posted by James Kuyper on Mon, 07 Apr 2003 14:10:31 GMT
View Forum Message <> Reply to Message

"M. Katz" wrote:
>
> b_gom@hotmail.com (Brad Gom) wrote in message
>> has anyone out there implemented a FFT routine that handles
>> nonuniformly gridded samples?
>
> If your samples aren't uniformly gridded, I'm not sure how you're
> going to get an "FFT" to work. ...

One popular technique I've seen is to take a non-uniform set of data,
and interpolate it using cubic splines to a uniformly spaced set of
points. Then an FFT is done on that. You lose some time and accuracy as
a result of doing the interpolation, but it's far faster than slow FFT,
and accurate enough for many purposes.

Subject: Re: nonuniform FFT
Posted by AJ on Mon, 07 Apr 2003 14:20:33 GMT
View Forum Message <> Reply to Message

Reading the online help, is the requested information not returned by the
keywords WK1 and WK2?

"Brad Gom" <b_gom@hotmail.com> wrote in message
news:bde24eff.0304041525.3cc5d5bb@posting.google.com...
> Hi All,
>
> has anyone out there implemented a FFT routine that handles
> nonuniformly gridded samples? The Numerical Recipes "fasper" routine
> seems to be one way to do it, but I don't want to write a DLM for it
> unless I have to. The internal IDL routine LNP_TEST is an
> implementation of the "fasper" code, but it returns only the maximum
> peak of the Lomb periodogram, and not the periodogram itself.
>
> Thanks
>
> Brad Gom

Subject: Re: nonuniform FFT
Posted by John Smith on Mon, 07 Apr 2003 18:47:26 GMT
View Forum Message <> Reply to Message

"Brad Gom" <b_gom@hotmail.com> wrote in message
news:bde24eff.0304041525.3cc5d5bb@posting.google.com...
> Hi All,
>
> has anyone out there implemented a FFT routine that handles
> nonuniformly gridded samples? The Numerical Recipes "fasper" routine
> seems to be one way to do it, but I don't want to write a DLM for it
> unless I have to. The internal IDL routine LNP_TEST is an
> implementation of the "fasper" code, but it returns only the maximum
> peak of the Lomb periodogram, and not the periodogram itself.
>
> Thanks

>
> Brad Gom

The maximum is exactly what the algorithm produces, and what it should be used for. You should not use it to reproduce the spectrum (even though it is a common approach by people who don't read the paper). The algorithm calculates confidence intervals for this maximum, which is quite useful.

If you carefully read the lomb scargle paper, you see that if performs a least squares fit to a SINGLE (complex) sinusoid. You can certainly apply that algorithm to a broad range of frequencies, but it is still only fitting one sinusoid at a time.
If you wanted to least squares fit the spectrum (all the sinusoids that an fft
would produce), you must do that simultaneously. (i.e. it is a "huge" matrix inversion).
The problem with that is it is quite likely to be singular.
Do not attempt a DFT instead of an FFT, as there is absolutely no difference between the two. A DFT has the same problem with gaps.

The problem is that when the sampling is not uniform, the "basis" sinusoids are no
longer orthogonal.

I would think the best approach is to either concentrate on a few (large amplitude)
sinusoids and employ the lomb scargle, or to interpolate the data (and perhaps
downsample).

The best solution is to take the data again, and sample regularly (which I would
guess is not possible).

Cheers
bob

Subject: Re: nonuniform FFT
Posted by b_gom on Mon, 07 Apr 2003 19:09:49 GMT
View Forum Message <> Reply to Message

Ah, right you are, AJ- I didn't look close enough. At any rate, I am concerned with speed issues. I know nothing will be as fast as the FFT, but the LNP_TEST function is too slow as it is (it looks like between 70 to a few hundred times slower). Also, I'm not exactly sure

how to interpret the output of the LNP as compared to an FFT.

It looks like I will need to decide between interpolating the data
onto a regular grid (with the errors this introduces and extra time it
takes) or implementing some reasonably fast nonuniform FT, which no
one seems to have done.

Thanks

Brad

"AJ" <a@nothing.com> wrote in message
news:<1049725229.195633@newsreader1.wirehub.nl>...
> Reading the online help, is the requested information not returned by the
> keywords WK1 and WK2?
>
> "Brad Gom" <b_gom@hotmail.com> wrote in message
> news:bde24eff.0304041525.3cc5d5bb@posting.google.com...
>> Hi All,
>>
>> has anyone out there implemented a FFT routine that handles
>> nonuniformly gridded samples? The Numerical Recipes "fasper" routine
>> seems to be one way to do it, but I don't want to write a DLM for it
>> unless I have to. The internal IDL routine LNP_TEST is an
>> implementation of the "fasper" code, but it returns only the maximum
>> peak of the Lomb periodogram, and not the periodogram itself.
>>
>> Thanks
>>
>> Brad Gom

Subject: Re: nonuniform FFT
Posted by stevenj on Mon, 07 Apr 2003 19:57:50 GMT
View Forum Message <> Reply to Message

b_gom@hotmail.com (Brad Gom) wrote in message
news:<bde24eff.0304041525.3cc5d5bb@posting.google.com>...
> has anyone out there implemented a FFT routine that handles
> nonuniformly gridded samples?

See: http://www.math.mu-luebeck.de/potts/nfft/

Subject: Re: nonuniform FFT
Posted by b_gom on Tue, 08 Apr 2003 18:04:49 GMT
View Forum Message <> Reply to Message

> The maximum is exactly what the algorithm produces, and what it
> should be used for. You should not use it to reproduce the spectrum
> (even though it is a common approach by people who don't read the
> paper). The algorithm calculates confidence intervals for this maximum,
> which is quite useful.

I've done some more careful reading and have come to the same conclusion..

> The best solution is to take the data again, and sample regularly (which I
> would guess is not possible).

Oh, if only it were possible. It seems like a large percentage of papers on this subject are written by astronomers, who have very limited access to the systems they are observing, and little opportunity to re-take data!

Thanks

Brad

Subject: Re: nonuniform FFT
Posted by b_gom on Tue, 08 Apr 2003 18:26:09 GMT
View Forum Message <> Reply to Message

Thanks for the link Steven. The page looks familiar- I must have seen it at some point and then completely forgotten about it. Unfortunately it is based on the FFTW library, which I have been having trouble compiling as a DLM for IDL. Do you know of anyone who has wrapped NFFT or FFTW for IDL?

I see that http://epsilon.nought.de/ has something in development, and I've tried the DLM posted by Stein Vidar Hagfors Haugan a couple years ago, but I get all sorts of strange crashes when I try to use it. I'm hoping someone with more windows C experience than me has already got it working!

Brad

stevenj@alum.mit.edu (Steven G. Johnson) wrote in message news:<27cfb406.0304071157.37f1b97b@posting.google.com>...
> b_gom@hotmail.com (Brad Gom) wrote in message news:<bde24eff.0304041525.3cc5d5bb@posting.google.com>...

>> has anyone out there implemented a FFT routine that handles
>> nonuniformly gridded samples?
>
> See: http://www.math.mu-luebeck.de/potts/nfft/
