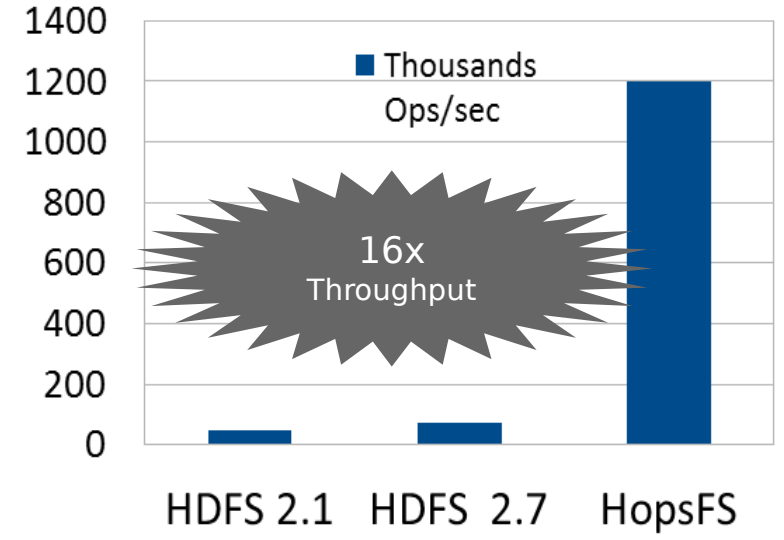



LOGICAL CLOCKS

Scaling Deep Learning to 100s of
GPUs on Hops Hadoop


Fabio Buso
Software Engineer
Logical Clocks AB

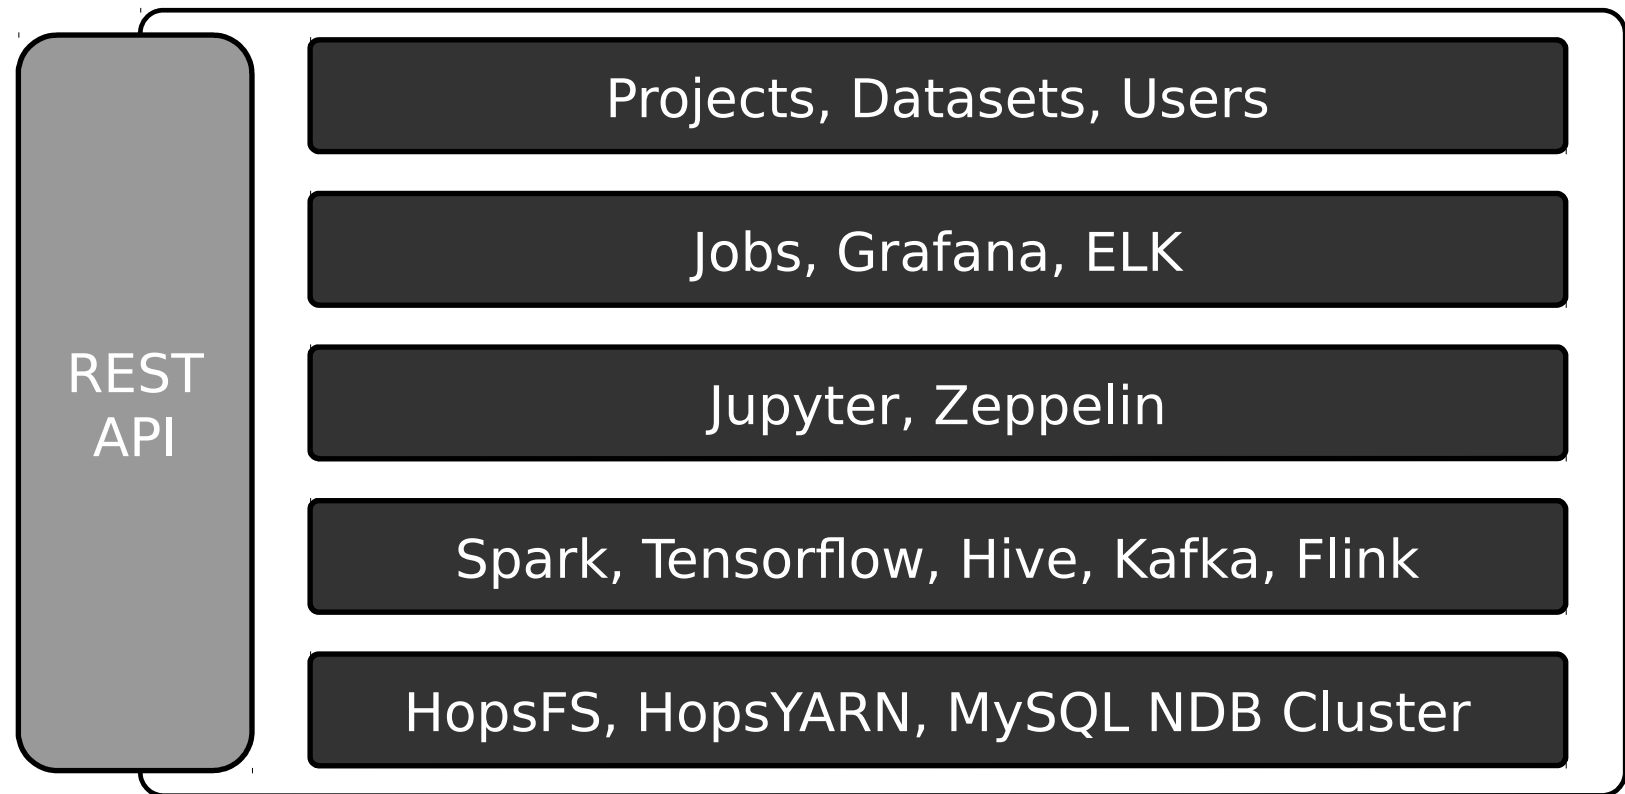
HopsFS: Next generation HDFS



*<https://www.usenix.org/conference/fast17/technical-sessions/presentation/niazi>
**<https://eurosys2017.github.io/assets/data/posters/poster09-Niazi.pdf>

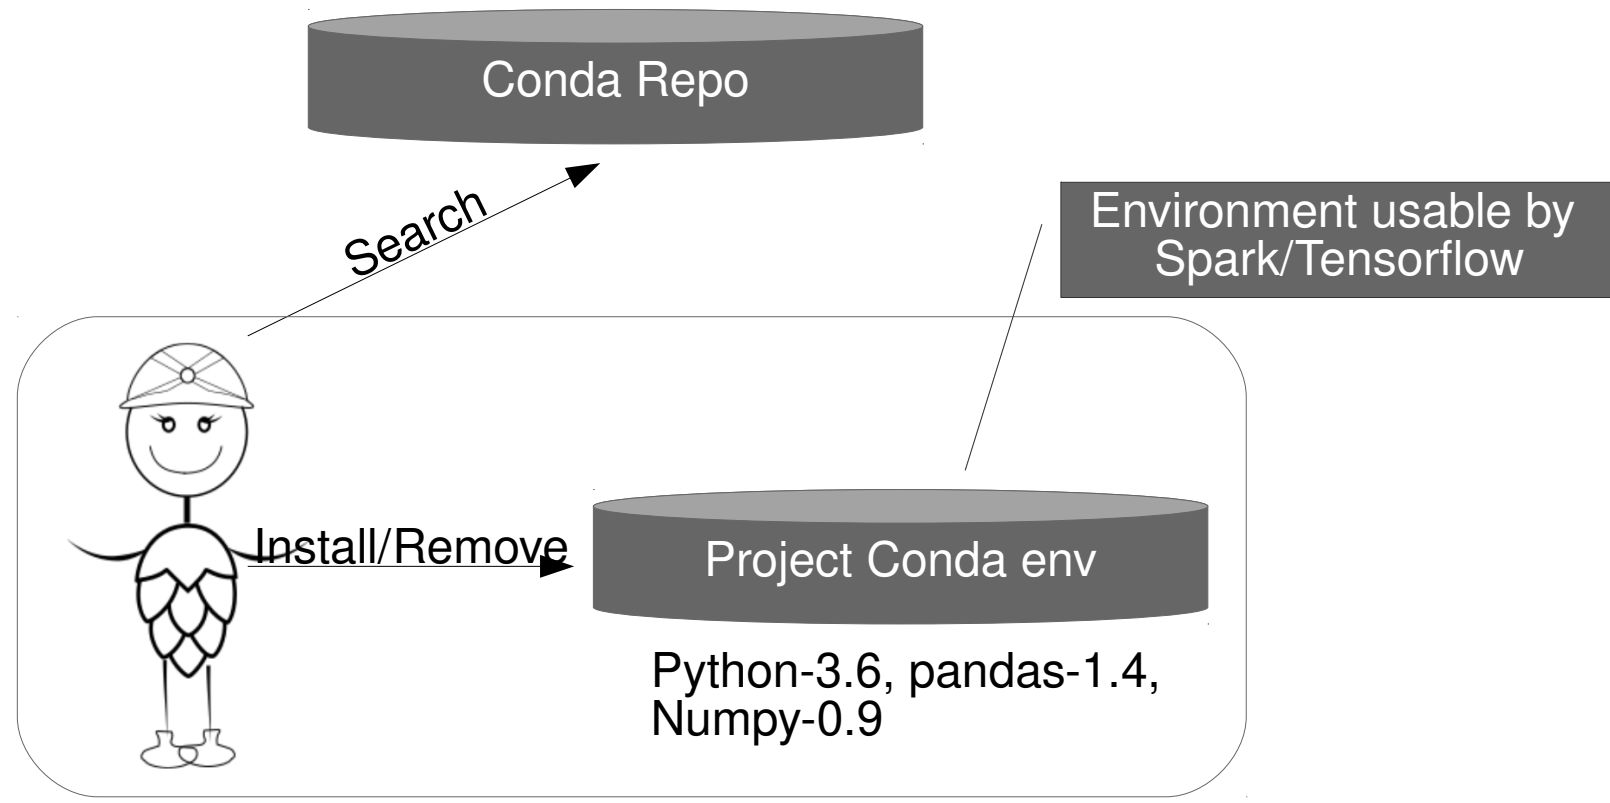
 **IEEE** Scale Challenge Winner (2017)





Version 0.3.0 just released!

Python first

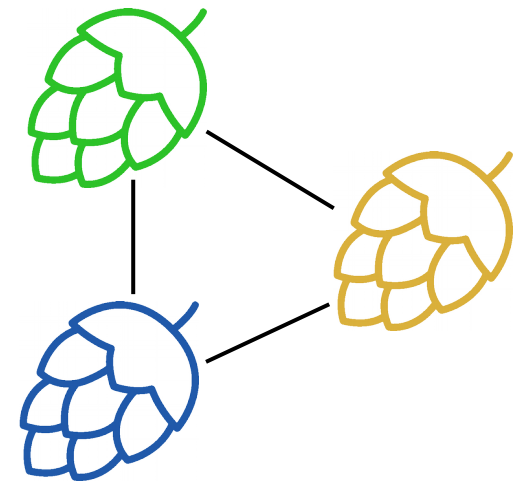


Hops python library: Make development easy

- Hyperparameter searching
- Manage Tensorboard lifecycle



- Discover, Share and experiment with interesting datasets
- p2p network of Hops Cluster
- ImageNet, YouTube8M, Reddit comments...
- Exploits unused bandwidth



[*http://ieeexplore.ieee.org/document/7980225/](http://ieeexplore.ieee.org/document/7980225/) (ICDCS 2017)

Scale out level: 1
Parallel Hyper parameter searching

Parallel Hyperparameter searching

```
def model(lr, dropout):  
    ...  
  
args_dict = {  
    'learning_rate': [0.001, 0.0005, 0.0001],  
    'dropout': [0.45, 0.7]}  
  
args_dict_grid = util.grid_params(args_dict)  
  
tflauncher.launch(spark, model, args_dict_grid)
```



Starts 6 parallel experiments

Scale out Level: 2 Distributed Training



- Distributed TensorFlow over Spark
- Runs on top of a Hadoop cluster
- PS/Workers executed inside Spark executors
- Uses Spark for resource allocations
 - Our version: exclusive GPUs allocations
 - Parameter server(s) do not get GPU(s)
- Manages Tensorboard

Run TFoS


```
def training_fun(argv, ctx):  
    .....  
    TFNode.start_cluster_server()  
    .....  
TFCluster.run(spark, training_fun, num_exec, num_ps...)
```

Full conversion guide:

<https://github.com/yahoo/TensorFlowOnSpark/wiki/Conversion-Guide>

Scale out level: Master of the dark arts
Horovod

PS server architecture doesn't scale




From: <https://github.com/uber/horovod>

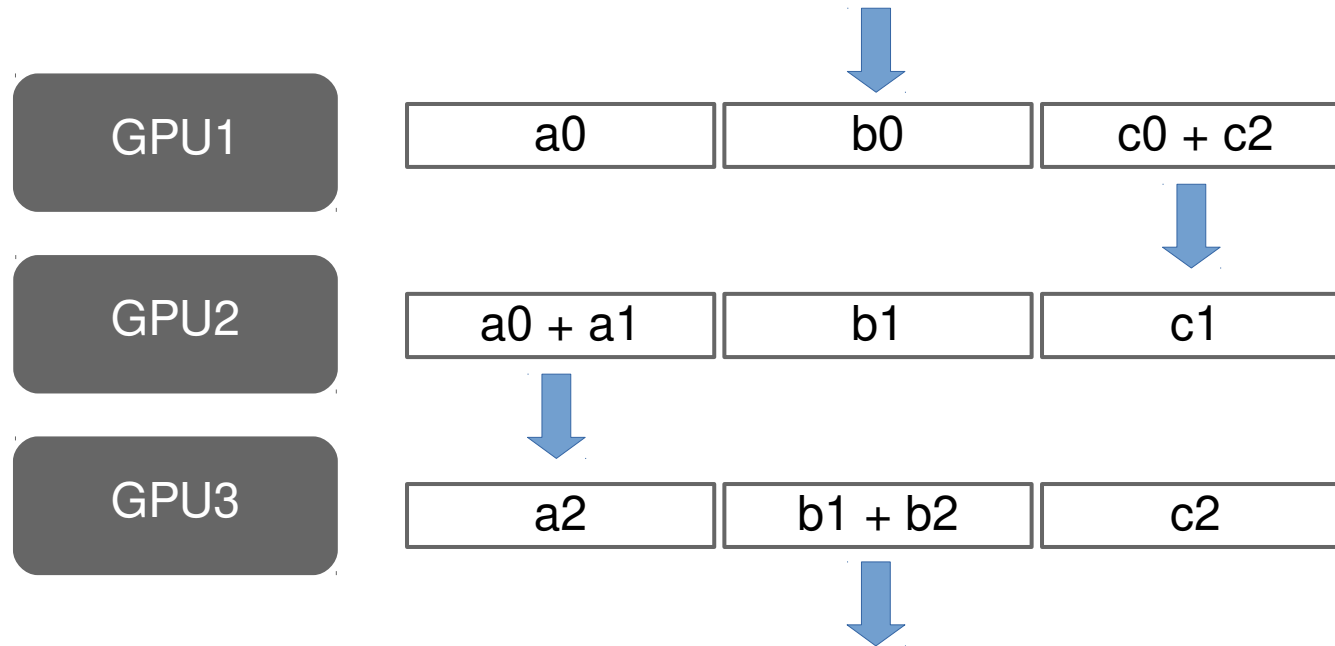


- Based on previous work done by Baidu
- Organize workers in a ring
- Gradients updates distributed using All-Reduce
- Synchronous protocol

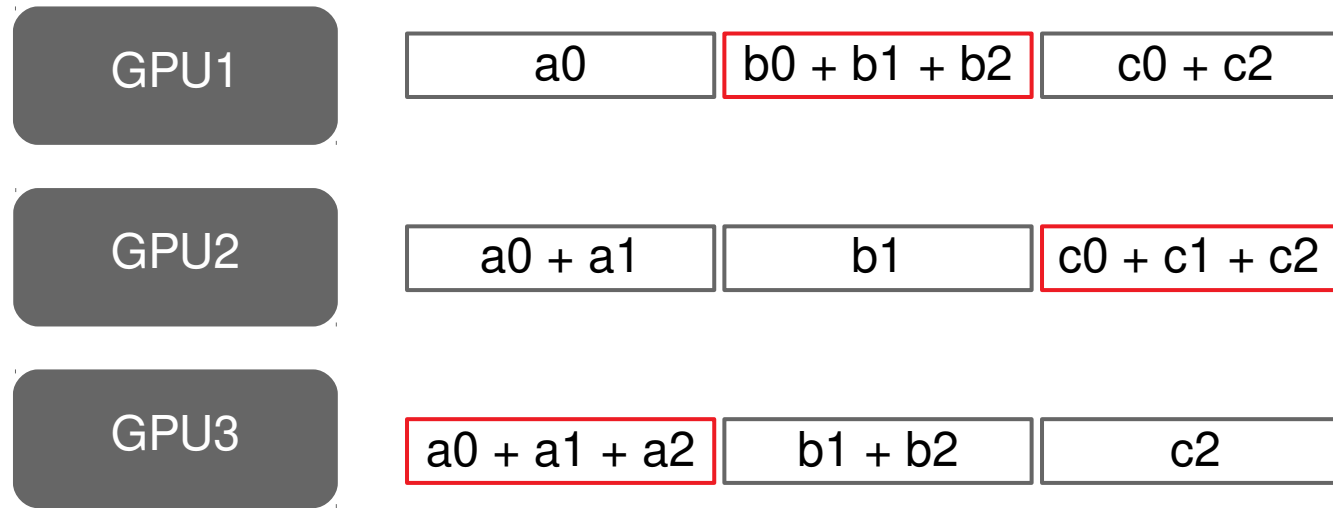
All-Reduce



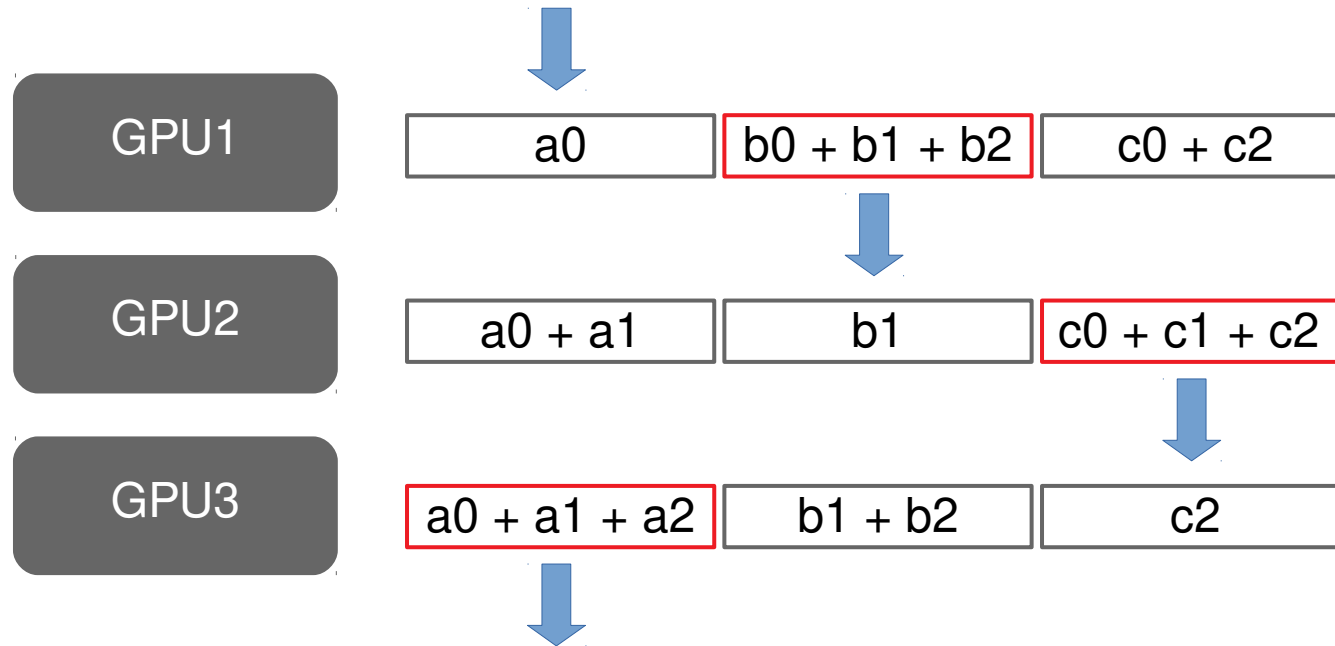
All-Reduce



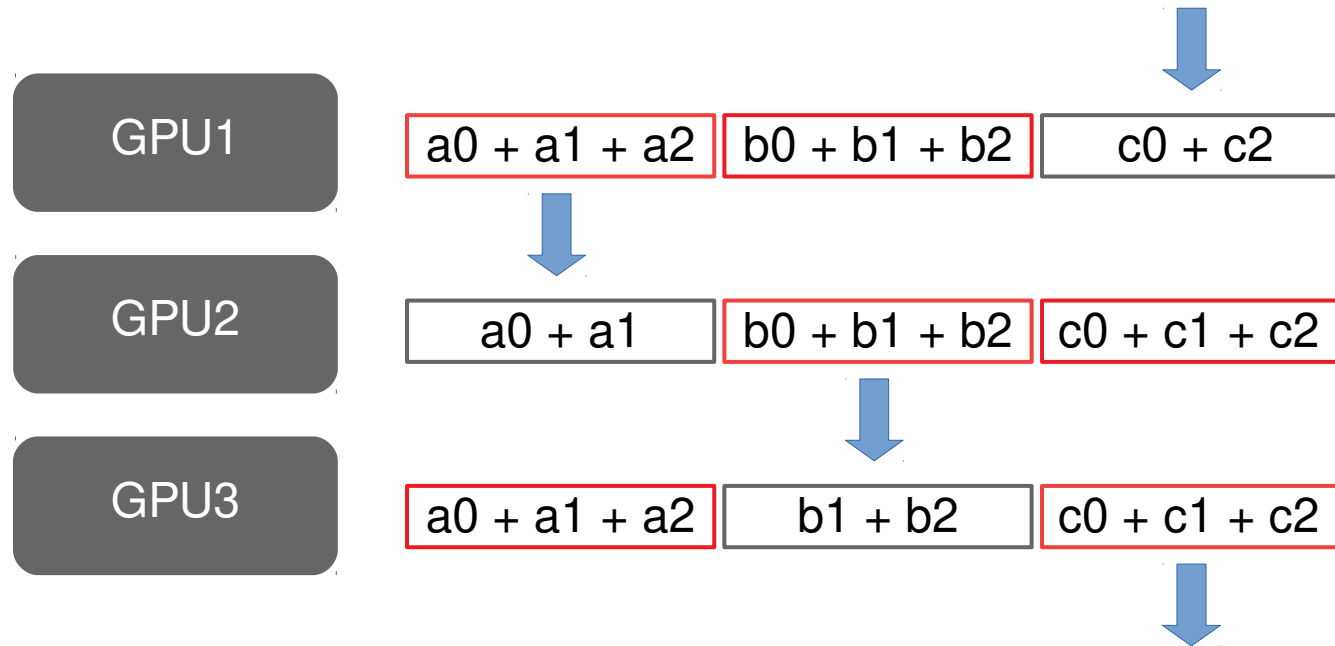
All-Reduce



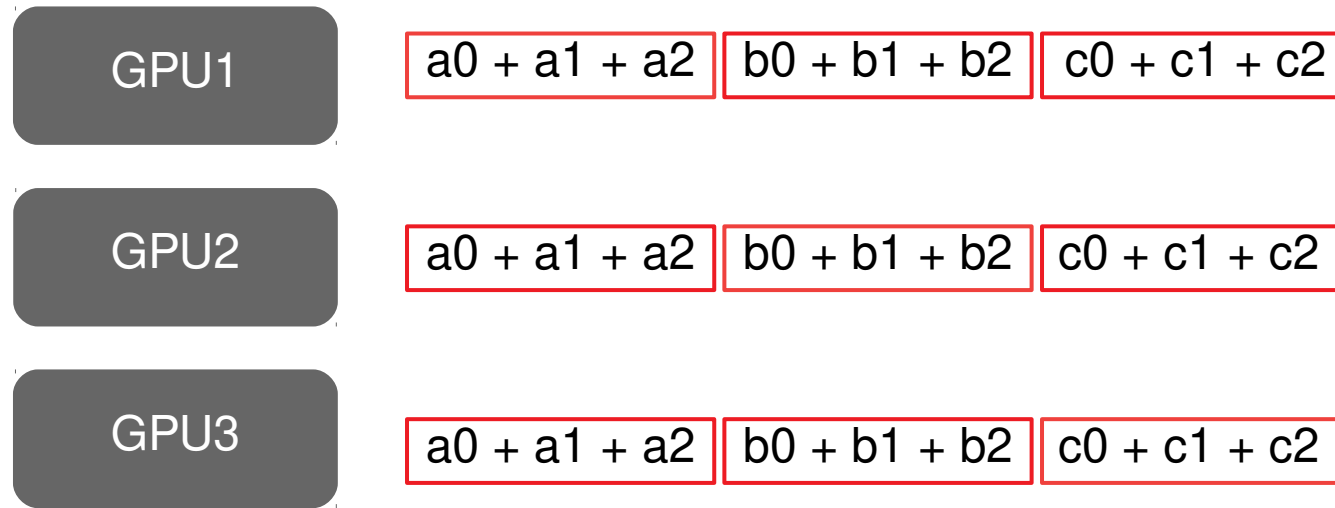
All-Reduce



All-Reduce




All-Reduce



Hops AllReduce

```
import horovod.tensorflow as hvd
def conv_model(feature, target, mode)
.....
def main(_):
    hvd.init()
    opt = hvd.DistributedOptimizer(opt)
    if hvd.local_rank() == 0:
        hooks = [hvd.BroadcastGlobalVariablesHook(0), ..]
        .....
    else:
        hooks = [hvd.BroadcastGlobalVariablesHook(0), ..]
        .....
    from hops import allreduce
    allreduce.launch(spark, 'hdfs:///Projects/
.../all_reduce.ipynb')
```

Demo time!



- Play with it → hops.io/?q=content/hopsworks-vagrant
- Doc → hops.io
- Star us! → github.com/hopshadoop
- Follow us! → [@hopshadoop](#)