




ÓBUDAI EGYETEM
ÓBUDA UNIVERSITY

Scaling up distributed deep learning on the SLICES infrastructure

Póra Krisztián


Experiment motivation

- Deep Learning
 - Neural networks with a large number of layers
 - Significant increase in trainable parameters
 - Better suited for solving complex tasks
- Highly resource intensive
 - Large volume of training data
 - Prolonged training time
 - A limiting factor for a long time
- Toolset and methodology
 - GPU cards
 - Distributed training
 - Frameworks



CPU
Multiple Cores

+



GPU
Thousands of Cores

<https://www.cherryservers.com/blog/gpu-vs-cpu-what-are-the-key-differences>





slices **RI**

umec

Transnational access

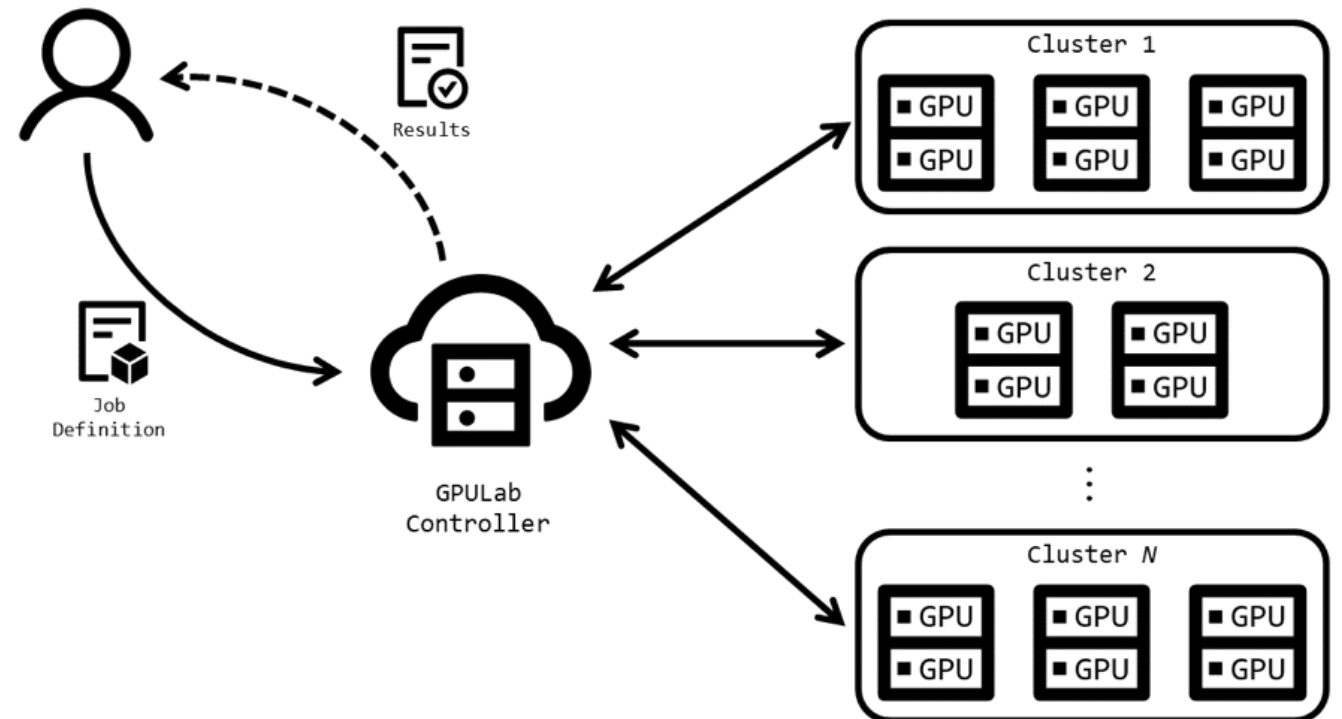
- Trip to Ghent, Belgium
 - Travel and accommodation
- Work at the IDLab office
- Presentations and tour
- Continuous support



- 550+ bare metal and GPU servers
- Fully configurable (nodes and network configuration)
- Fast context switching
- Remote access
- Networking, cloud, big data and application testing and research



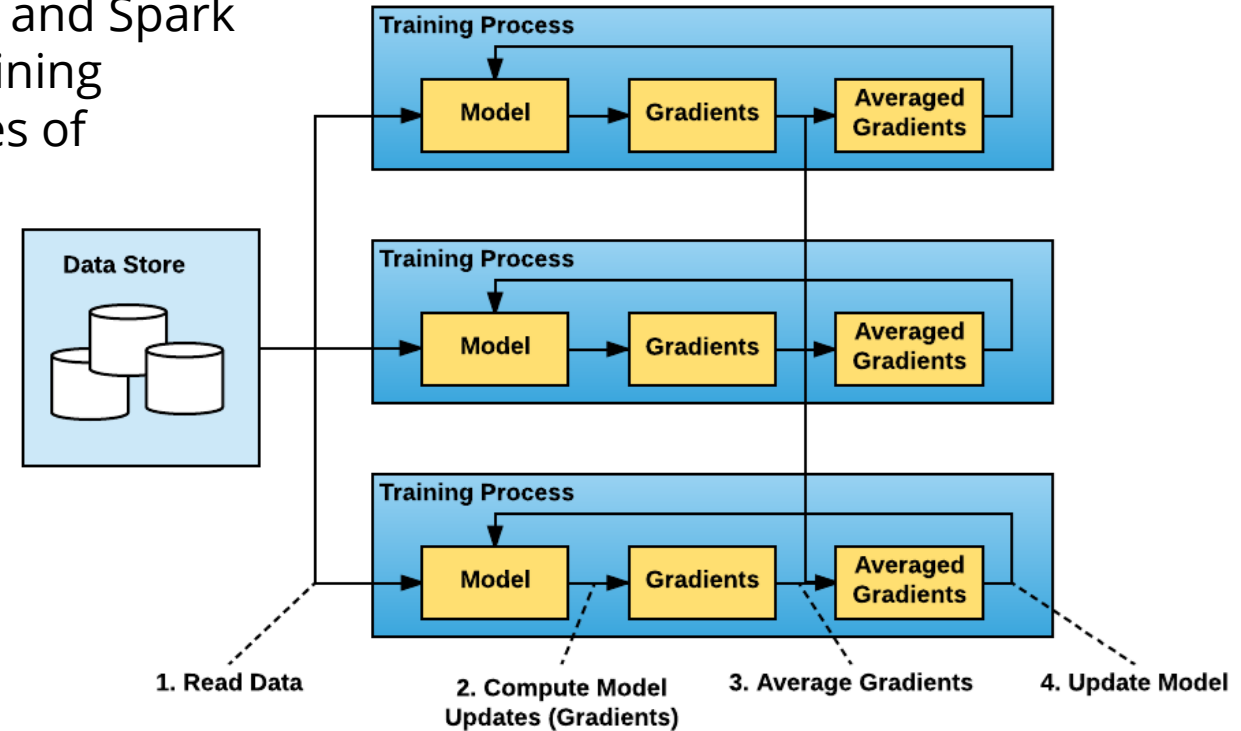
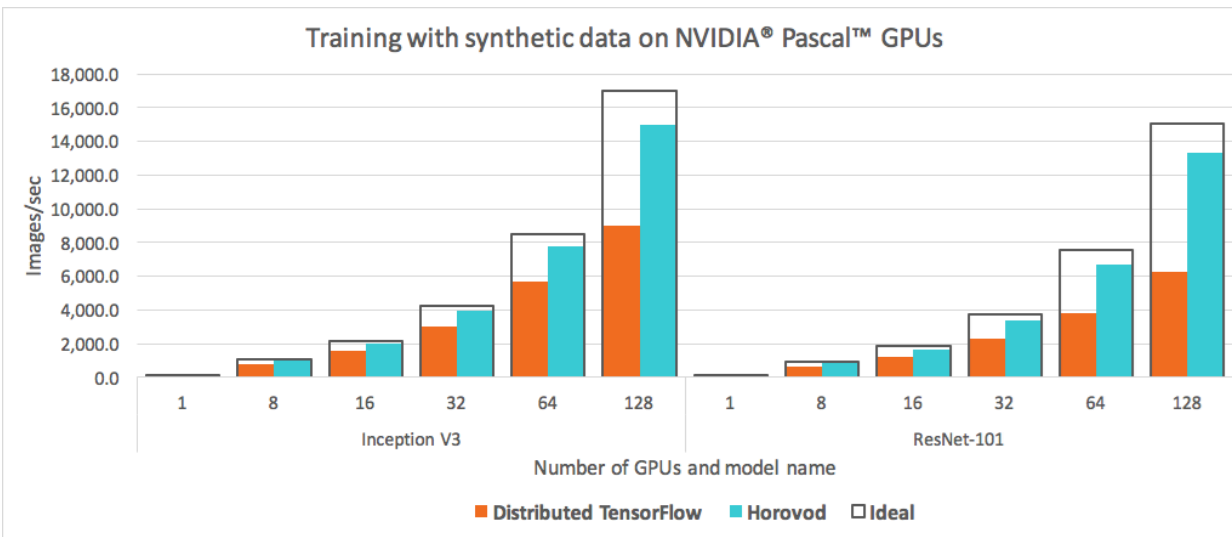
- Heterogenous clusters
- 150+ GPUs
 - 700.000 CUDA cores
 - 3TB GPU RAM
- Job-based GPU processing
- Containerized execution
- JupyterHub access
- AI and data processing research



GPULab

Introduction to Horovod

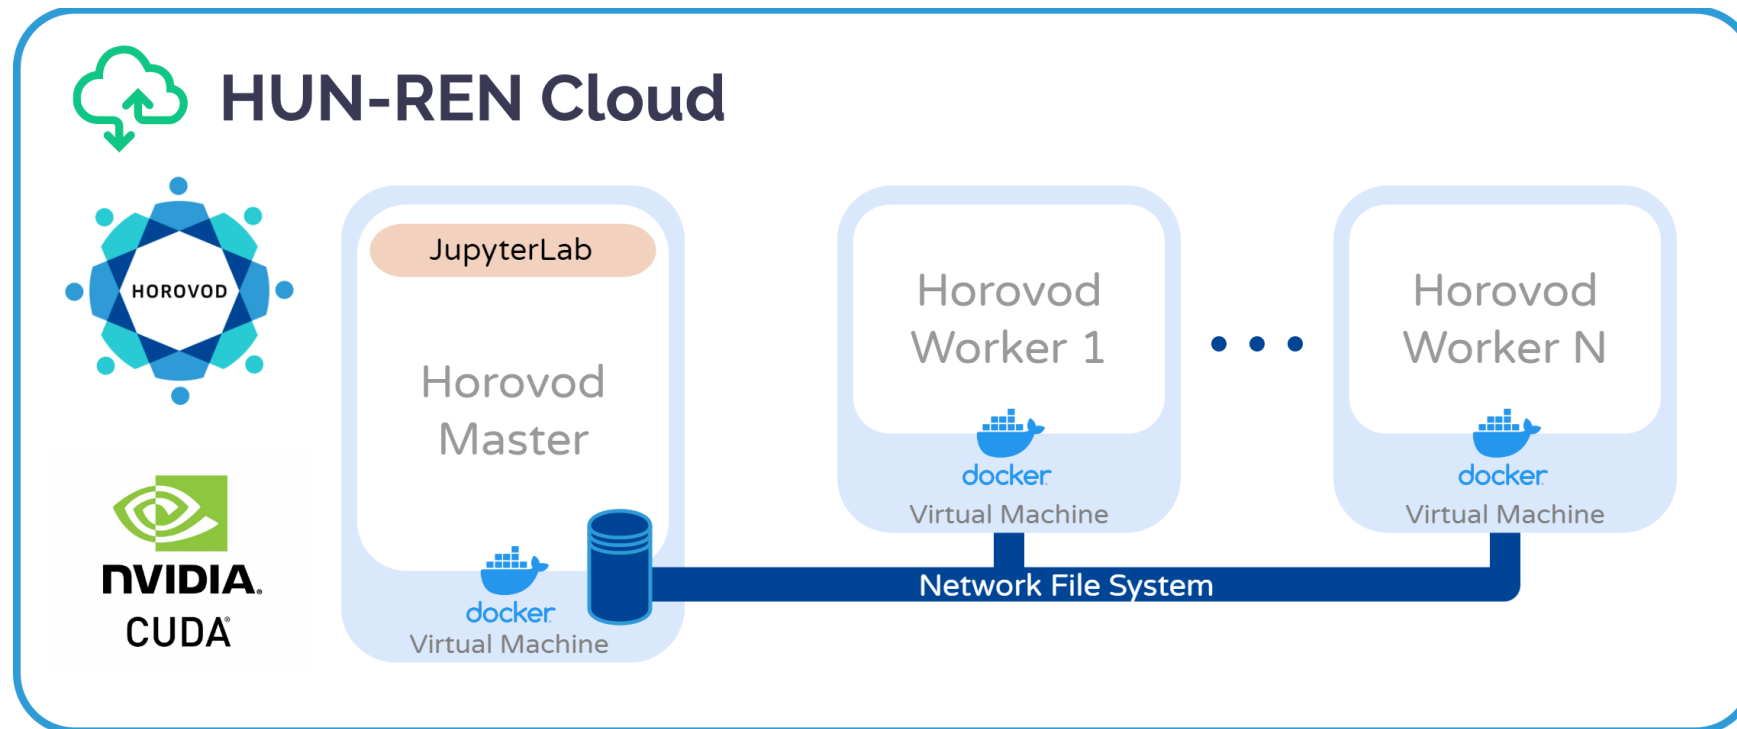
- Open-source distributed deep learning framework by Uber
- Supports TensorFlow, Keras, PyTorch, Apache MXNet and Spark
- Provide an easy-to-use framework for distributed training
 - Execute on hundreds of GPUs with just a few lines of additional code
- Data parallel execution



Sergeev, A., & Balso, M.D. (2018). Horovod: fast and easy distributed deep learning in TensorFlow. ArXiv, abs/1802.05799.

Horovod Reference Architecture

- Support for distributed deep learning applications
- JupyterLab development environment
- Network based file sharing between nodes
- Utilization of GPU resources





HashiCorp

Terraform

- **Provision infrastructure**
 - Virtual machines
 - Network settings
 - Firewall rules
 - Execute tasks
 - Invoke Ansible



ANSIBLE

- **Configure nodes**
 - Install packages
 - Start services
 - Run Docker containers

Deployment – Virtual Wall



HashiCorp

Terraform

- **Provision infrastructure**
 - Virtual machines
 - Network settings
 - Firewall rules
 - Execute tasks
 - Invoke Ansible

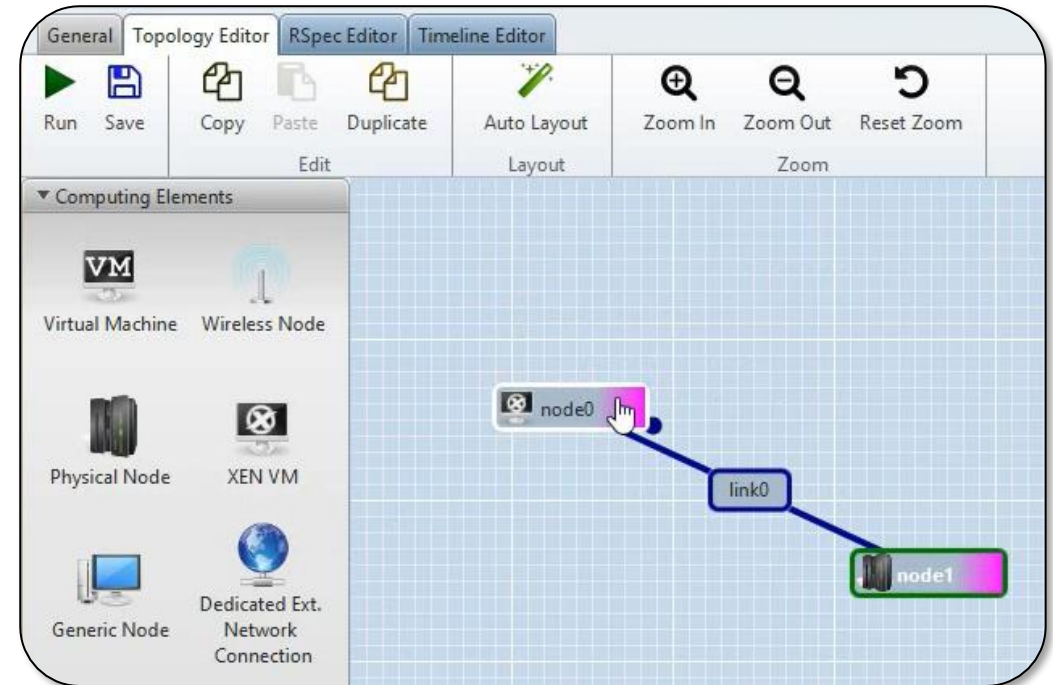


jFed



ANSIBLE

- **Configure nodes**
 - Install packages
 - Start services
 - Run Docker containers





HashiCorp

Terraform

- **Provision infrastructure**
 - Virtual machines
 - Network settings
 - Firewall rules
 - Execute tasks
 - Invoke Ansible



ANSIBLE

- **Configure nodes**
 - Install packages
 - Start services
 - Run Docker containers

Create a job

Job Definition

Core info

* Name:

Description:
Describe what your experiment does (optionally)

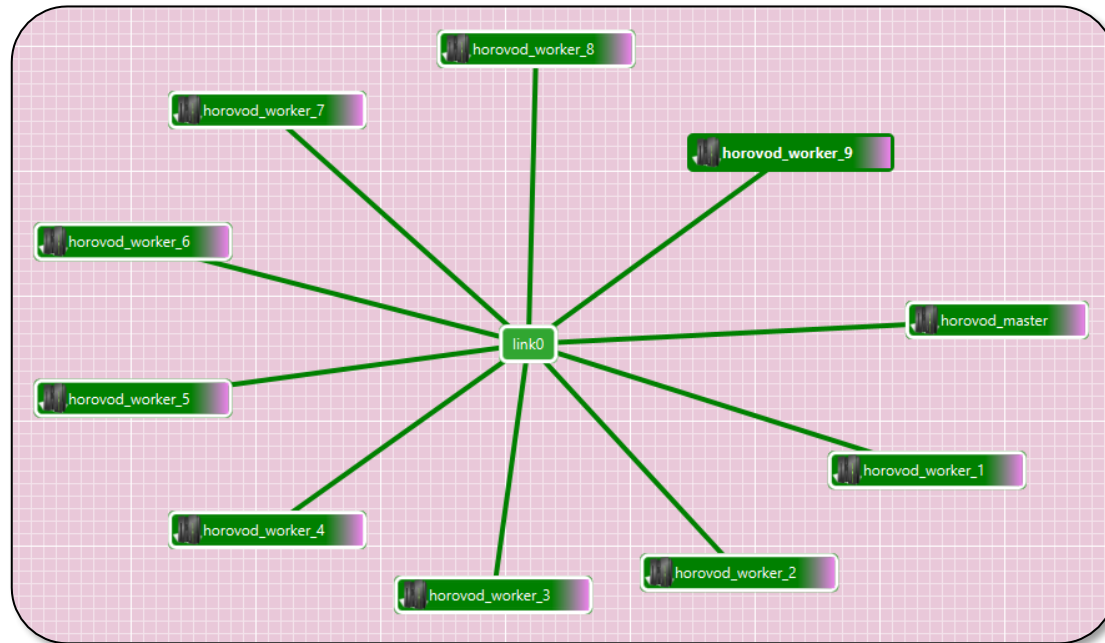
* Cluster ID:
Uncheck this to run this job on any compatible cluster

* Docker Image:

Resulting jobDefinition:

```
{
  "jobDefinition": {
    "name": "NVIDIA SMI"
    "description": "Writes the output of the command 'nvidia-smi' to the log and exits"
    "dockerImage": "gpulab.ilabt.imec.be:5000/sample:nvidia-smi"
    "command": ""
    "resources": {
      "gpus": 1
      "systemMemory": 2048
      "cpuCores": 2
      "minCudaVersion": 10
    }
  }
}
```

Virtual Wall



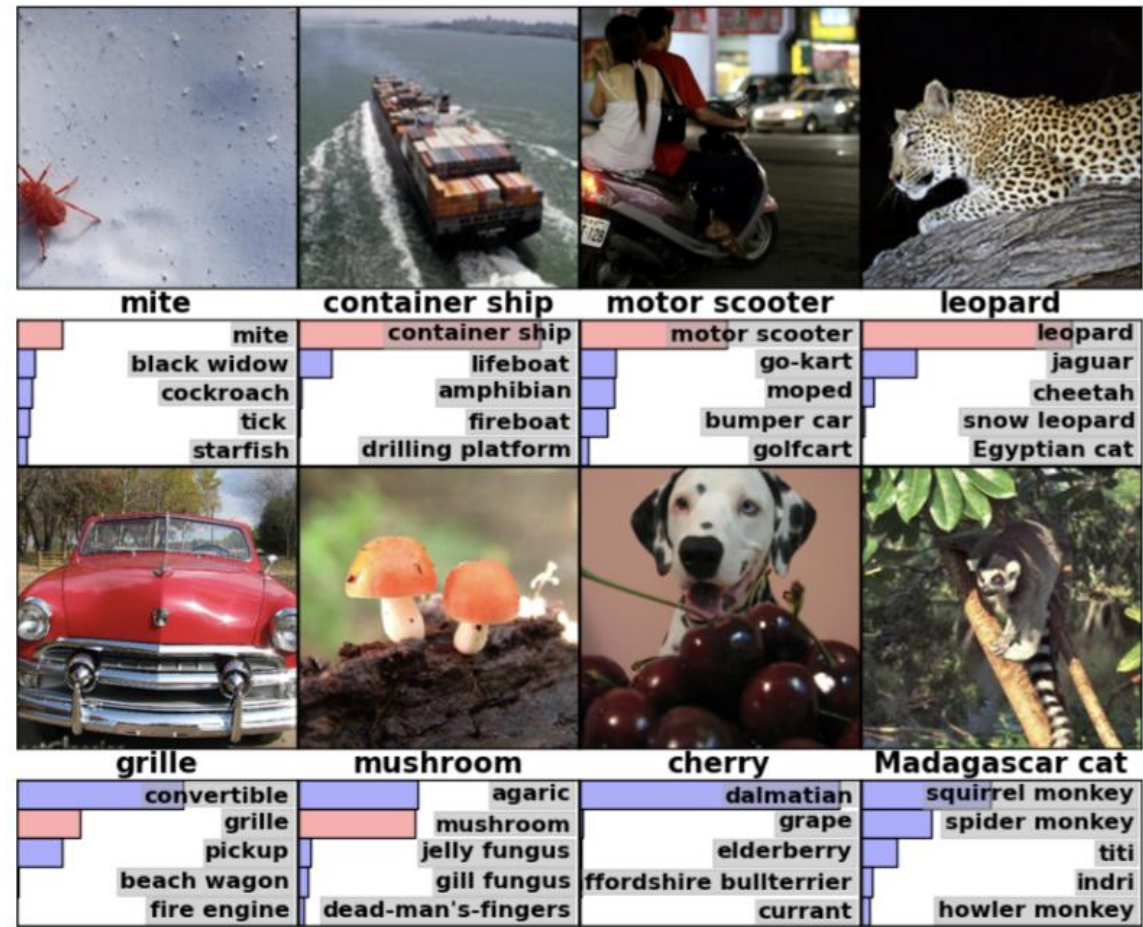
GPULab

A screenshot of the GPULab interface showing a grid of running jobs. The jobs are organized into three rows and four columns. Each job card displays the job name, ID, resource allocation (1GPU, 2CPUs, 24GB RAM), and a 'Running' status with a 'Cancel' button. The jobs are:

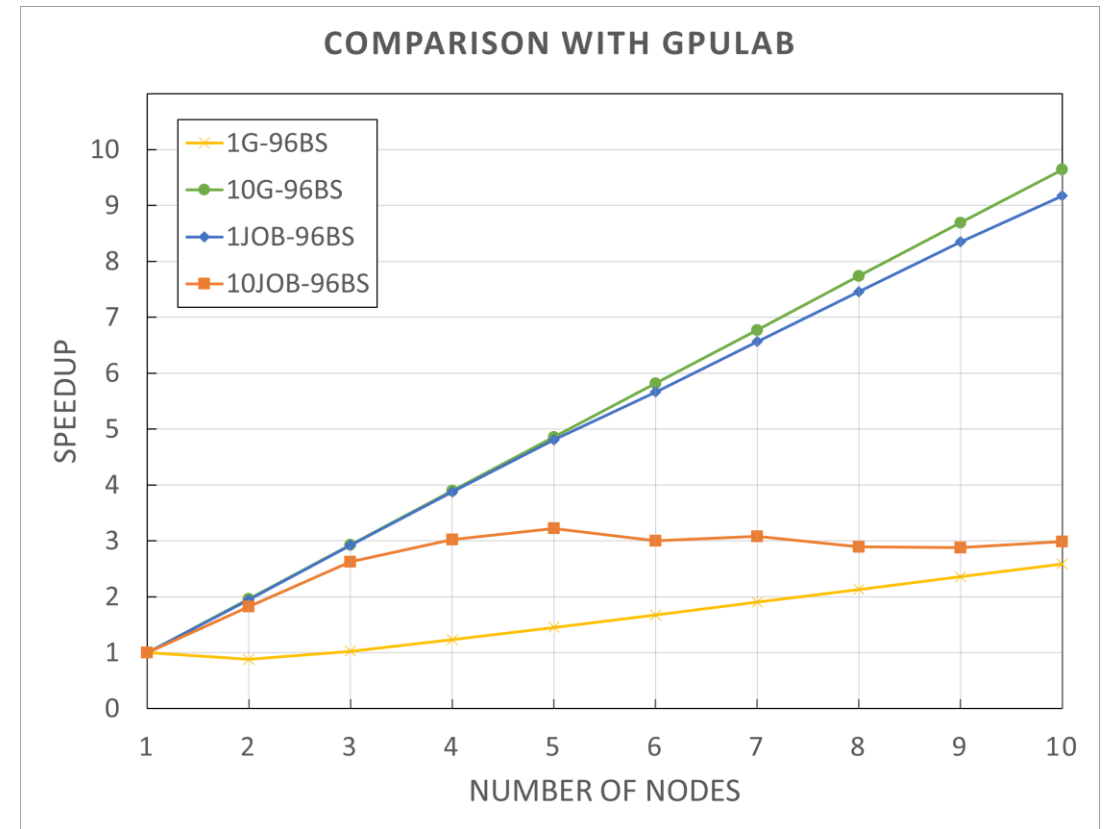
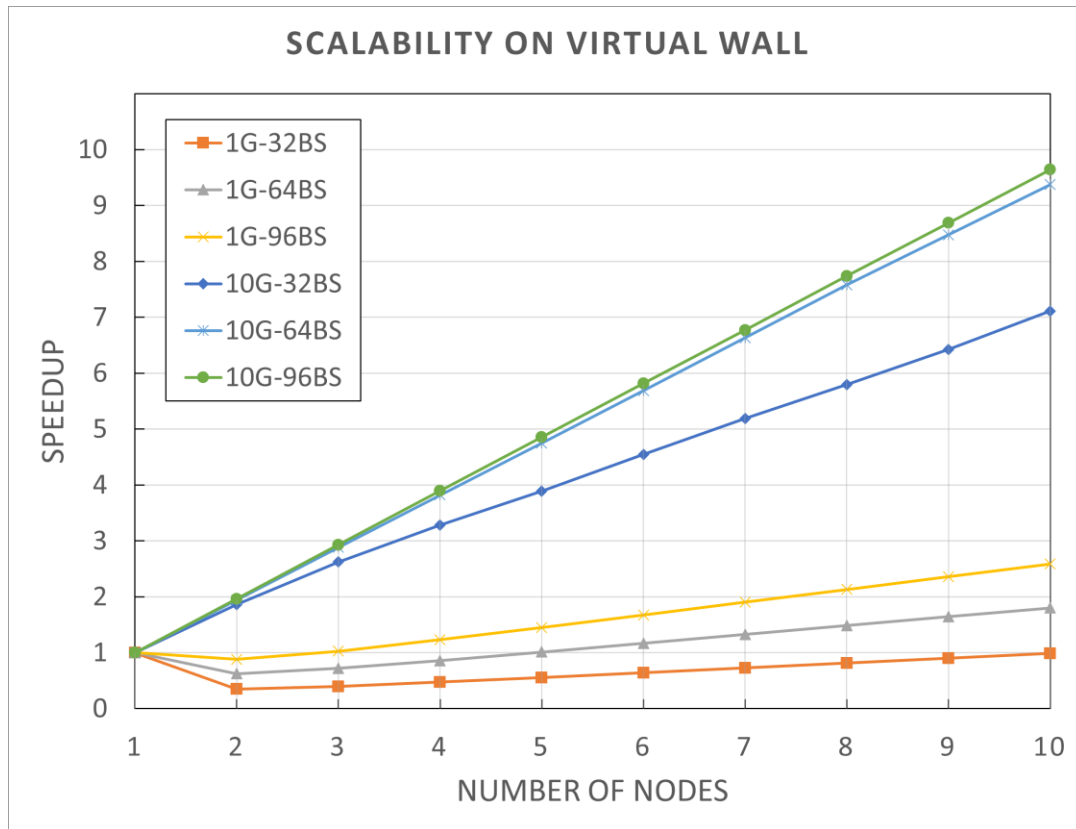
- Row 1: horovod-worker9, horovod-worker8, horovod-worker7, horovod-worker6
- Row 2: horovod-worker5, horovod-worker4, horovod-worker3, horovod-worker2
- Row 3: horovod-worker1, horovod-master

Benchmark parameters

- Training dataset
 - ImageNet
- Batch size
 - 32, 64, 96
- Number of nodes (1-10)
- Measured parameters
 - Processing performance
 - Scaling efficiency
- Environment
 - Virtual Wall – 1 Gbit/s
 - Virtual Wall – 10 Gbit/s
 - GPULab – 1 job with 10 GPUs
 - GPULab – 10 jobs with 1 GPU each



Benchmark results





ÓBUDAI EGYETEM
ÓBUDA UNIVERSITY

Thank you for your attention!