

### Making environmental science environmentally friendly Follow the merm(AI)d

Dr. Francesco Bongiovanni





## HARDWARE SOFTWARE



## ARCHITECTURE





## **ESG IMPACT**





# ARCHITECTURE



LuxProvide offers a unique platform that combines data science and supercomputing resources to help organizations increase the ROI of their most challenging innovation projects





•

ICT Business Partner of the Year



LA ONMOLNONLAONLAONN NOMOLNONMLAONNLAON NOMLAONMLA©NMLAONNLA

<u>onnomnomnonnomnonno</u>

A CALMON A AMAGA AMAGA AMAGA AMAGA AMA AMAGA AMA

A ONE NOMERON NO VANONAL ON COMPLONAL COMPANY ON THE COMPLEX ON TH

MONTONONNONNONNONNONNON



## The EUROHPC Supercomputers

- 5 operational systems in Luxembourg, Slovenia, Czechia, Bulgaria & Finland;
- 3 systems underway in Italy, Portugal & Spain;
- New hosting entities selected in June 2022, to host new mid-range systems and the first European exascale supercomputer.





#### Luxembourg's national supercomputer Meluxina





90.000 5.120 **HPC CPU Cores** 

800 **GPU-AI** accelerators

20 Petabytes high performance storage Cloud vCPUs

Infiniband HDR 200 Gb/s

300+**Tailored software** packages



## **World-class supercomputing**







#### Performance

Compute 18 FP64 PetaFlops 500 Al PetaFlops

Data **500+** <sub>GB/s</sub> Scratch storage **190+** <sub>GB/s</sub> Project storage

Interconnect 200 Gb/s on Cluster nodes 400 Gb/s on Accelerator and large memory nodes





June 2021, Accelerator Module

The





Accelerated HPC, HPDA & AI workloads

Mixed HPC, HPDA & BigData workloads

#### Purpose:

TT

## MeluXina User Software Environment

Compilers, Languages & Performance Eng.

Numerical & data libraries

Parallelization tools, MPI suites & acceleration libraries

Frameworks, runtime & platform tools

End user applications



Container system on MeluXina

- Docker & OCI compatible •

- Support for creating and running encrypted • containers
- verified

#### Bring-your-own software stack

- REPRODUCIBILITY
- Users create tooling and pipeline on their infrastructure --- RUN PIPELINE ON MELUXINA



#### **Singularity-CE**



Non-privileged mode for improved security Support for GPU accelerated applications

• Support for trusted containers: PGP signed &



Enable users to easily control complete stack ----

#### PG













**TensorFlow** 





**O** PyTorch











## **Green TierIV Data Center**

- PUE of the Data Centers constantly measured and monitored
- Continuous improvement plan for Energy Efficiency supervised by the government agency Klima-Agence
- Waste heat from servers used to heat office space and preheat diesel generators





- Optimised use of Free Cooling
- **Biomass** recycling representing a yearly reduction in CO<sub>2</sub> in excess of 27,000 metric tons





100% Green Electricity supplied from hydroelectric power sources















### Know your hardware



7H12								
Cores	64							
Base clock	2.6 Ghz							
Max. Boost clock	3.3 Ghz							
_3 cache	256 MB							
Memory	DDR4							
Memory channel	8							
Memory Bandwidth	204.8 GB/s per socket							
TDP	280W							

A100 SXM									
FP64	9.7 TFLOPS								
FP64	19.5 TFLOPS								
Tensor Core									
FP32	19.5 TFLOPS								
TF32	156 TFLOPS								
BFLOAT16	312 TFLOPS								
FP16	312 TFLOPS								
INT8	624 TOPS								
Memory	40GB HBM2								
Memory Bandwidth	1,555 GB/s								
Interconnect	600GB/s								
TDP	250W								







#### Stratix 10 MX 2100

#### Know your hardware

PROGRAMMING EFFORTS













#### **Al-oriented profilers: Scalene**

#### scalene python3 train\_1gpu.py

									Memo	ory usage train_	e: 1gpu
Time spent in	r	Line	Time Python	native	syst	GPU	Memory Python	peak	timeline/%	Copy (MB/s)	tra
		1 2 3 4									Most
Time spent in		5									Niko
native code		7	7% ንԳ	12%	48%		63%	83M	5%	41	impo
		9 10 11	20		140		100%			•	impo impo impo
Time spent on		12 13									impo impo
GPU		14 15									
		16 17 18 19									def
Data movement		20 21								╊	
		22 23 24									
		25 26	2%	1%	3%		76%	295M	34%	5	
		27 28									
		29 30									
		31 32	2%		1%		79%	64M	<sup>7%</sup>	3	
		33 34									
		35 36									
		37									



py: % of time = 100.00% (2m:54.443s) out of 2m:54.443s.

n\_1gpu.py

```
ly based on the official pytorch tutorial https://pytorch.org/tutorials/beginner/blitz/cifar10_tutorial.html
 ied for educational purposes.
 las, AI Summer 20222
rt torch
rt torchvision
 t torchvision.transforms as transforms
 t torch.nn as nn
rt torch.nn.functional as F
 t torch.optim as optim
rt time
create_data_loader_cifar10():
transform = transforms.Compose(
    transforms.RandomCrop(32),
    transforms.RandomHorizontalFlip(),
   transforms.ToTensor(),
   transforms.Normalize((0.5, 0.5, 0.5), (0.5, 0.5, 0.5))])
batch_size = 64
trainset = torchvision.datasets.CIFAR10(root='./data', train=True,
                                          download=True, transform=transform)
trainloader = torch.utils.data.DataLoader(trainset, batch_size=batch_size,
                                          shuffle=True, num_workers=10, pin_memory=True)
testset = torchvision.datasets.CIFAR10(root='./data', train=False,
                                     download=True, transform=transform)
testloader = torch.utils.data.DataLoader(testset, batch_size=batch_size,
                                          shuffle=False, num_workers=10)
return trainloader, testloader
```

#### **Al-oriented profilers: Scalene**

#### Zooming on GPU usage



Time spent on GPU



```
def train(net, trainloader):
   print("Start training...")
   criterion = nn.CrossEntropyLoss()
   optimizer = optim.SGD(net.parameters(), lr=0.001, momentum=0.9)
   epochs = 1
   num of batches = len(trainloader)
   for epoch in range(epochs): # loop over the dataset multiple times
       running loss = 0.0
       for i, data in enumerate(trainloader, 0):
           inputs, labels = data
           images, labels = inputs.cuda(), labels.cuda()
           optimizer.zero grad()
           outputs = net(images)
           loss = criterion(outputs, labels)
           loss.backward()
           optimizer.step()
           running loss += loss.item()
       print(f'[Epoch {epoch + 1}/{epochs}] loss: {running loss / num of batches:.3f}')
```

```
print('Finished Training')
```

#### **Al-oriented profilers: Scalene**

#### scalene --json --outfile profile\_scalene.json python3 train\_1gpu.py

Using the scalene GUI (<u>http://plasma-</u> mass.org/scalene-gui/), You can view the .json profile and ask ChatGPT for some optimizations 🙂

show all hide all only display profiled lines 🗹

▼train\_1gpu.py: % of time = 78.6% (1m:19.256s) out of 1m:40.815s.

TIME	MEMORY	MEMORY	MEMORY	MEMORY	COPY	<u>GPU</u>	<u>GPU</u>	LIN	IE PROFILE (click to reset order)
	average	peak	timeline	activity		util.	memory	tr	ain_1gpu.py
	1	1	_		42			7	4-import torch
					17			8	4 import torchvision
								16 💢	<pre>4 def create_data_loader_cifar10():</pre>
					6			26	<pre>trainset = torchvision.datasets.CIFAR10(root='./data', train=True,</pre>
1	1	1		'	4			31	<pre>testset = torchvision.datasets.CIFAR10(root='./data', train=False,</pre>
								38 💢	<pre>def train(net, trainloader):</pre>
								44 💢	<pre>for epoch in range(epochs): # loop over the dataset multiple times</pre>
							•	47 💢	<pre>for i, data in enumerate(trainloader, 0):</pre>
								51	<pre>images, labels = inputs.cuda(), labels.cuda()</pre>
								54	<pre>optimizer.zero_grad()</pre>
					30		)	57	<pre>outputs = net(images)</pre>
								58	<pre>loss = criterion(outputs, labels)</pre>
						•		59	<pre>loss.backward()</pre>
								60	<pre>4 optimizer.step()</pre>
						1		63	<pre>4 running_loss += loss.item()</pre>
								70 💢	<pre> def test(net, PATH, testloader): </pre>
					1			71	<pre>net.load_state_dict(torch.load(PATH))</pre>
						1		77 💢	for data in testloader:
								80	<pre>images, labels = images.cuda(), labels.cuda()</pre>
						1		82	<pre>outputs = net(images)</pre>
								86	<pre>correct += (predicted == labels).sum().item()</pre>
1			1		10			98	<pre>net = torchvision.models.resnet50(False).cuda()</pre>
								103	<pre>torch.save(net.state_dict(), PATH)</pre>
TIME	MEMORY	MEMORY	MEMORY	MEMORY	<u>COPY</u>	GPU	GPU	FU	NCTION PROFILE (click to reset order)
	average	peak	timetine		10	utu.	memory	46	ain_igpu.py
_		-		_	10			10	trein
				•	30			38	
1					1			70	test





hover over bars to see breakdowns; click on COLUMN HEADERS to sort.

#### **Al-oriented profilers:** NVidia DLProf

Prc sol	ofiling with Nvidia DLProf requires some light modifications to the urce code
1.	Import the library
	import nvidia_dlprof_pytorch_nvtx as nvtx
2.	Modify the main function add the following initialization code :
	nvtx.init(enable_function_stack=True)
	Wrap the train function with this code:
	with torch.autograd.profiler.emit_nvtx():
3.	Once you have done the modifications, just run:
	dlprof python3 train_1gpu.py

This will create a couple of files, two sqlite files and one nsys-rep.



[fbongiovann	i@
-rw-r	1
-rw-rr	1
[fbongiovann	i@
-rw-rw-r	1
[fbongiovann	i@



```
92 if __name__ == '__main__':
93 start = time.time()
```

```
import torchvision
```

```
PATH = './cifar_net.pth'
trainloader, testloader = create_data_loader_cifar10()
net = torchvision.models.resnet50(False).cuda()
nvtx.init(enable_function_stack=True)
start_train = time.time()
with torch.autograd.profiler.emit_nvtx():
    train(net, trainloader)
end_train = time.time()
# save
torch.save(net.state_dict(), PATH)
# test
test(net, PATH, testloader)
end = time.time()
seconds = (end - start)
seconds_train = (end_train - start_train)
print(f"Total elapsed time: {seconds:.2f} seconds, \
```

```
login02 pytorch-ddp]$ ls -lah *.sqlite
fbongiovanni lxp 101M Mar 22 15:15 dlprof_dldb.sqlite
fbongiovanni lxp 650M Mar 22 15:12 nsys_profile.sqlite
login02 pytorch-ddp]$ ls -lah *.nsys-rep
fbongiovanni lxp 54M Mar 22 15:11 nsys_profile.nsys-rep
login02 pytorch-ddp]$
```

#### **Al-oriented profilers:** NVidia DLProf







	Direction	Ор Туре	Calls	TC Eligible	Using TC
rd/ResNet::_forward_impl/Sequential::_call_impl[8]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C _conv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[8]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C _conv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[8]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C _conv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[8]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C onv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[8]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C onv_forward/conv2d	bprop	conv2d	196	~	~
	bprop	backward	196	<b>~</b>	~
rd/ResNet::_forward_impl/Sequential::_call_impl[8]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C onv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[7]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C onv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[7]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C onv_forward/conv2d	bprop	conv2d	196	~	~
rd/ResNet::_forward_impl/Sequential::_call_impl[7]/Sequential::forward/Bottleneck::_call_impl/Bottleneck::forward/C onv_forward/conv2d	bprop	conv2d	196	~	*

#### **Al-oriented profilers:** NVidia DLProf

System Config	
GPU Count	4
GPU Name(s)	NVIDIA A100-SXM4-40GB NVIDIA A100-SXM4-40GB NVIDIA A100-SXM4-40GB NVIDIA A100-SXM4-40GB
CPU Model	AMD EPYC 7452 32-Core Processor
GPU Driver Version	515.65.01
Framework	PyTorch 1.12.0
CUDA Version	11.7
NSys Version	2022.4.1.21-0db2c85
DLProf CLI Version	v1.8.0
DLProf DB Version	1.8.0
DLProf Viewer Version	1.8.0

	Problem	Recommendation
<b>»</b>	2269 ops were eligible to use tensor cores but none are using FP16	Try enabling AMP (Automatic Mixed Precision). For more information: https://developer.nvidia.com/automatic-mixed-preci
	The following GPU device IDs are not being used: 1 2 3	Make sure to pass the necessary options to docker, the framework, and the model to use all available GPUs
	12.1% of the aggregated run was spent in the dataloader while not simultaneously running on the GPU	Focus on reducing time spent in the training data input process. This could be time spent in file reading, preprocessing an in the dataloader to enable asynchronous data loading. Consider using NVIDIA DALI, a library that is a high performance Learn more here: https://developer.nvidia.com/DALI
<b>»</b>	The aggregated iteration range of 0 to 196 contains a lot of variation	Try limiting the iteration range to a steady range by rerunning with thedatabase option and settingiter_start=27iter_
	Convolution operations were detected but torch.backends.cudnn.benchmark was not enabled.	Try setting torch.backends.cudnn.benchmark = True in your network. For best performance, the input shapes should be r
	GPU Memory is underutilized: Only 12% of GPU Memory is used	Try increasing batch size by 4x to increase data throughput

#### **Ops and Kernel Summaries**

🖌 Ops														
Click an op to see its kernels below														
Show 10 V entries	export to: Excel PDF Clipboard CSV JSON													
GPU Time (ns) 🛛 🔻	CPU Time (ns)	Op ID	Op Name	Direction 🔶	Ор Туре	Calls	♦ TC Eligible	Using TC 👙	Kernel Calls 👙	Data Type 🔶	Stack Trace			
Search GPU Time	Search CPU Time (r	Search Op ID	Search Op Name	Search Direction	Search Op Type	Search Calls	0 or 1	0 or 1	Search Kernel Ca	Search Data Typ	Search Stack Trace			
324,433,581	3,109,669,488	CONV2D_53_BPROP	/module/train/ResNet::_call_impl/ResNet::forward/ResNet:: _forward_impl/Sequential::_call_impl[8]/Seq See more ✔	bprop	conv2d	196	~	~	977	float32	/mnt/tier2/project/lxp/fbongiovanni/python_project s/pytorch-ddp/train_1gpu_prof.py:103 See more ~			
324,310,584	31,528,107	CONV2D_46_BPROP	/module/train/ResNet::_call_impl/ResNet::forward/ResNet:: _forward_impl/Sequential::_call_impl[8]/Seq See more ✔	bprop	conv2d	196	~	~	977	float32	/mnt/tier2/project/lxp/fbongiovanni/python_project s/pytorch-ddp/train_1gpu_prof.py:103 See more V			
323,863,592	31,719,988	CONV2D_50_BPROP	/module/train/ResNet::_call_impl/ResNet::forward/ResNet:: _forward_impl/Sequential::_call_impl[8]/Seq See more ❤	bprop	conv2d	196	~	~	977	float32	/mnt/tier2/project/lxp/fbongiovanni/python_project s/pytorch-ddp/train_1gpu_prof.py:103 See more v			
317,166,280	31,151,502	CONV2D_48_BPROP	/module/train/ResNet::_call_impl/ResNet::forward/ResNet:: _forward_impl/Sequential::_call_impl[8]/Seq See more ❤	bprop	conv2d	196	~	~	979	float32	/mnt/tier2/project/lxp/fbongiovanni/python_project s/pytorch-ddp/train_1gpu_prof.py:103 See more ❤			



#### ision

nd augmentation or file transfer. Set num\_workers > 0 alternative to built-in data loaders and data iterators.

stop=83

relatively stable.

#### Guidance

Understanding GPU utilization and timing details of the operations is the first step in profiling your model.

- To learn more about Tensor cores and Mixed Precision training, visit this site:https://developer.nvidia.com/tensor cores
- You will find resources on how to train networks with mixed precision and make full use of Tensor cores for Tensorflow models here: https://docs.nvidia.com/deeplearning/sdk/mixed-precisiontraining/index.html#training\_tensorflow
- Note that if there are multiple kernels being observed on single op, these are likely performing data transposes to prepare the data for efficient use by tensorcores. Such transposes themselves would not use tensor cores.















#### Destination Earth









## **Destination Earth (DestinE)**

DestinE will provide unique digital modeling capabilities of the Earth to enhance the EU's ability to monitor and model environmental changes, predict extreme events, and adapt EU actions and policies to climate-related challenges.

> Destination Earth factsheet digital-strategy.ec.europa.eu/en/library/destination-earth-factsheet



#### **UNDERSTAND THE PAST, PREDICT THE FUTURE**

Fed by real-world observations, these digital twins let us understand what has happened on Earth – and what will happen in the decades ahead.



#### FUTURE

Predict future changes and test how we might intervene

Simulations

#### **O3b mPOWER: Reimagining Satellites Terabit-level System for Cloud Everywhere on Earth**

#### **O3b MEO Constellation**



- 20-satellites
- 10 user beams per satellite
- Up to 1.6 Gbps per beam
- 9 SES data gateways globally

- ▲ From tens of Mbps to multiple Gbps per beam







#### **O3b mPOWER Constellation**

- 11-satellites
- ▲ Thousands of beams per satellite
- ▲ Flexible data gateways





## **Running Flood Depth Maps at X10 speeds**









#### **Climate change-aware risk analysis**

#### Flood risk commercial maps from:

- **3rd party Earth Observation data** •
- 2-D flood modeling (+ HPC)

Showing flooding risk likelihood of **today** + including climate change effects:

- Likelihood 1:10
- Likelihood 1:100
- Likelihood 1:1000
- Customized

**10x speed** with CPU model version on MeluXina GPU nodes



## Thanks for your attention

luxprovide.lu Follow us in