CS 491 Undergraduate Research

Sergio Servantez
Computer Science Department
Illinois Institute of Technology
Chicago, Illinois 60616
sservantez@hawk.iit.edu
Advisor: Dr. Zhiling Lan

Report 3: 3/17/19

Over the last couple weeks I have been able to submit several DNN training jobs on Theta, but have ran into several issues. Early on I ran into some issues with the data science modules not loading correctly, but those issues have been resolved. The main issue has been the OS run-time error displayed below. After inspecting the TensorFlow code, I have been able to locate the problem. In the DNN training code I am loading the MNIST dataset via the Keras library. However, the Keras library is attempting to fetch the dataset from an external Google server. Theta compute nodes do not allow users to connect to external URLs. To resolve this problem, I have obtained the MNIST dataset and installed it locally on Theta. I am now working on changing the code to consume this local dataset and properly slice the data among the ranks. Regarding the selection of a profiling tool for internode communication, Dr. Lan was able to confirm that Horovod has a built-in profiler. We will be using this tool as we move forward. Over the next couple weeks, my main goal is to resolve the MNIST dataset issue so we can begin collecting data on internode communication.

```
Intel Python 3.5 version 2017.0.035 loaded
-------------------------------
Tensorflow version 1.10.1 loaded
Intel Python 3.5 version 2017.0.035 loaded
Horovod version 0.13.11 loaded
Intel Python 3.5 version 2017.0.035 loaded
Intel Python 3.5 version 2017.0.035 loaded
-------------------------------
Tensorflow version 1.10.1 loaded
Traceback (most recent call last):
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/urllib/request.py", line 1254, in do_open
    h.request(req.get_method(), req.selector, req.data, headers)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 1106, in request
    self._send_request(method, url, body, headers)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 1151, in _send_request
    self.endheaders(body)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 1102, in endheaders
    self._send_output(message_body)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 934, in _send_output
    self.send(msg)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 877, in send
    self.connect()
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 1252, in connect
    super().connect()
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/http/client.py", line 849, in connect
    (self.host,self.port), self.timeout, self.source_address)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/socket.py", line 711, in create_connection
    raise err
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/socket.py", line 697, in create_connection
    sock = socket(af, socktype, proto)
  File "/opt/intel/python/2017.0.035/intelpython35/lib/python3.5/socket.py", line 134, in __init__
    _socket.socket.__init__(self, family, type, proto, fileno)
OSError: [Errno 97] Address family not supported by protocol
```