



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

[Home](#) > only using one thread on the GPU

only using one thread on the GPU

Mon, 07/22/2024 - 08:32

#1

mrhansenable2

only using one thread on the GPU

hi,

I am trying out the 2DTM from cisTEM-2.0.0-alpha-183-e549844. I have a node on our cluster with four 2080Ti cards and 96 CPU. I set up a test run with a single micrograph, without modifying the settings at all. For the single micrograph it estimates it should run 4 hours.

It also gives this output, which I'm not sure how to interpret. I've tried changing settings "# of threads per copy" but it had no effect.

Warning, you are only using one thread on the GPU. Suggested minimum is 2. Check compute saturation using nvidia-smi -l 1

Any help is appreciated.

Also, do my reference and micrographs need to have the same pixel size or will cisTEM automatically scale the ref?

thanks

Jesse

jojoelfe

Changing the # threads per

Changing the # threads per copy should do the trick, just make sure you click the "save" button on the bottom right of the settings page.

To use all four GPUs you will have to setup 4 different commands, described here: <https://github.com/timothygrant80/cisTEM/issues/482>

And yes, reference and micrographs need to have the same pixelsize. In fact 2DTM is very sensitive to changes in the pixelsize, so make extra sure the calibration for the micrograph pixelsize is correct. We often make templates with slightly varying pixelsizes (say 0.97,0.98,0.99,1.00,1.01,1.02,1.03 for a nominal pixelsize of 1.00) whenever we get data from a new scope.

Tue, 07/23/2024 - 09:38 (Reply to #2)

mrhansenable2

thanks! I followed the

thanks! I followed the github thread but I think I'm doing it wrong. Let me explain.

I am not modifying the "manager command" at all (the top), but rather I click the "edit" button for the command and edit there. I have # of copies set to 1, # of threads set to 4. Then in the "command to run" I paste in the below text:

```
CUDA_VISIBLE_DEVICES=0 shared/cistem/2.0.0-alpha/$command  
CUDA_VISIBLE_DEVICES=1 shared/cistem/2.0.0-alpha/$command  
CUDA_VISIBLE_DEVICES=2 shared/cistem/2.0.0-alpha/$command  
CUDA_VISIBLE_DEVICES=3 shared/cistem/2.0.0-alpha/$command
```

It does indeed run, but I think it's still only using 1 GPU. The estimated time to complete the job (200 micrographs) is 3700 hours. Is that normal? When I have with 1 GPU or my attempt at 4 GPUs in each case it says 3700 hours.

thanks again

Jesse

Mon, 03/24/2025 - 13:57 (Reply to #3)

huwl

Same problem of ~5K hours running on 4 GPUs

Jesse,

I met the same problem as yours, did you solve this problem at the end? do you have any suggestions? I tried to just install the latest version in github, but I failed again. It shows this error **"/bin/sh: line 20: cd: src: No such file or directory" when I tapped "make". And about the bfactor usage, do you have any clue? Sorry for my whining.**

Wanlong

Tue, 03/25/2025 - 04:43 (Reply to #4)

huwl

writing an issue on github (527)

I wrote an issue on github with more details (<https://github.com/timothygrant80/cisTEM/issues/527>), any suggestion could be wrote there. Thanks!

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

Source URL: <https://cistem.org/only-using-one-thread-gpu?page=0>