



Published on *cisTEM* (<https://cistem.org>)

[Home](#) > Time estimate for movie alignment?

Time estimate for movie alignment?

Fri, 02/16/2018 - 01:42

#1

Rick_Baker

Time estimate for movie alignment?

Hi,

we recently got a new machine with enough CPUs and memory to run cisTEM on large datasets (32 CPUs, 512 GB RAM).

I tried to do movie alignment on 500 micrographs (K2, super res) but it seems to be going very slowly.

When the job is first submitted, the GUI freezes for a few minutes. After 10-15 minutes it starts to show periodic updates on the progress, but the initial estimates are ~22 hours for 500 micrographs.

Is this normal?

When I submit a single movie for alignment, the GUI still freezes. After a minute or two, it refreshes and the job is done.

Where could we be encountering problems? Or, is this typical behaviour?

Thanks!

Rick

timgrant

Hi Rick,

Hi Rick,

The time update for the movie alignment is only updated when the whole alignment is done, that means that at the beginning you need to wait until at least one result is sent back for this updates (btw the gui shouldn't freeze, it just won't update the progress bar?).

If you submit a single movie, and it takes a couple of minutes, then under ideal conditions it should take $(500 * 2) / 32 = \sim$ half an hour to align all the movies. Obviously, no conditions are perfect, but 22 hours sounds way too long. The most likely suspect here is disk access speed. What disk are the movies located on?

Also, the default run profile is set to the size of the number of detected cores, if you have hyperthreading, this will be 64 (+ 1 for the master = 65) in your case. This is overkill for the movie alignment, as it will vastly increase required memory and disk speed, so in this case you should set this down to 33. Do you have 32 physical cores, or 32 with hyperthreading?

Cheers,

Tim

Rick_Baker

Hi Tim,

Hi Tim,

We have 32 physical cores. Initially cisTEM created a profile with 65 threads, but I changed it to 33 before I did any of the analysis.

As for disk access, we have several Linux boxes that have a shared file system. I am not sure how our largest machine is connected to this filesystem, but it does have a 2 TB scratch disk. I copied a few movies to this drive and it seemed to run better.

1. 17 movies aligned in 3.5 minutes
2. 32 movies aligned in 5.1 minutes

What's the optimal disk location for the movies? All on scratch? If that's so, my movies will be too large for the scratch disk. The FAQ mentions using the mrc2tif program to compress the non gain corrected movies. If I do this, should I also compress the gain reference?

Thanks!

Rick

timgrant

Hi Rick,

Hi Rick,

No you don't need to compress the gain reference. We save the data as uncorrected compressed tif from the beginning, and basically never decompress it ourselves, it is all done in cisTEM, so you don't need huge disk space. We have machines with local SSD scratch, and we keep our data for active projects there for disk access speed.

Cheers,

Tim

[Log in](#) or [register](#) to post comments

Source URL: <https://cistem.org/time-estimate-movie-alignment>