Homework 2

Please provide the report of your study:

Provide system information of your computer system (CPU and dedicated GPU). If you do not have it, please try to get access to the machine that have:

-CPU model
-CPU Memory size
-GPU model
-GPU Memory size
-PCIe generation

Calculate theoretical performance of your GPU on the following categories:

- Peak Flops
- Peak memory bandwidth
- Peak PCI bus bandwidth
- Energy for 1 day

Run the following benchmark and report the actual performance of your system:

- The Babel Stream benchmark

1. **git clone git@github.com:UoB-HPC/BabelStream.git**
2. **make -f CUDA.make**
3. **./cuda-stream**

- Roofline performance model of GPU

1. Get the roofline toolkit

   git clone https://bitbucket.org/berkeleylab/cs-roofline-toolkit.git

2. **cd cs-roofline-toolkit/Empirical_Roofline_Tool-1.1.0**
3. **cp Config/config.voltar.uoregon.edu Config/config.V100_gpu**
4. edit Config/config.V100_gpu and change the following
   a) ERT_RESULTS Results.V100_gpu
   b) ERT_PRECISION FP64
   c) ERT_NUM_EXPERIMENTS 5
5. Run tests ./ert Config/config.V100_gpu
6. View Results.config.V100_gpu/Run.001/roofline.ps
7. **cp Config/config.odinson-ocl-fp64.01 Config/config.Vega20_gpu**
8. edit Config/config.Vega20_gpu
   a) ERT_RESULTS Results.Vega20_gpu
   b) ERT_CFLAGS -O3 -x c++ -std=c++11 -Wno-deprecated-declarations -I<path to OpenCL headers>
   c) ERT_LDLIBS -L<path to OpenCL libraries> -lOpenCL
9. Run tests ./ert Config/config.Vega20_gpu
10. View Results.config.Vega20_gpu/Run.001/roofline.ps

And make the roofline plot (see example below).

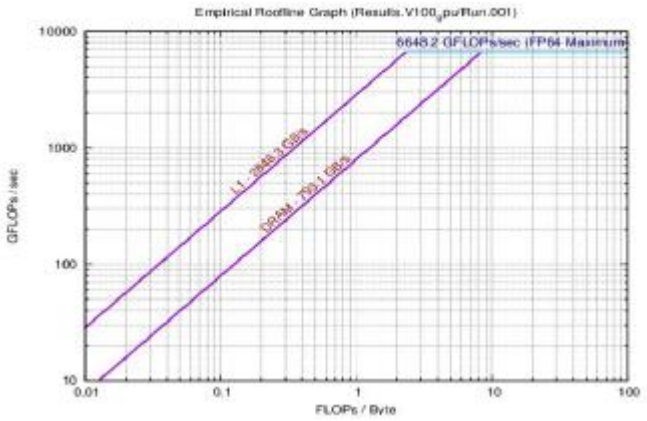

- Mixbench performance tool to obtain the peak flop rate

1. Get the mixbench code

   git clone https://github.com/ekondis/mixbench.git

2. Check for CUDA or OpenCL and install if necessary
3. **cd mixbench;** edit Makefile
   a) fix the path to the CUDA and/or OpenCL installations
   b) set the executables to build
4. you can override the path to the CUDA installation with
   a) **make CUDA_INSTALL_PATH=<path>**
5. Run either **./mixbench-cuda-ro or ./mixbench-ocl-ro**

   And plot the graph below:

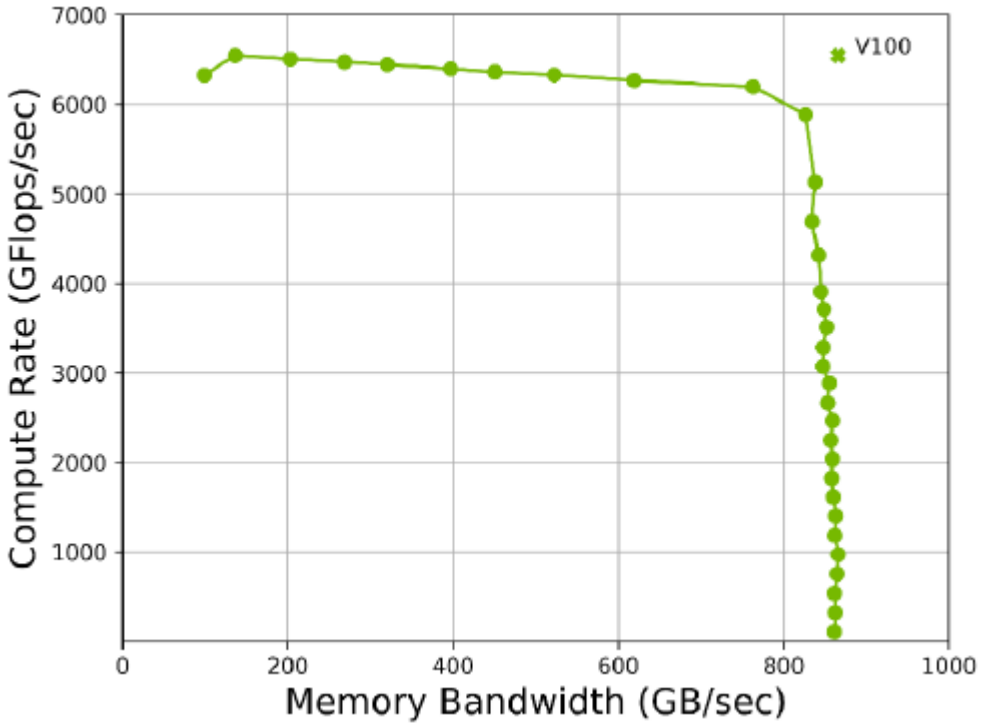

- Monitor GPU Consumption over application lifetime (you can choose to run any GPU applications, but please provide program information)

```
nvidia-smi dmon -i 0 --select pumct -c 65 --options DT --filename gpu_monitoring.log
&
```

Below is the example of the output

```
#Date     Time        gpu   pwr gtemp mtemp    sm   mem enc dec     fb bar1 mclk pclk rxpci txpci
#YYYYMMDD  HH:MM:SS    Idx     W     C     C     %     %   %   %     MB   MB  MHz  MHz  MB/s  MB/s
  20191221  21:36:47    0    64    43    41    24    28   0   0      0    0  877 1530     0     0
  20191221  21:36:48    0   176    44    44    96   100   0   0  11181    0  877 1530     0     0
  20191221  21:36:49    0   174    45    45   100   100   0   0  11181    0  877 1530     0     0
. . . and so on
```

Please plot the graph to report the final results in this format: