

This notebook is designed to show the problem with interchip communication on nengo-loihi

If any of the cells time out interrupt the cell and start the notebook over to see.

Start out by importing necessary modules

```
In [1]: from tensorflow.keras.layers import *
from nengo_loihi.hardware.allocators import PartitionInterchip
import tensorflow as tf
import nengo_dl
import nengo
import numpy as np
import nengo_loihi
import sys
```

```
2022-01-11 05:32:32.888271: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcudart.so.11.0'; dlerror: libcudart.so.11.0: cannot open shared object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-01-11 05:32:32.888311: I tensorflow/stream_executor/cuda/cudart_stub.cc:29] Ignore above cudart dlerror if you do not have a GPU set up on your machine.
2022-01-11 05:32:49.539888: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcudart.so.11.0'; dlerror: libcudart.so.11.0: cannot open shared object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-01-11 05:32:52.552207: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcuda.so.1'; dlerror: libcuda.so.1: cannot open shared object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-01-11 05:32:52.552263: W tensorflow/stream_executor/cuda/cuda_driver.cc:269] failed call to cuInit: UNKNOWN ERROR (303)
```

Construct a simple convolutional network

```
In [2]: # define a convolutional network
def simple_neural_network():
    inputs = Input(shape=(64, 64, 3))
    spiking_input = Activation(tf.nn.elu)(inputs) # this will be replaced by SpikingRectified
    conv_out1 = Conv2D(32, (5, 5), strides=(2, 2), padding='valid', activation=tf.nn.relu,
    conv_out2 = Conv2D(64, (3, 3), strides=(2, 2), padding='valid', activation=tf.nn.relu,
    conv_out3 = Conv2D(128, (2, 2), strides=(2, 2), padding='valid', activation=tf.nn.relu
    conv_out4 = Conv2D(256, (2, 2), strides=(2, 2), padding='valid', activation=tf.nn.relu

    flat_out = Flatten()(conv_out4)
    output = Dense(4, activation=None, name="dense", use_bias = False)(flat_out)

    model = tf.keras.Model(inputs=inputs, outputs=output)
    return model
```

convert network to spiking implementation

```
In [3]: ann_model = simple_neural_network()
nengo_converter = nengo_dl.Converter(
    ann_model,
    swap_activations={tf.nn.relu: nengo_loihi.LoihiSpikingRectifiedLinear(), # this is ou
```

```
)
tf.nn.elu: nengo.SpikingRectifiedLinear()},
```

```
2022-01-11 05:32:57.467913: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcuda.so.1'; dlerror: libcuda.so.1: cannot open share
d object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-01-11 05:32:57.467962: W tensorflow/stream_executor/cuda/cuda_driver.cc:269] failed c
all to cuInit: UNKNOWN ERROR (303)
2022-01-11 05:32:57.467993: I tensorflow/stream_executor/cuda/cuda_diagnostics.cc:156] ker
nel driver does not appear to be running on this host (ncl-edu): /proc/driver/nvidia/versi
on does not exist
```

Specify first layer as running off chip

```
In [4]: with nengo_converter.net as net:
        nengo_loihi.add_params(net) # allow on_chip to be set
        net.config[nengo_converter.layers[ann_model.layers[1]].ensemble].on_chip = False
```

Define our input to the SNN

This is fake imagery data with extremely low spike rates

```
In [6]: nengo_input = nengo_converter.inputs[ann_model.layers[0]]
        with nengo_converter.net as net:
            nengo_input.output = nengo.processes.PresentInput(
                np.random.random((10, 64, 64, 3)) * 300, presentation_time=.25
            )
```

Specify our block sizes

If a block shape of (n, m, c) is specified it means that for that specific layer output of the neural network it will try to allocate neurons with this block shape. The following inequality holds due to hardware constraints $n \cdot m \cdot c \leq 1024$

By decreasing the dimensions of n, m, or c we can force our neural network to be partitioned inefficiently and use more cores.

In the next cell I use efficient block shapes so you can see that the neural network successfully functions

```
In [7]: block_sizes = [None, (16, 16, 4), (8, 8, 16), None]
        conv_layers = [layer for layer in ann_model.layers if "conv" in str(layer).lower()]
        for layer, block_size, layer_idx in zip(conv_layers, block_sizes, range(len(conv_layers))):
            if block_size == None: #
                continue
            output_shape = tuple(layer.output.shape[1:])
            with nengo_converter.net as net:
                net.config[net.ensembles[layer_idx+1]].block_shape = nengo_loihi.BlockShape(block_size)
```

Try Inference on the network for a set number of epochs

If this cell fails with an srun timeout interrupt and restart it until it works

```
In [8]:
```

```
nengo_loihi.hardware.interface.HostSnip.recv_timeout = 10.0 # Default is 0.01 (10 ms)
nengo_loihi.hardware.interface.HostSnip.recv_retries = 100 # Default is 10
with nengo_loihi.Simulator(net, remove_passthrough=False, precompute = True, target="loih
print(f"Precompute = {loihi_sim.precompute}")
# iterate over the timesteps
loihi_sim.run(.25 * 10)
# for i in range(10): # iterate over the batches
#     loihi_sim.clear_probes() # prevents the buffer from filling up and crashing our
#     loihi_sim.run(.25 * 10)
```

77

```
INFO:DRV: SLURM is being run in background
INFO:DRV: Connecting to 10.212.98.108:35235
INFO:DRV: Host server up.....Done 0.55s
INFO:DRV: Encoding axons/synapses....Done 12.28s
INFO:DRV: Compiling Embedded snips....Done 0.39s
INFO:DRV: Compiling MPDS Registers....Done 0.76ms
INFO:HST: Args chip=0 cpu=0 /homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/si
te-packages/nxsdk/driver/compiler/../../../../temp/1641908099.4190736/launcher_chip0_lmt0.bi
n --chips=1 --remote-relay=1
INFO:HST: Nx...
INFO:DRV: Booting up.....Done 0.69s
INFO:DRV: Encoding probes.....Done 3.41ms
Precompute = True
INFO:DRV: Transferring probes.....Done 0.03s
INFO:DRV: Configuring registers.....Done 2.55s
INFO:DRV: Transferring spikes.....Done 226.36s
INFO:DRV: Executing.....Done 31.48s
INFO:DRV: Processing timeseries.....Done 1.24s
INFO:DRV: Executor: 2500 timesteps.....Done 261.67s
INFO:HST: chip=0 cpu=0 halted, status=0x0
```

In [10]: `loihi_sim.data[nengo_converter.net.probes[-1]].shape`

Out[10]: (2500, 4)

Now we redo the block size to encourage the neural network to fit on two chips. As you will see however, it will error instead

```
In [11]: block_sizes = [None, (8, 8, 4), (4, 4, 16), None]
conv_layers = [layer for layer in ann_model.layers if "conv" in str(layer).lower()]
for layer, block_size, layer_idx in zip(conv_layers, block_sizes, range(len(conv_layers)))
    if block_size == None: #
        continue
    output_shape = tuple(layer.output.shape[1:])
    with nengo_converter.net as net:
        net.config[net.ensembles[layer_idx+1]].block_shape = nengo_loihi.BlockShape(block
```

```
In [12]: nengo_loihi.hardware.interface.HostSnip.recv_timeout = 1.0 # Default is 0.01 (10 ms)
nengo_loihi.hardware.interface.HostSnip.recv_retries = 100 # Default is 10
with nengo_loihi.Simulator(net, remove_passthrough=False, precompute = True, target="loih
print(f"Precompute = {loihi_sim.precompute}")
# iterate over the timesteps
loihi_sim.run(.25 * 10)
```

```

76
73
INFO:DRV: SLURM is being run in background
INFO:DRV: Connecting to 10.212.98.108:46755
INFO:DRV: Host server up.....Done 0.26s
INFO:DRV: Encoding axons/synapses.....Done 12.53s
INFO:DRV: Compiling Embedded snips....Done 1.61s
INFO:DRV: Compiling MPDS Registers....Done 0.84ms
INFO:HST: Args chip=0 cpu=0 /homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/compilers/../../../../temp/1641908859.0948048/launcher_chip0_lmt0.bin --chips=2 --remote-relay=1
INFO:HST: Args chip=1 cpu=0 /homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/compilers/../../../../temp/1641908859.0948048/launcher_chip1_lmt0.bin --chips=2 --remote-relay=1
INFO:HST: Nx...
INFO:DRV: Booting up.....Done 1.23s
INFO:DRV: Encoding probes.....Done 4.49ms
Precompute = True
INFO:DRV: Transferring probes.....Done 0.03s
INFO:DRV: Configuring registers.....Done 2.35s
INFO:DRV: Transferring spikes.....Done 225.47s
INFO:HST: srun: Force Terminated job 1350086
INFO:HST: srun: Job step aborted: Waiting up to 32 seconds for job step to finish.
INFO:HST: slurmstepd: error: *** STEP 1350086.0 ON ncl-ext-ghrd-04 CANCELLED AT 2022-01-11T06:07:46 DUE TO TIME LIMIT ***
INFO:DRV: Executing.....Error 736.79s
INFO:DRV: Executor: 2500 timesteps.....Error 964.66s
INFO:HST: srun: error: ncl-ext-ghrd-04: task 0: Terminated

```

```

-----
_InactiveRpcError Traceback (most recent call last)
/tmp/ipykernel_1093282/4203283615.py in <module>

```

```

6
----> 7     loihi_sim.run(.25 * 10)
8

```

```

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/simulator.py in run
(self, time_in_seconds)

```

```

348     )
--> 349     self.run_steps(steps)
350

```

```

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/simulator.py in run
_steps(self, steps)

```

```

361
--> 362     self._runner.run_steps(steps)
363     self._n_steps += steps

```

```

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/simulator.py in loi
hi_precomputed_host_pre_and_host(self, steps)

```

```

558     self._host2chip(self.loihi)
--> 559     self.loihi.run_steps(steps, blocking=True)
560     self._chip2host(self.loihi)

```

```

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardware/interface.
py in run_steps(self, steps, blocking)

```

```

218     # start the board running the desired number of steps
--> 219     self.nxsdk_board.run(steps, aSync=not blocking)
220

```

```

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/arch/base/nxboard.py in r
un(self, numSteps, aSync, maxTimeInterval, generateCfg, cfgPath, partition)

```

```

282         with setEnvWithinContext(PARTITION=partition):
--> 283             self._run(
284                 numSteps=numSteps,

```

```

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/arch/base/nxboard.py in _
run(self, numSteps, aSync, traceDirectory)
    256             self, traceDirectory=traceDirectory)
--> 257         self.executor.start(numSteps, aSync)
    258

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/executor.py in sta
rt(self, numSteps, aSync)
    83         if not aSync:
---> 84             self.finish()
    85

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/executor.py in fin
ish(self)
    120         if self._state in (ExecutionState.RUNNING, ExecutionState.PAUSED):
--> 121             self._wait()
    122             self._notifyListeners(ExecutionEventEnum.POST_EXECUTION)

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/executor.py in _wa
it(self)
    127         with timedContextLogging("Executing", NxSDKLogger.NXDRIVER):
--> 128             self._executor_service.waitExecution(empty)
    129

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/grpc/_channel.py in __call__(se
lf, request, timeout, metadata, credentials, wait_for_ready, compression)
    945             wait_for_ready, compression)
--> 946         return _end_unary_response_blocking(state, call, False, None)
    947

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/grpc/_channel.py in _end_unary_
response_blocking(state, call, with_call, deadline)
    848         else:
--> 849             raise _InactiveRpcError(state)
    850

```

```

_InactiveRpcError: <_InactiveRpcError of RPC that terminated with:
  status = StatusCode.UNAVAILABLE
  details = "Socket closed"
  debug_error_string = "{"created":"@1641910066.927447291","description":"Error rece
ived from peer ipv4:10.212.98.108:46755","file":"src/core/lib/surface/call.cc","file_lin
e":1069,"grpc_message":"Socket closed","grpc_status":14}"
>

```

During handling of the above exception, another exception occurred:

```

_InactiveRpcError                                Traceback (most recent call last)
/tmp/ipykernel_1093282/4203283615.py in <module>
     5     # iterate over the timesteps
     6
----> 7     loihi_sim.run(.25 * 10)
     8
     9

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/simulator.py in __e
xit__(self, exc_type, exc_value, traceback)
    227     def __exit__(self, exc_type, exc_value, traceback):
    228         for sim in self.sims.values():
--> 229             sim.__exit__(exc_type, exc_value, traceback)
    230         self.close()
    231

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardware/interface.
py in __exit__(self, exc_type, exc_value, traceback)
    123
    124     def __exit__(self, exc_type, exc_value, traceback):

```

```

--> 125         self.close()
      126
      127     @classmethod

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardware/interface.py in close(self)
      160
      161         if self.nxsdk_board is not None:
--> 162             self.nxsdk_board.disconnect()
      163             self.nxsdk_board = None
      164

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/arch/base/nxboard.py in disconnect(self)
      343         """
      344         BasicSpikeGenerator.isSpikeGenProcessConfigured = False
--> 345         self.executor.stop()
      346         self._executor = None
      347

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/executor.py in stop(self, force)
      95             _force.force = force
      96             self._executor_service.stopExecution(_force)
--> 97             self._notifyListeners(ExecutionEventEnum.ON_STOP)
      98             self._host_coordinator.stop()
      99             self._state = ExecutionState.UNDEFINED

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/executor.py in _notifyListeners(self, event)
      157             listener.postExecution()
      158             elif event == ExecutionEventEnum.ON_STOP:
--> 159                 listener.onStop()
      160             else:
      161                 raise Exception("Invalid event {}".format(event))

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/listeners/lakemont_orchestrator.py in onStop(self)
      39     def onStop(self) -> None:
      40         """Stops the lakemont driver"""
--> 41         self.stopNxDriver(empty)

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/grpc/_channel.py in __call__(self, request, timeout, metadata, credentials, wait_for_ready, compression)
      944         state, call, = self._blocking(request, timeout, metadata, credentials,
      945                                     wait_for_ready, compression)
--> 946         return _end_unary_response_blocking(state, call, False, None)
      947
      948     def with_call(self,

~/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/grpc/_channel.py in _end_unary_response_blocking(state, call, with_call, deadline)
      847         return state.response
      848     else:
--> 849         raise _InactiveRpcError(state)
      850
      851

```

```

_InactiveRpcError: <_InactiveRpcError of RPC that terminated with:
  status = StatusCode.UNAVAILABLE
  details = "failed to connect to all addresses"
  debug_error_string = "{"created":"@1641910066.932258110","description":"Failed to pick subchannel","file":"src/core/ext/filters/client_channel/client_channel.cc","file_line":3158,"referenced_errors":[{"created":"@1641910066.932256728","description":"failed to connect to all addresses","file":"src/core/lib/transport/error_utils.cc","file_line":147,"g

```

```
rpc_status":14}}"  
>
```

In []: