

March 2020

March 4 2020


Trying to get tesla m40 into z620

UEFI is required!! I converted the z620 machine (mwanafunzi) from legacy boot to uefi boot and switched the gpu from legacy to efi support. [Here](#) is the pastebin output from that.

This worked very well, with the tesla drivers, I got the m40 gpu to show up with all 24 GB of VRAM.

```
      .-/+00SSSS00+/- .
    `:+SSSSSSSSSSSSSSSSSS+:`
  +-SSSSSSSSSSSSSSSSSSyySSSS+-
  .OSSSSSSSSSSSSSSSSSSdMMMNySSSSO.
  /SSSSSSSSSSShdmNNmmyNMMMMhSSSSSS/
  +SSSSSSSSShmydMMMMMMMMNddddySSSSSSSS+
  /SSSSSSSSShNMMMyhhyyyyhmNMMNhSSSSSSSS/
  .SSSSSSSSdMMMNhSSSSSSSSShNMMMdSSSSSSSS.
+SSSSShhyNMMNySSSSSSSSSSSSyNMMMySSSSSSSS+
OSSyNMMNyMMhSSSSSSSSSSSSShmmhSSSSSSSSO
OSSyNMMNyMMhSSSSSSSSSSSSShmmhSSSSSSSSO
+SSSSShhyNMMNySSSSSSSSSSSSyNMMMySSSSSSSS+
.SSSSSSSdMMMNhSSSSSSSSShNMMMdSSSSSSSS.
/SSSSSSShNMMMyhhyyyhdNMMNhSSSSSSSS/
+SSSSSSSSdmydMMMMMMMMdddySSSSSSSS+
/SSSSSSSSShdmNNNNmyNMMMMhSSSSSS/
.OSSSSSSSSSSSSSSSSSSdMMMNySSSSO.
  +-SSSSSSSSSSSSSSSSSSyySSSS+-
    `:+SSSSSSSSSSSSSSSSSS+:`
      .-/+00SSSS00+/- .

root@mwanafunzi
-----
OS: Ubuntu 18.04.4 LTS x86_64
Host: HP Z620 Workstation
Kernel: 5.3.0-40-generic
Uptime: 25 secs
Packages: 3166
Shell: bash 4.4.20
Terminal: run-parts
CPU: Intel Xeon E5-2650 0 (16) @ 2.800GHz
GPU: NVIDIA Tesla M40
GPU: NVIDIA GeForce GTX 1070
Memory: 716MiB / 48123MiB
```



I tested this out with allennlp and it worked quite well. The m40 trained an rnn about as fast as my 1070's. This is the worst case for this comparison because the 1070 has a higher clock rate but fewer cuda cores and rnn's are difficult to parallelize.

The additional vram allowed me to go up to 128 for the batch size.

However, my cooling solution for the m40 was insufficient. After about 1 epoch of training (5 minutes of heavy usage) the temp exceeded 80 C and I had to end the workload. I was using a single NF 4x20A fan from noctua. However, these only provide 5 CFM of airflow. I purchased a 2 pack of delta 40mm fans that achieve 10 CFM each. This should be enough airflow for operation. While the noise level is going up, it is only increasing from 17 dba to 35 dba (per the documentation for the respective products). since the z620 case fans are about 35 dba, this noise difference shouldn't be very noticable.

AMD GPU build (RX580)

batch_size 32 tensorflow 1.14

resnet50: 49.93 images per second

resnet 152: 20.83 images per second

inception v3: 20.03 images per second

Revision #4

Created Wed, Mar 4, 2020 4:40 PM by [kenneth](#)

Updated Mon, Mar 9, 2020 1:36 AM by [kenneth](#)