

Using OpenCL on paragon.cs.fit.edu

About the machine

Paragon is available to all students and faculty. Users are limited to 100mb of disk space. Files for users may be cleared out after 1 semester of inactivity.

SSH Access

The primary way to access this machine is via SSH. Log in with your tracks username and password. This is accessible from both on and off campus machines. You may also use tools such as SFTP to move files back and forth between the machine. I recommend [PuTTY](#) for connecting via SSH and [Filezilla](#) for doing SFTP.

Samba Access

You may also use Samba(Windows file sharing) to access your files for editing. This will only work after you have logged in via SSH at least once. It is not accessible from off-campus. Please use SFTP or SSH when off-campus.

Windows

- Click Start->Run
- Type [\\paragon.cs.fit.edu](https://paragon.cs.fit.edu)
- You may need to enter your username as 'fltech\username' and use your tracks password

OSX

- Finder->Go->Connect to Server.
- Enter smb://paragon.cs.fit.edu/
- When prompted enter your username as 'fltech\username' and your tracks password.

A Simple Sample Project

There are some sample files you can use to test the setup and as a base for building your own programs. You can copy them from the /opt/opencl/sample/ directory.

```
kjohns07@paragon:~$ cp -R /opt/opencl/sample ~/sample
kjohns07@paragon:~$ cd sample
kjohns07@paragon:~/sample$ ls
kernel.cl  main.cpp  Makefile
```

A brief description of the files

main.cpp	The Main Program
kernel.cl	The OpenCL kernel(this is code that runs on the GPU)
Makefile	A set of instructions for building and submitting your code; used by make

Compile the code

To compile the code, simply run 'make' from the sample directory.

```
kjohns07@paragon:~/sample$ make
g++ main.cpp -lOpenCL -o main

Run your job on the GPU with 'make submit'
```

Run your program

Running with 'make submit' afterwards results in:

```
kjohns07@paragon:~/sample/sample$ make submit
Platform number is: 1
Platform is by: NVIDIA Corporation
Hello World
```

This Makefile can ease submission of your program into the job queue. Running 'make submit' will send your job to a different machine that contains the actual GPU your program will use to run. Any output will be sent to your terminal. Attempting to run your program directly will result in failure since paragon does not actually have a GPU present.

```
kjohns07@paragon:~/sample/sample$ ./main
ERROR: cl::Platform::get (-1)
```

Detailed Configuration/Information

There is some information you should be aware of if you do not wish to build off of the sample project.

Compiling

When compiling your code, you need to tell it to link against OpenCL. This is done with the '-lOpenCL' linker command. For example, this is the command used to compile the sample program:

```
g++ main.cpp -lOpenCL -o main
```

Execution

In order to run your code, you need to submit it to the job queue so that it can run on the host containing a real GPU. By default a job has a time-limit of 1 minute. To submit a job use the `srun` command. A sample invocation would look like this:

```
srun ./main
```

Checking the status of a job

To check the status of a job once it has been submitted, you can use the 'squeue' command. Running it with no arguments will show the current status of the job queue.

A simple CUDA example

Copy the files from `/opt/opencl/cuda_sample/`. This uses a similar makefile to the one above, but compiles with `nvidia cuda` support. Please note: the `nvidia-toolkit` installed on this machine is version 4.0(Somewhat outdated)

If it is working correctly, it should print "Hello World!". If you run it on paragon(i.e. without using `srun` or `make submit`), it will output "Hello Hello".

Questions/Comments

If you have any questions, comments, concerns, or anything else relating to this system please email `sysadmin@cs.fit.edu`. Please include as many details as possible if you are reporting an error to save time.