

Install pytorch-rocm on bare metal opensuse Tumbleweed

These instructions are adopted from section 4 of [this page](#).

pytorch commit 9d1138afec26a4fe0be74187e4064076f8d45de7 worked for some stuff but pieces of allennlp are incompatible because this is pytorch v1.7.0a0+9d1138a.

I tried with 1.5 and 1.5.1 several times but it wasn't working on opensuse. I could have sworn it worked on ubuntu though.

Add rocm repos

Using the files provided by AMD for SLES SP1 per [the official instructions](#).

```
sudo zypper install dkms
sudo zypper cleansudo zypper addrepo --no-gpgcheck http://repo.radeon.com/rocm/zyp/zypper/rocm
sudo zypper ref
sudo zypper install rocm-dkms
sudo reboot
```

Modify `/etc/modprobe.d/10-unsupported-modules.conf` to have

```
allow_unsupported_modules 1
```

Then run the following, though it's probably not strictly necessary on tumbleweed

```
sudo modprobe amdgpu
```

Add your user to the video group

```
usermod -a -G video <username>
```

Verify everything is working by examining the output of `rocm_info` and make sure your gpu is listed.

Create a virtual environment

```
virtualenv -p python3 ~/venvs/torch  
source ~/venvs/torch/bin/activate
```

Install pytorch prerequisites

```
sudo zypper in glog-devel python3-pip libopenblas-devel libprotobuf-devel libnuma-devel  
libpthread-stubs0-devel libopencv-devel git gcc cmake make lmdb-devel libleveldb1 snappy-  
devel hiredis-develsudo zypper in rocm-dev rocm-libs miopen-hip hipsparse rocthrust hipcub  
rccl roctracer-dev
```

Fix issues with cmake files for rocm

```
sed -i 's/find_dependency(hip)/find_dependency(HIP)/g'  
/opt/rocm/rocspase/lib/cmake/rocspase/rocspase-config.cmake sed -i  
's/find_dependency(hip)/find_dependency(HIP)/g' /opt/rocm/rocfft/lib/cmake/rocfft/rocfft-  
config.cmake sed -i 's/find_dependency(hip)/find_dependency(HIP)/g'  
/opt/rocm/miopen/lib/cmake/miopen/miopen-config.cmake sed -i  
's/find_dependency(hip)/find_dependency(HIP)/g' /opt/rocm/roclblas/lib/cmake/roclblas/roclblas-
```

```
config.cmake sed -i 's/find_dependency(hip)/find_dependency(HIP)/g'  
/opt/rocm/rccl/lib/cmake/rccl/rccl-config.cmake sed -i  
's/find_dependency(hip)/find_dependency(HIP)/g'  
/opt/rocm/hipsparse/lib/cmake/hipsparse/hipsparse-config.cmake
```

Clone the repo

```
git clone https://github.com/pytorch/pytorch.git  
cd pytorch  
git checkout v1.5.0  
git submodule update --init --recursive
```

Build

This process will take a while (2-3 hours)

```
export RCCL_DIR="/opt/rocm/rccl/lib/cmake"  
python tools/amd_build/build_amd.pyUSE_ROCM=1 USE_LMDB=1 USE_OPENCV=1 MAX_JOBS=4 python  
setup.py install
```

Install allennlp

```
pip install allennlp  
pip uninstall torch  
# rebuild torch (very fast this time)USE_ROCM=1 USE_LMDB=1 USE_OPENCV=1 MAX_JOBS=4 python  
setup.py install
```

Verify installation

Make sure you get a non-zero value (should correspond to the number of GPUs).

```
python
>>import torch
>>torch.cuda.device_count()
```

Revision #5

Created Thu, Jul 2, 2020 1:22 PM by [kenneth](#)

Updated Thu, Jul 9, 2020 12:17 AM by [kenneth](#)