

## [core22 0.0.13 in BETA testing : CUDA support at last!](#)

by [JohnChodera](#) » Sun Sep 20, 2020 9:22 am

After a lot of work from a large number of awesome folks, we've just rolled out core22 0.0.13 to BETA!

Currently, we're only testing project 17102, which is a collection of different RUNs (RUN0 through RUN17) that test different workloads in short WUs for benchmarking and stability analysis. As a reminder, you can set your CODE:

`SELECT ALL`client-type

to CODE:

`SELECT ALL`beta

to run through BETA projects (like 17102).

After that, we'll test more broadly in BETA projects before the full announce.

Full release notes are below:

### New features

- 0.0.13 adds support for CUDA for NVIDIA GPUs! This enables NVIDIA GPUs to run ~25% faster on normal projects and up to 50-100% faster for COVID Moonshot free energy sprints. Expect to see noticeably lower time per frame (TPF) numbers for these projects with significantly higher points per day (PPD). The core will fall back to OpenCL if CUDA cannot be configured.
- More useful debugging information is printed in the client logs when the CUDA or OpenCL platforms cannot be configured.
- The core now bundles dynamic libraries instead of compiling libraries statically to permit various platforms and plugins to be loaded dynamically. This was essential for enabling CUDA support.
- The OpenCL driver version is now reported in the client logs
- Checkpoints are now reported in the client logs
- The final integrator state is returned at the end of the WU, allowing for more complex simulation schemes that employ OpenMM CustomIntegrator

### Bugfixes

- Several OpenMM bugfixes for AMD GPUs were incorporated.
- Temperature monitor warnings have been removed.
- Requesting the core to shut down during core initialization should now be handled more gracefully
- On older clients, `-gpu`` is interpreted as a synonym for `-opencl-device``, and a warning to upgrade the client is issued

### Known issues

- We are noticing that some NVIDIA linux beta drivers sometimes fail when using the OpenCL platform with the Windows Subsystem for Linux 2 (WSL2) with CODE: `SELECT ALL`ERROR:exception: Error initializing context: clCreateContext (218) and are actively working with NVIDIA to track down and fix this issue. In the meantime,

we hope the CUDA support will provide a work-around for this issue. Please post here if you encounter this issue so we can work with you to more quickly test a fix.

- Older NVIDIA GPUs (Fermi and older) may not be able to make use of CUDA, but will fall back to OpenCL instead

## Acknowledgments

Enormous thanks to the following people for their help in producing the latest core:

- Peter Eastman, lead OpenMM developer (Stanford)
- Joseph Coffland, lead Folding@home developer (Cauldron Development)
- Adam Beberg, Principal Architect, Distributed Systems (NVIDIA) (and original co-architect of Folding@home nearly 21 years ago!)

Huge thanks to the extremely patient Folding@home volunteers who helped us test many builds of this core, especially PantherX, Anand Bhat, Jesse\_V, bruce, toTOW, davidcoton, mwroggenbuck, artoar\_11, rhavern, hayesk, muziqaz, Zach Hillard, \_r2w\_ben, bollix47, joe\_h, ThWuensche, and everyone else who tested the core and provided feedback.

### [Re: core22 0.0.13 in BETA testing : CUDA support at last!](#)

by [Jesse\\_V](#) » Sun Sep 20, 2020 9:39 am

[size=1.3em]John deserves his share of the credit too for an enormous amount of time and effort getting this version to run smoothly, reliably, and ready for beta testing. We have all been working on this and testing different builds for several months now.

During our testing, we've also been able to perform some benchmarking between Windows -> Linux and OpenCL -> CUDA. While a ~20-25% performance boost is fairly common for Linux when compared to Windows, there's a very big boost when comparing OpenCL on Windows to CUDA on Linux. You can see this in the time per frame (TPF) for specific projects. For example, on COVID Moonshot free energy calculation projects:

p13426 at 10% in Windows 10 on GTX 1080 Ti running OpenCL: TPF of 1:57

p13426 at 10% in Windows 10 on GTX 1080 Ti running CUDA: TPF of 1:16

p13426 at 10% in Debian on GTX 1080 Ti running OpenCL: TPF of 1:34

p13426 at 10% in Debian on GTX 1080 Ti running CUDA: TPF of 0:59

So if Windows-OpenCL is our baseline, then on this project, the speed improvements are:

Windows-CUDA: +53%

Debian-OpenCL: +24%

Debian-CUDA: +98%

CUDA average: +75%

The core will attempt to use CUDA when available. This translates into a very significant gain in PPD, as Adam Beberg found for another project across different Nvidia hardware in Linux:

TPF 73s - GTX 1080Ti running OpenCL/ 1.554 M PPD

TPF 57s - GTX 1080Ti running CUDA / 2.253 M PPD

TPF 49s - RTX 2080Ti running OpenCL/ 2.826 M PPD

TPF 39s - RTX 2080Ti running CUDA / 3.981 M PPD

TPF 36s - RTX 3080 running OpenCL / 4.489 M PPD

TPF 31s - RTX 3080 running CUDA / 5.618 M PPD

Finally, I'd like to highlight that core22 v0.0.13 currently seems incompatible with the `pocl-opencl-icd`

package in Ubuntu, which is depend on libpocl2. These package is not necessary to run F@h, but if you have them installed, please remove them before running the core. Other than that, if the core segfaults at startup or fails to start or run in any way, please post below!

### [Re: core22 0.0.13 in BETA testing : CUDA support at last!](#)

by [JohnChodera](#) » Mon Sep 21, 2020 1:38 pm

[size=1.3em]Heads up that we've just moved the longer 1343x projects to BETA and enabled them for core22 0.0.13 so that we can test full-length projects to make sure failure rates are within expected margins.

The screenshot displays the Folding@home client interface for a client named 'client0' which is 'Online' and 'Running'. The interface is divided into several sections:

- Client Information:** Client: client0 Online Running. Status, System Info, and Log tabs are visible.
- Folding Power:** A progress bar shows the current power level, set to 'Medium' between 'Light' and 'Full'. Controls for Fold, Pause, and Finish are present.
- Identity:** Name: husg, Team: 3213. Points Per Day: 2836072.
- Folding Slots:** A table showing the status of GPU slots. Slot 00 is 'Running' on 'gpu:0:GP102 [GeForce GTX 1080 Ti] 11380'.
- Work Queue:** A table showing the status of work units. Unit 01 is 'Running' with 49.30% progress and an ETA of 1 hour 21 mins. Credit is 316294 and PRCG is 1343.
- Selected Work Unit:** PRCG 13432 (253, 3, 29). Folding Slot ID 00, Work Queue ID 01, Status Running, Progress 49.30%, ETA 1 hours 21 mins, Base Credit 70660, Estimated Credit 316294, Estimated PPD 2817300, Estimated TPF 1 mins 37 secs, Project 13432, FahCore 0x22, Waiting On Attempts 0, Next Attempt Unknown, Assigned 2020-09-21T11:31:40Z, Timeout 2020-09-23T11:31:40Z, Expiration 2020-09-24T11:31:40Z, Work Server 18.188.125.154, Collection Server 140.163.4.231.

At the bottom left, there are '+ Add' and '- Remove' buttons. The bottom status bar shows 'Per Day: 7371220' and 'UTC: 2020-09-21T12:51:29Z'.

Client: local **Online** Running

Address

127.0.0.1:3633

192.168.1.100:

Status System Info Log

▶ Fold
⏸ Pause
▶▶ Finish

Folding Power

Light

Medium

Full



Identity

Name husq Team 3213

Points Per Day

4479203

Folding Slots

ID	Status	Description
00	<b>Running</b>	gpu:0:TU102 [GeForce RTX 2080 Ti] M 13448

Selected Work Unit

PRCG 13432 (113, 2, 36)  
 Folding Slot ID 00  
 Work Queue ID 01  
 Status **Running**  
 Progress  37.69%  
 ETA 1 hours 12 mins  
 Base Credit 70660  
 Estimated Credit 362802  
 Estimated PPD 4478019  
 Estimated TPF 1 mins 10 secs  
 Project 13432  
 FahCore 0x22  
 Waiting On Attempts 0  
 Next Attempt Unknown  
 Assigned 2020-09-21T11:57:53Z  
 Timeout 2020-09-23T11:57:53Z  
 Expiration 2020-09-24T11:57:53Z  
 Work Server 18.188.125.154  
 Collection Server 140.163.4.231

Work Queue

ID	Status	Progress	ETA	Credit	PRCG
01	<b>Running</b>	37.69%	1 hours 12 mins	362802	13432

+ Add - Remove

er Day: 7317750

UTC: 2020-09-21T12:48:07Z