

GPU memory resource management for Graph Neural Network (GNN) training on large graphs

14^{ème} Journées Informatiques IN2P3/IRFU (14-17 novembre 2022)

Sylvain Caillou

(On behalf of L2IT « Computing, Algorithms and Data » Team)


The rise of geometric ML and representation learning

⇒ Geometric and graph-based ML methods have become one of the hottest fields of AI research




⇒ Graph Neural Networks (GNNs) capture deep geometric and structural patterns in data represented as graph

What does 2022 hold for Geometric & Graph ML?

Michael Bronstein



CNNs can be seen as a specific use case of GNN on regular grid graph



Transformers architectures in Natural Language Processing operate on fully connected graph



In 2021 triumph of Geometric ML and a paradigm shift in structural biology: Breakthrough in prediction of the 3D folding structure of a protein by AlphaFold 2 (deepmind)

Graph-based ML in HEP


Since 2020 Graph-based ML algorithms are becoming increasingly popular for a large number of LHC physics tasks, including:

- Track Reconstruction
- Vertex Reconstruction
- Calorimeter Clustering
- Jet Clustering
- Event Reconstruction (Pileup Rejection, Particle Flow, Jet Assignment)

Now working to integrate Graph ML in production for in-line and off-line data processing systems

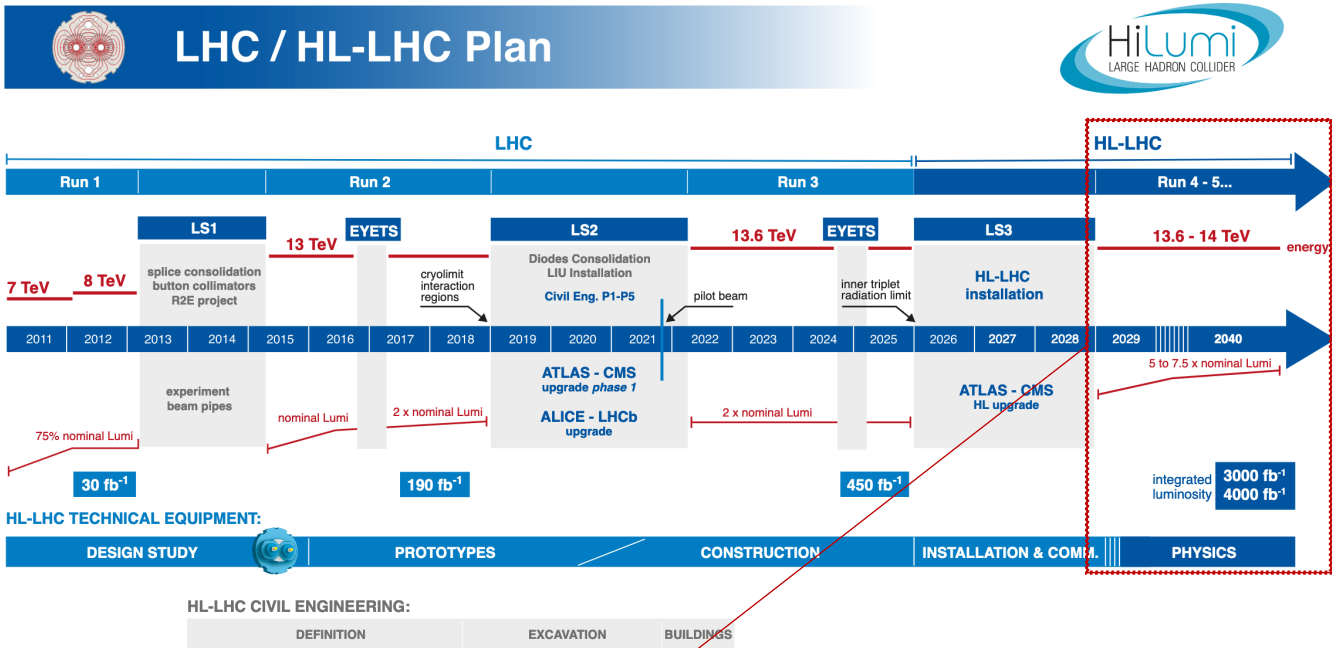
⇒ Starting to be use in Astrophysics, cosmology, nuclear physics, ...

⇒ *Nature of physic experimental data brings to large graph representation and computing challenges for Graph-based ML algorithms*



September/October 2021

HL-LHC and unprecedented challenges for software and computing



⇒ In 2028, the rate of collisions provided by the LHC to the ATLAS detector will be increased by an order of magnitude.

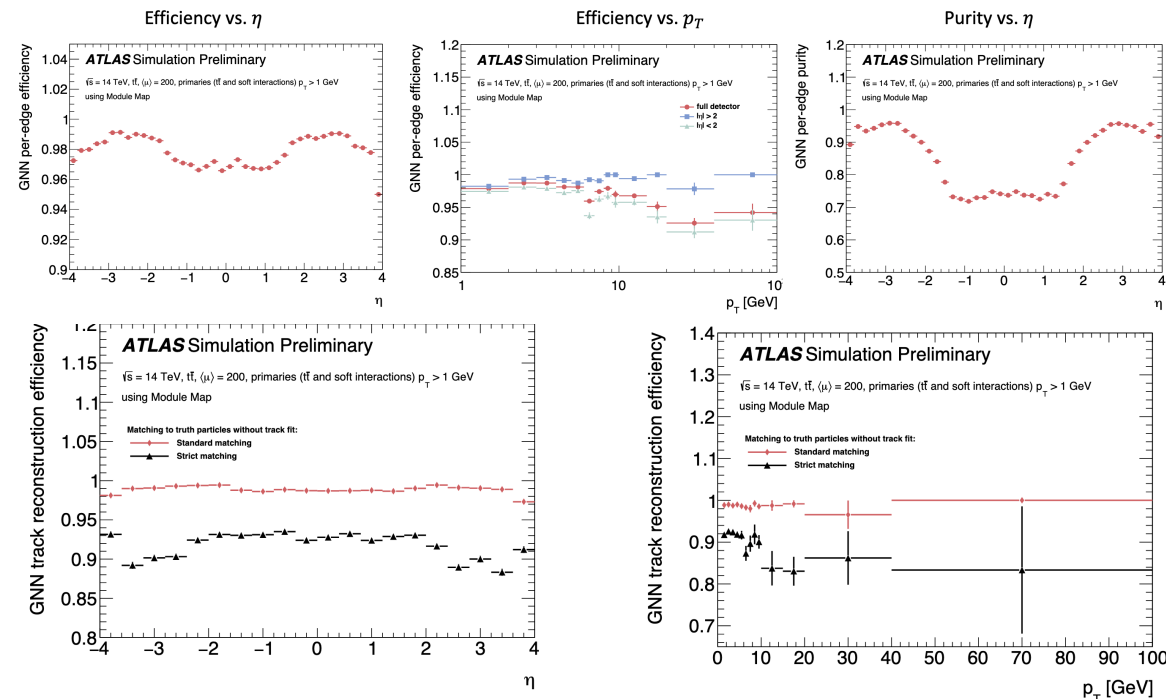

⇒ Existing tracks reconstruction algorithms based on *Kalman Filter* won't fit the combinatorics and will require unrealistic amount of CPU resources

⇒ **The HL-LHC physics program could not be completed (obviously not an option !)**

GNN for charged particles tracking in ATLAS ITk

Graph Neural Networks (GNNs) perform pretty well to learn geometric pattern of the tracks

Close collaboration L2IT & ExatTrkX started in 2021 to construct a GNN-based track reconstruction algorithm for ATLAS ITk



[ATLAS ITk Track Reconstruction with a GNN-based pipeline, C.Rougier et al., CTD 2022](#)
[Graph Neural Network track reconstruction for ATLAS ITk, D. Murnane et al., IML 2022](#)

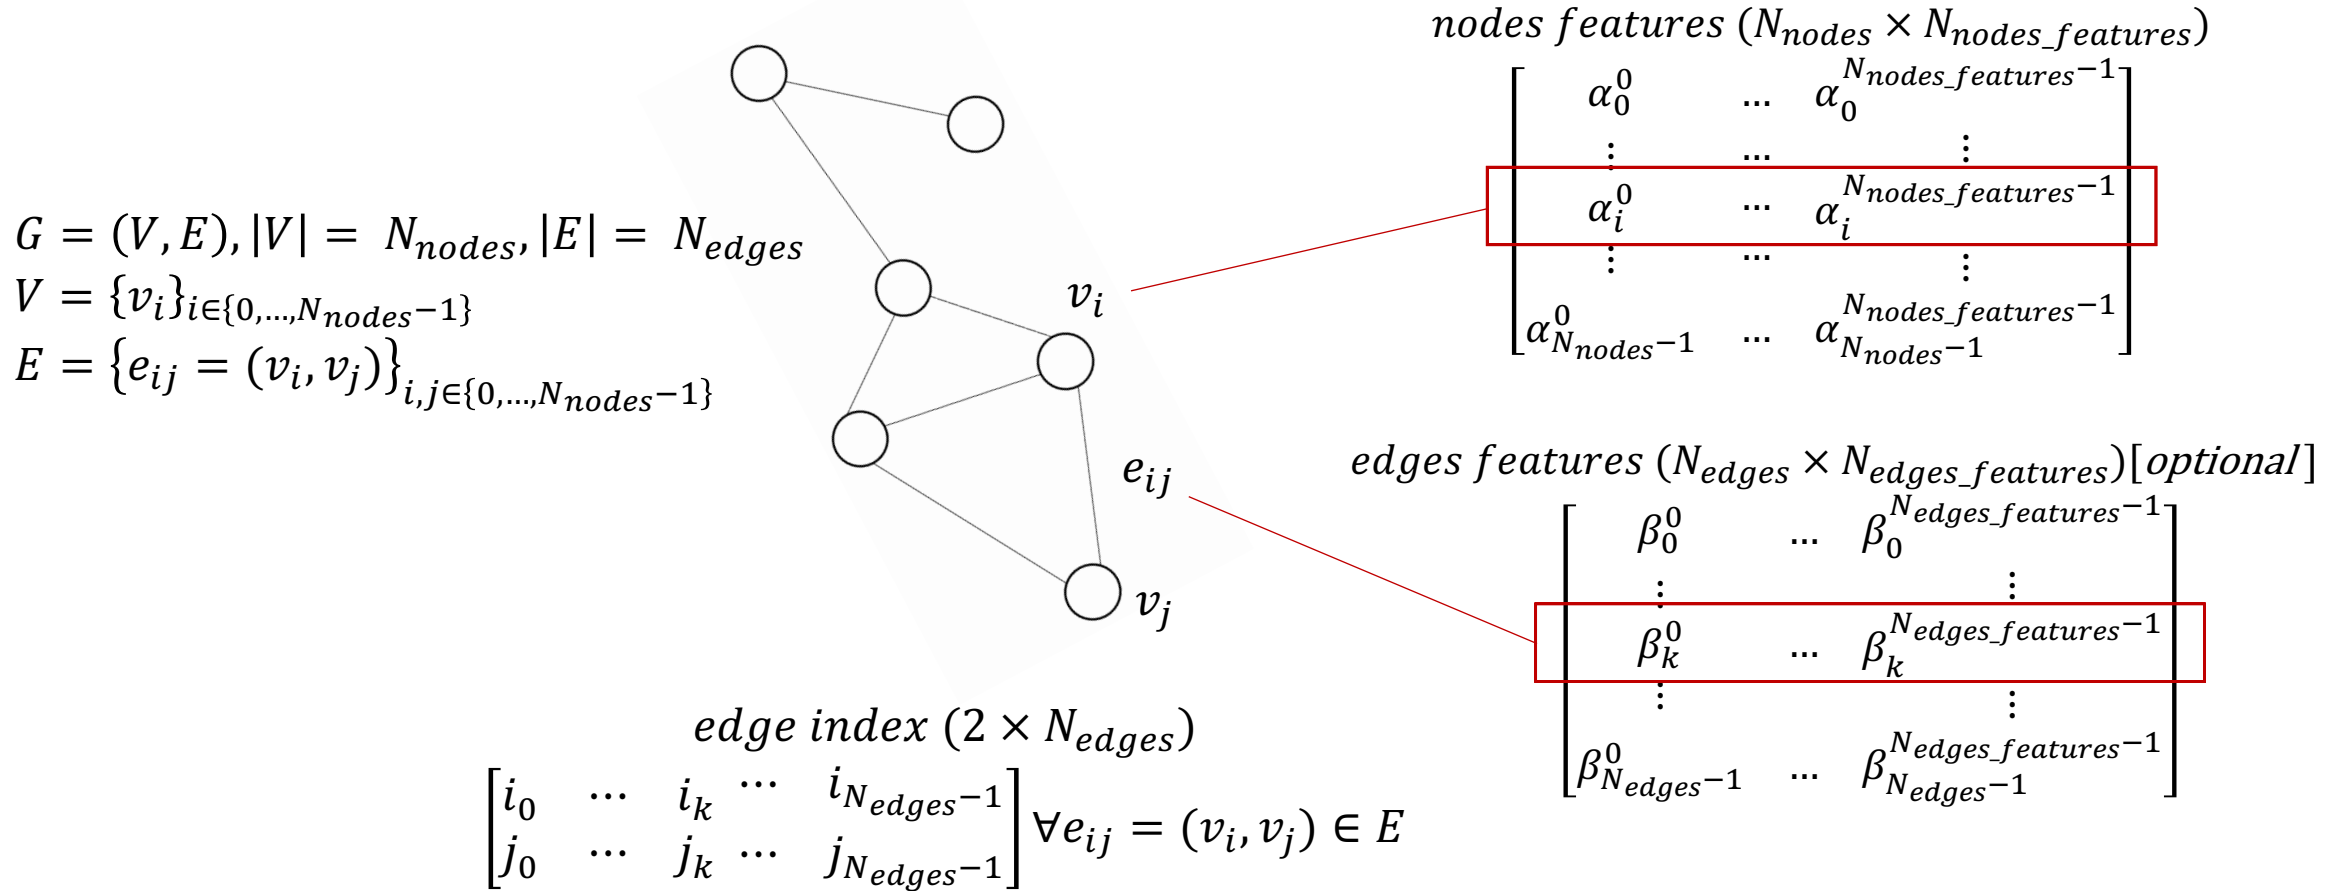
First results on ITk published in 2022 more than encouraging

⇒ GNN-based algorithms now appear as competitive solutions for the future generation of charged particle track reconstruction algorithms which will have to be put into production for the HL-LHC

⇒ But to get this results we have to train our GNNs on very large graphs...

Data representation as graph


A graph represents the relations (*edges*) between a collection of entities (*nodes*).




Graph connectivity is represented as an Adjacency matrix or in COO format

Graph Neural Networks models (Message Passing paradigm)

Update nodes from local node neighborhood



- 1) Update edges from source and destination nodes
- 2) Update nodes from incoming and outgoing edges




Node features $\xrightarrow{\text{Embedding}}$ $\{h_i^0\}$

Edge features $\xrightarrow{\text{Embedding}}$ $\{e_{ij}^0\}$

Edge index $\begin{bmatrix} i_0 & \dots & i_k & \dots & i_{N_{edges}-1} \\ j_0 & \dots & j_k & \dots & j_{N_{edges}-1} \end{bmatrix}$


Encoder layer

Layer $l: \{h_i^l\}, \{e_{ij}^l\}$



GNN^l

Layer $l+1: \{h_i^{l+1}\}, \{e_{ij}^{l+1}\}$



$L \times GNN$ layer

$h_i^l \xrightarrow{MLP}$ Node predictions


$\frac{1}{N_{nodes}} \sum_{i=0}^{N_{nodes}} h_i^l \xrightarrow{MLP}$ Graph predictions

$e_{ij}^l \xrightarrow{MLP}$ Edge predictions

Prediction layer

Graph Neural Networks models (Message Passing paradigm)

Update nodes from local node neighborhood



- 1) Update edges from source and destination nodes
- 2) Update nodes from incoming and outgoing edges



It is necessary to have the whole graph in one piece to compute a complete update

Node features $\xrightarrow{\text{Embedding}} \{h_i^0\}$

Edge features $\xrightarrow{\text{Embedding}} \{e_{ij}^0\}$

Edge index $\begin{bmatrix} i_0 & \dots & i_k & \dots & i_{N_{edges}-1} \\ j_0 & \dots & j_k & \dots & j_{N_{edges}-1} \end{bmatrix}$

Encoder layer

Layer $l: \{h_i^l\}, \{e_{ij}^l\}$

Layer $l+1: \{h_i^{l+1}\}, \{e_{ij}^{l+1}\}$

GNN^l

$L \times GNN$ layer

$h_i^l \xrightarrow{\text{MLP}} \text{Node predictions}$

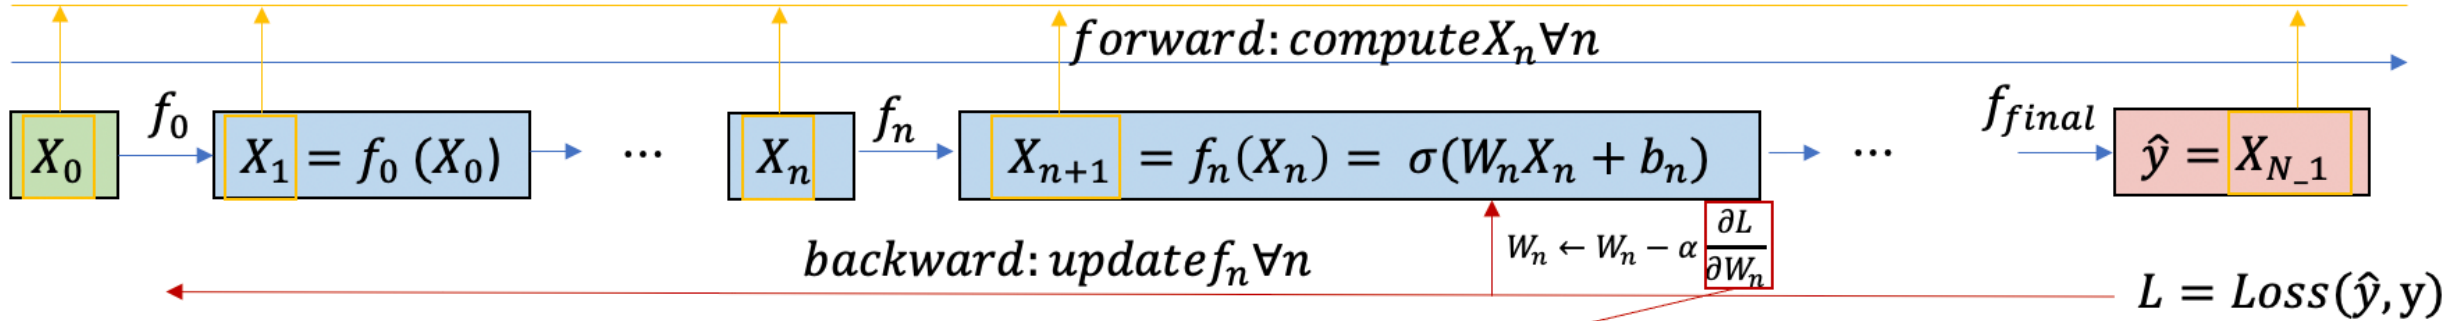
$\frac{1}{N_{nodes}} \sum_{i=0}^{N_{nodes}} h_i^l \xrightarrow{\text{MLP}} \text{Graph predictions}$

$e_{ij}^l \xrightarrow{\text{MLP}} \text{Edge predictions}$

Prediction layer

Neural Networks training and memory consumption

stored for backward gradients computation

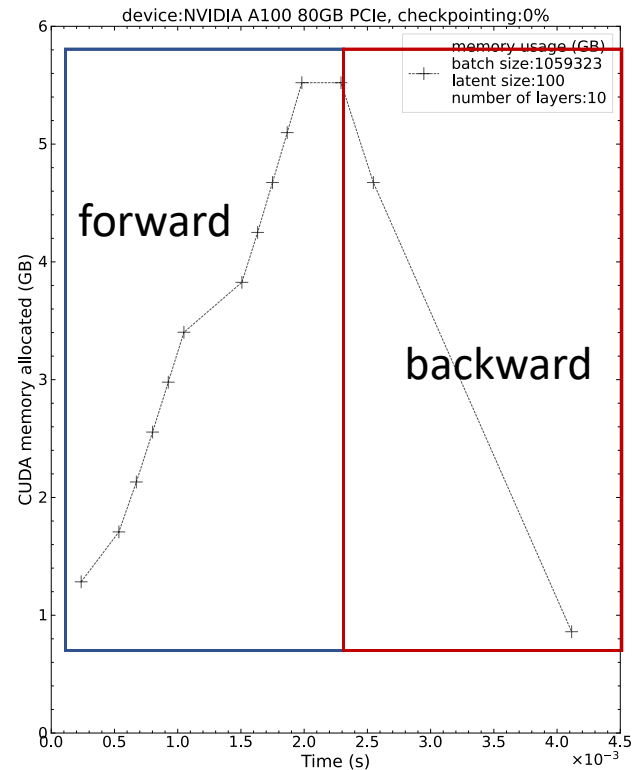


$$\frac{\partial L}{\partial W_n} = \frac{\partial L}{\partial X_{n+1}} \frac{\partial X_{n+1}}{\partial W_n} \rightarrow \frac{\partial \sigma(W_n X_n + b_n)}{\partial W_n} = X_n$$

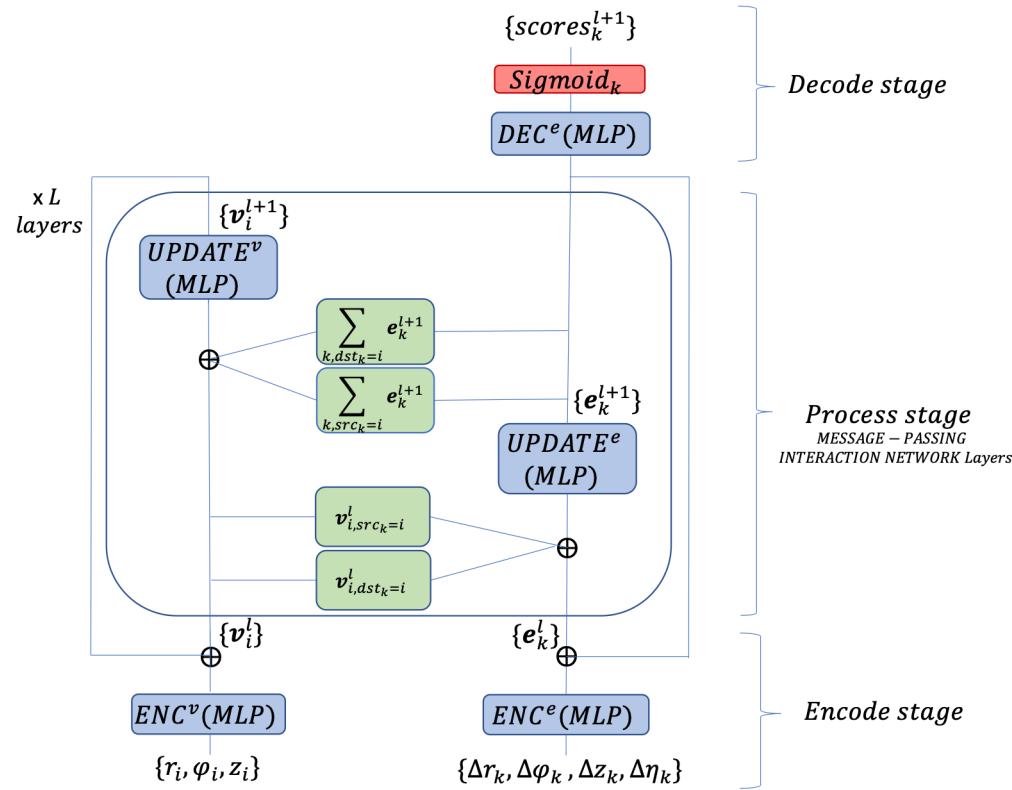
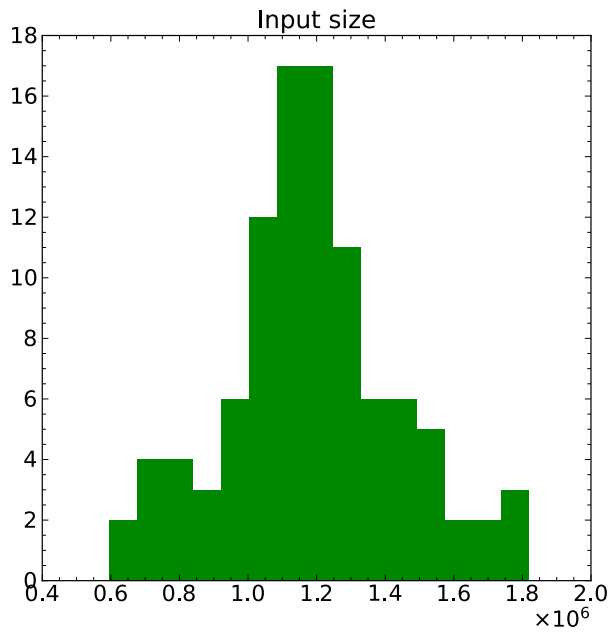
upstream gradient *local gradient*

- ⇒ Store inputs for each layers as hidden states during forward pass
 - ⇒ Consume these stored states during backward pass to compute the gradient
- X_n : *input size* × *latent size* get huge in memory for large input size and/or latent size
- total storage* = *input size* × *latent size* × *numbers of layers*

Example : $(2^{e6} \times 100 \times 100 \times 4) / 1^{e9} \Rightarrow O(80 \text{ GB})$




The never ending battle of the GPU memory consumption



Memory consumption

$$O(10^2 \text{ GB})$$


Neural Networks training and memory consumption




Specific configuration CC support




Host CPU: 1 TB RAM
GPU: 2 x Nvidia RTX8000 48 GB



Host CPU: 500 GB RAM
GPU: 1(+1) x Nvidia A100 80 GB



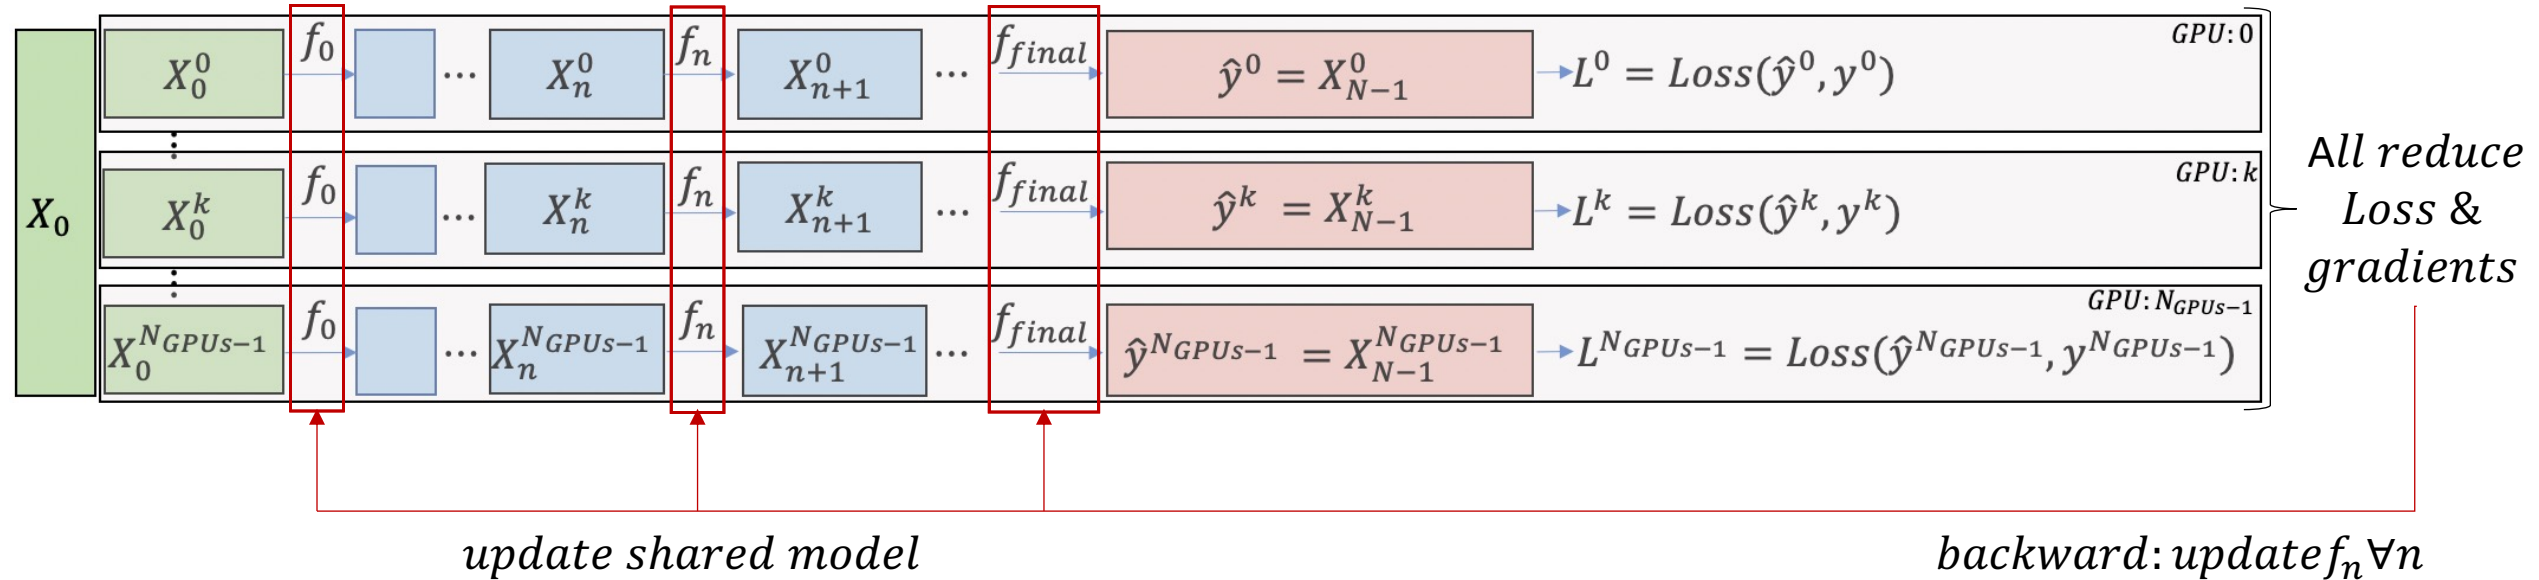
Host CPU: 500 GB RAM
GPU: 2 x Nvidia A100 80 GB



Special thanks to:
Benoît DELAUNAY
Vanessa HAMAR
Xiaomei NIU
Bertrand RIGAUX
...

Distributed Data Parallel

- ⇒ Distribute data
- ⇒ Share model
- ⇒ Computation operate on distributed data between GPUs with copy of the same model (same weights)




- ⇒ Useful to distribute inputs between GPUs
- ⇒ Can be used for GNN models if the graphs are small (i.e. X_0 contains several small graphs)
- If the graphs are large (i.e. X_0 contains only one graph) and it cannot be cut

⇒ So Distributed Data Parallel approach is not a solution in case of large graph which won't fit in only one GPU

Distributed Model Parallel

- ⇒ Distribute model
- ⇒ Share data
- ⇒ Computation of share data is distributed between GPUs: Each GPU compute a part of the model



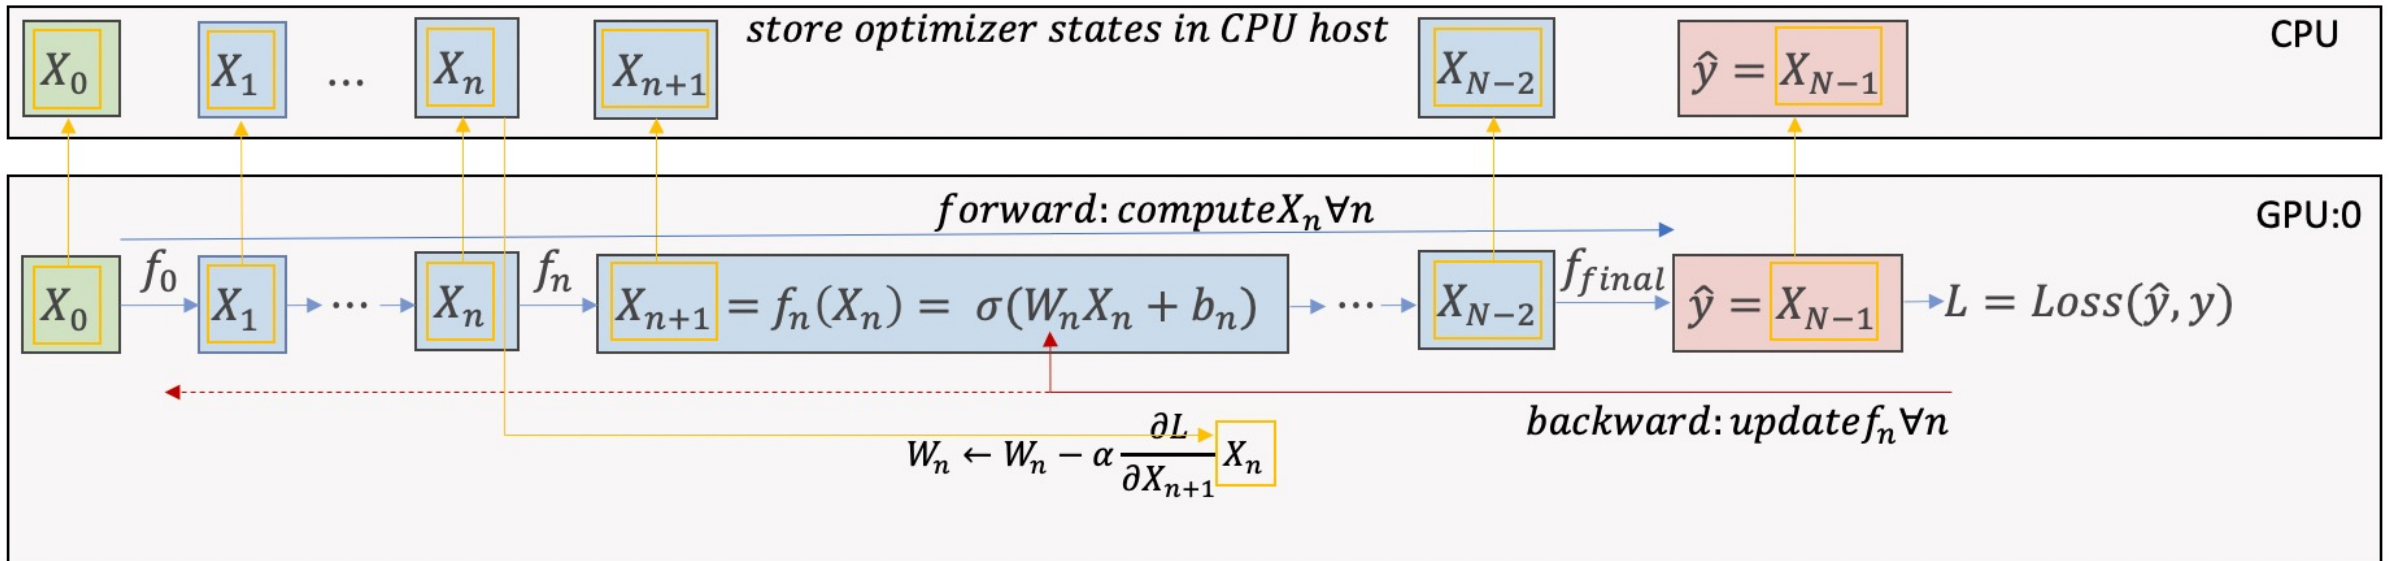
Usefull for models with a very large number of parameters (language model actually up to $O(10^{12})$ parameters)

- ⇒ Distribute between GPUs:
 - ⇒ Memory to store model parameters
 - ⇒ Computation time

⇒ Useless to solve GNN GPU memory consumption on large graph (whatever the number of model parameters)

Optimizer states and gradient offload

- ⇒ Swap on CPU all memory which is not immediately needed for computation
- ⇒ Swap back from CPU to GPU tensors when they are needed




IBM: [Tensorflow Large Model Support \(TFLMS\)](#) and [Pytorch Large Model Support \(PLMS\)](#)

DeepSpeed: [ZeRO-Offload](#)

Pros: **Allows to train very large model on large data**

Cons: **Slow down (a lot) the training because of I/O swapping between CPU host and GPU**

Checkpointing



⇒ Checkpoint layers during forward pass

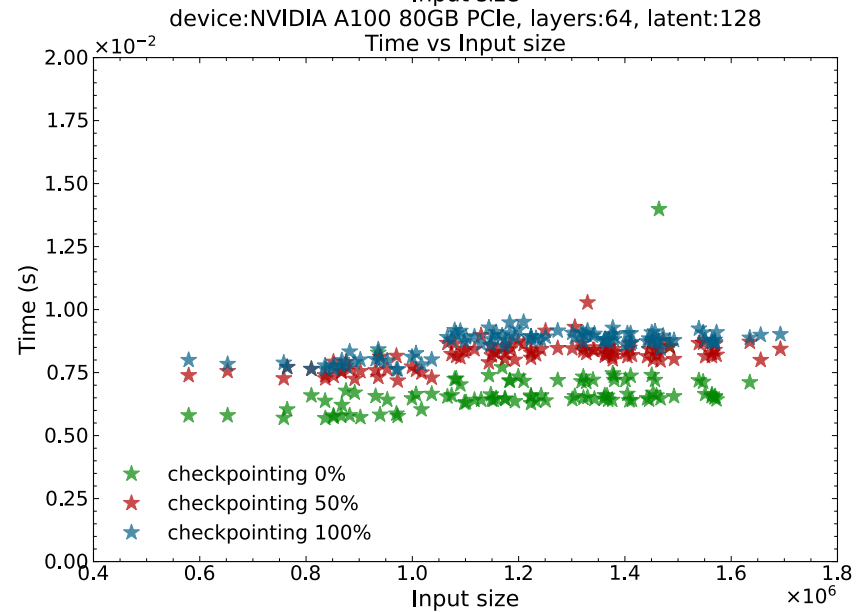
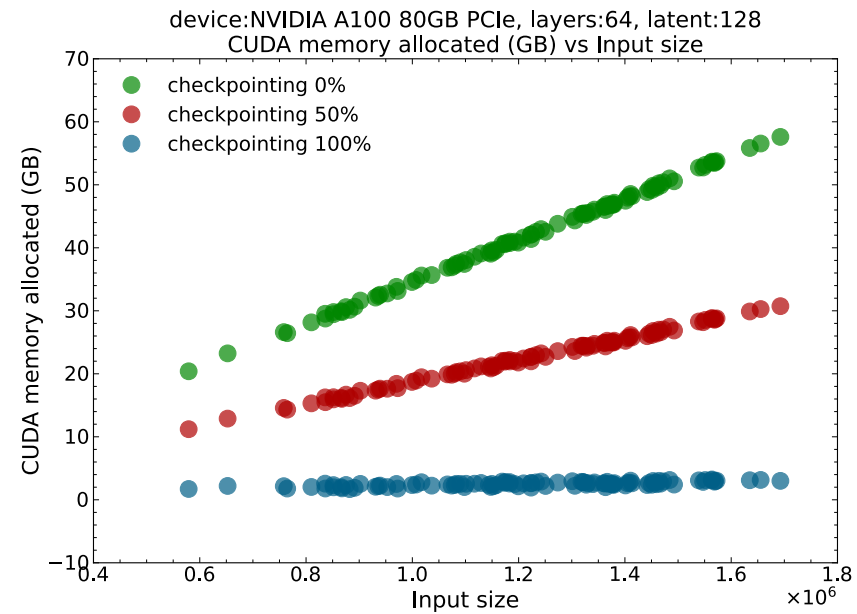
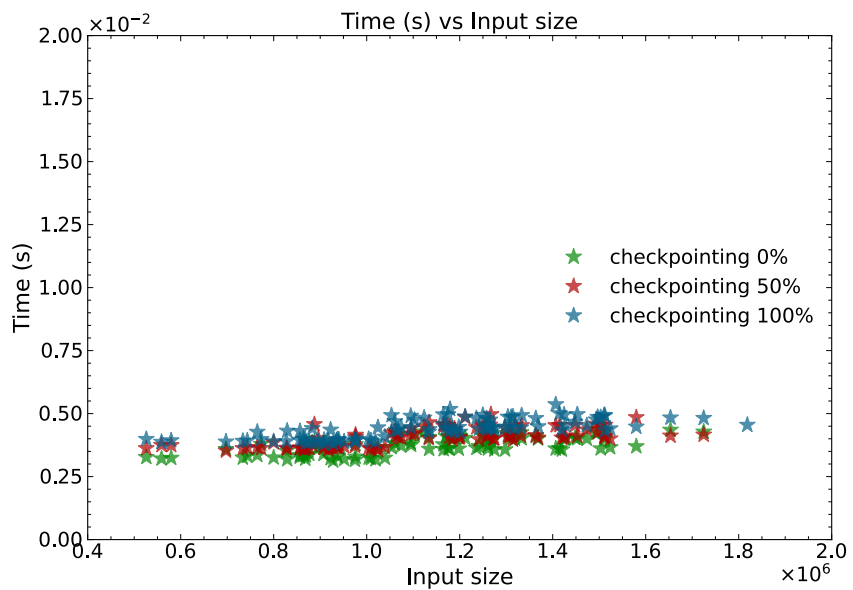
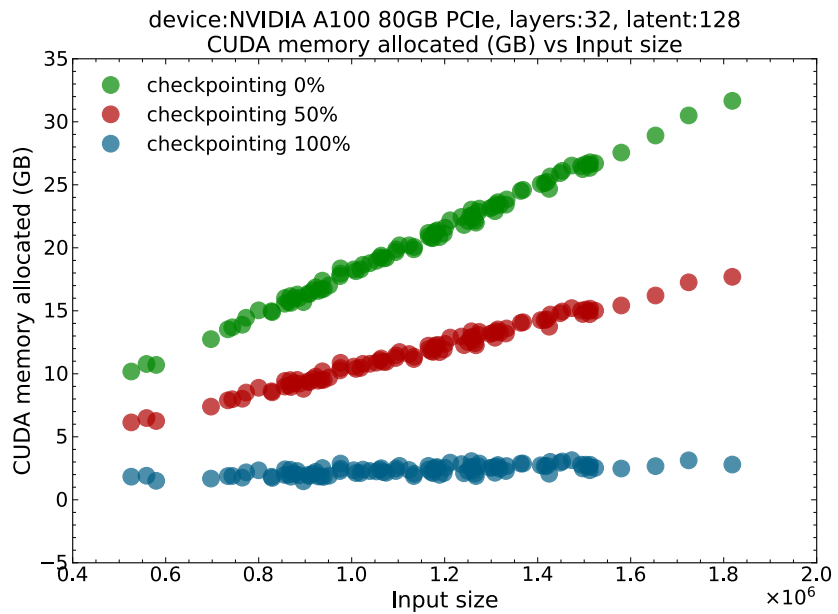
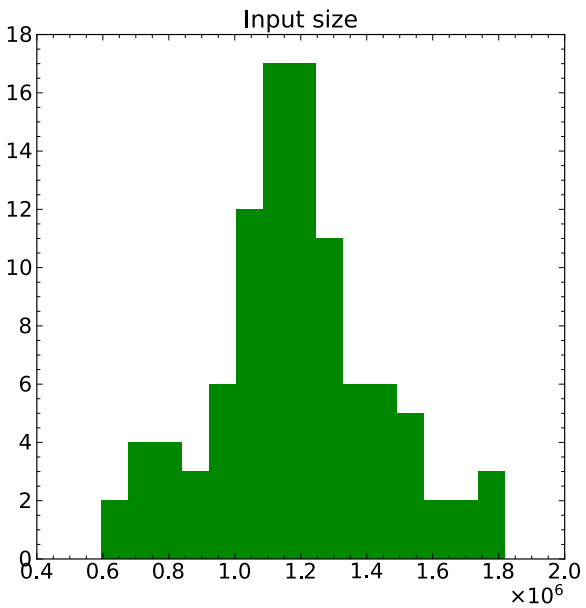
⇒ Intermediate states of checkpointed layers won't be store

⇒ Intermediate states will have to be recompute during backward pass

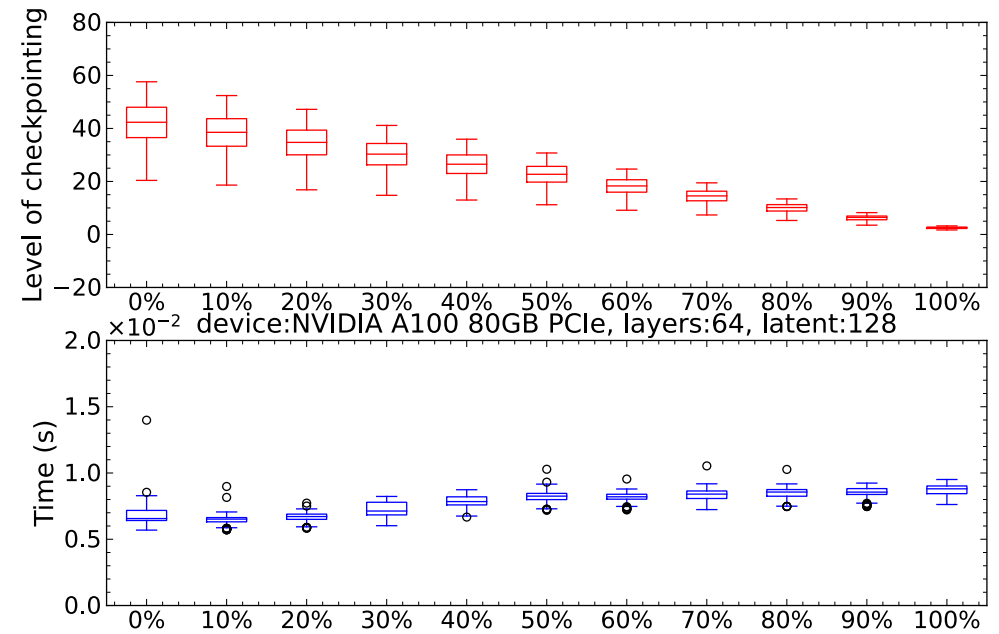


Pros: Save memory (depends of level of checkpointing)

Cons: Come with a cost in time

Checkpointing



Checkpointing



What about sobriety ?

Applying checkpointing (in our use case)
⇒ We divide by $O(10)$ the time for training

Merci de votre attention



Backup

100 layers...

