Martin Erdmann
Peter Fackeldey
Benjamin Fischer
Dennis Noll

Vectorised Neutrino Reconstruction by Computing Graphs

ACAT 2021

Vectorised Computations and Graphs

- Many HEP analyses moving to vectorised computations
- Some computations challenging to vectorise:
  - E.g. Neutrino reconstruction (event-wise fits)
- Can be solved using computing graphs!

Typical Executable:
- Event-by-Event processing

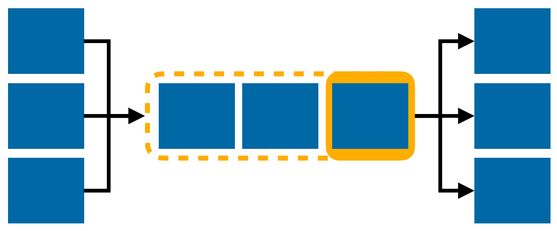

Parallel Computing Graph:
- Processing of many events in parallel
- Many graph libraries optimised for parallel execution

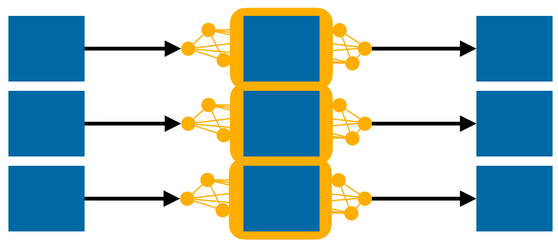

Physics Case: Neutrino Reconstruction

- Analytic reconstruction of longitudinal neutrino momentum:
  - Using Event Kinematics (Lepton, $\not{E}_T$) and W mass
  - Depending on kinematics: real or complex (fit) solution

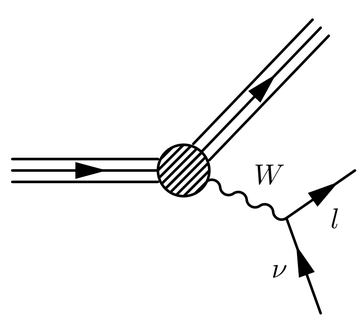

Computing Graph

- Key functionalities are part of the graph:
  - Analytical operations
  - Optimiser & placeholders for fitting
  - Fitting operations:
    - Loop with 100 iterations
    - Adaptive optimisation (Adam)
- All tensors evaluated during execution
- Conditions (choices) guide flow
- Implementation based on Tensorflow
- Graph model highly portable

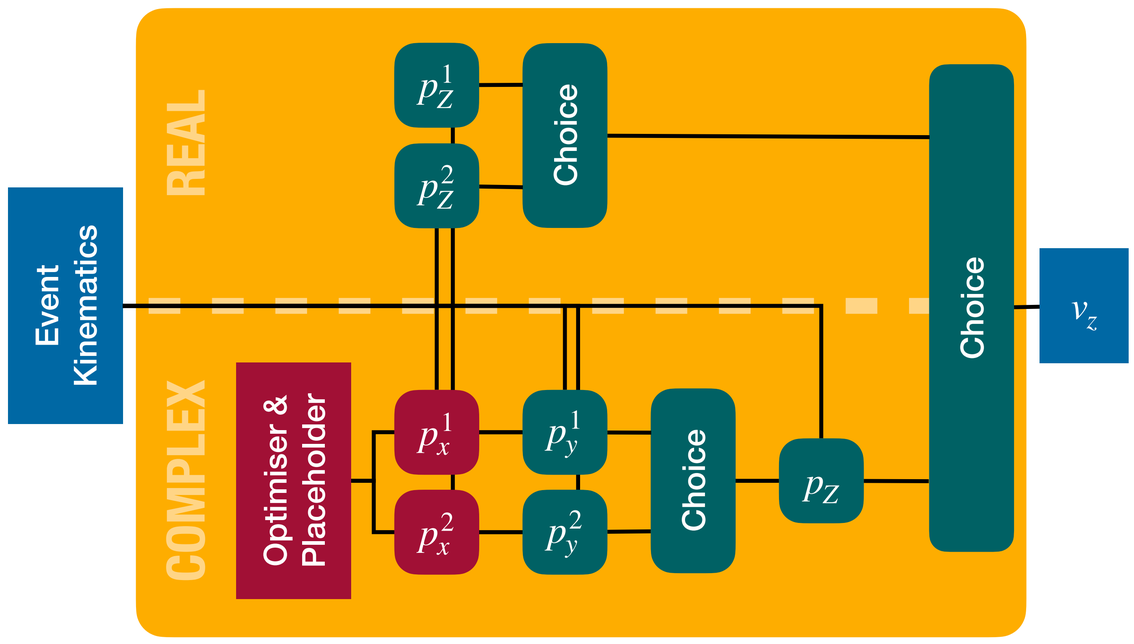

Legend:

Array	Analytical Operator
Tensor	Fitting Operator

Experimental Setup

- Two different hardware setups:

On-board:
- One GPU on each processing node

Server:
- One GPU Server
- One or many processing nodes
- Connected via LAN

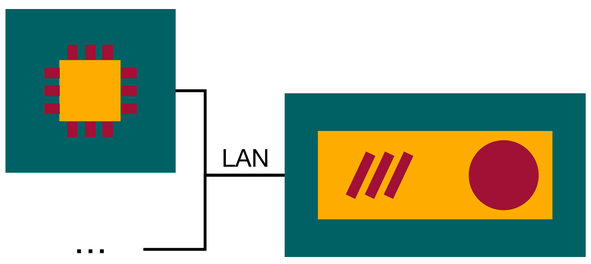

- Pre-heating the GPU:
  - Constant time offset for loading and building the graph
  - Testing two different setups:
    - **Cold**: Setup during computation
    - **Warm**: Setup before computation
- Technical specifications:
  - CPU: 2x Intel Xeon Silver 4216
  - GPU: NVIDIA Quadro RTX 5000
  - Network (LAN): 4x 1GBit

Results

- For a typical analysis with O(100M) events:
  - Parallel computing graph >100x faster than typical executable
  - Speed: On-board GPU > Server GPU = On-board CPU
- All graph calculations done in batches of 100k events

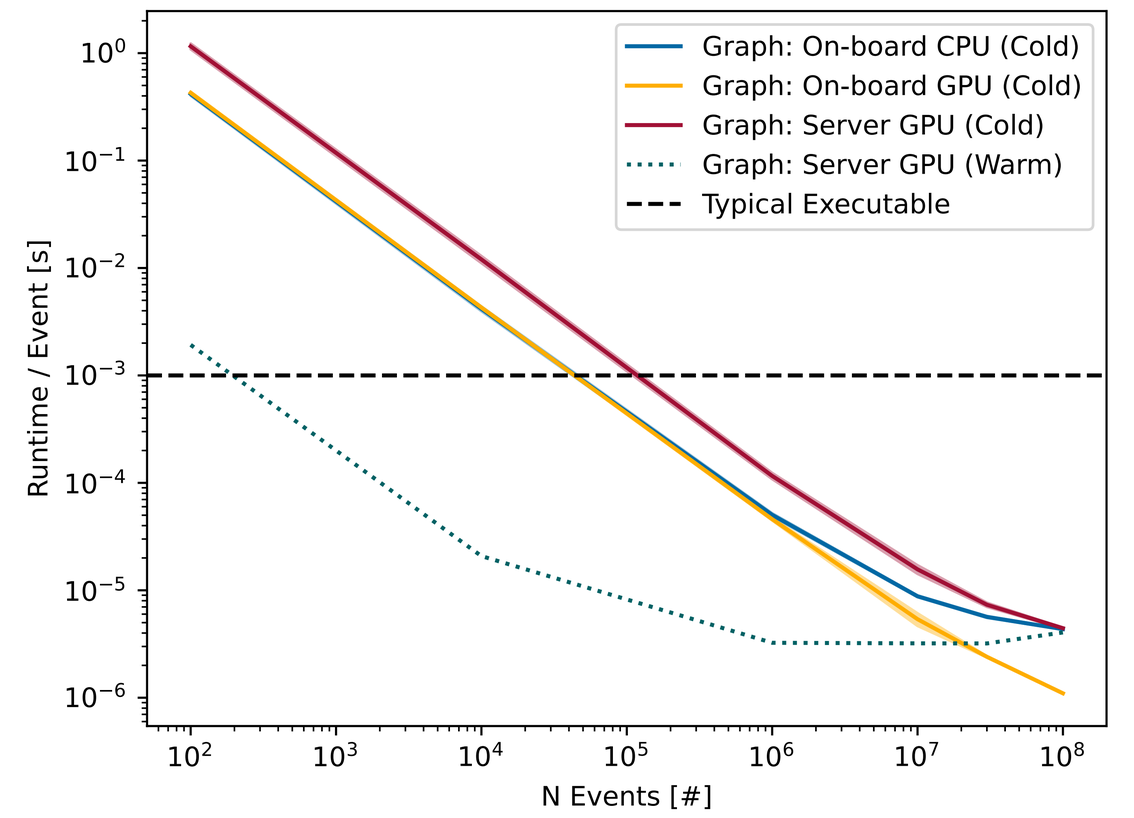

Conclusion

- Even complex HEP computations can be transferred to Parallel Computing Graphs
- Large speedups possible!

Try on binder
https://bndr.it/8yw3z

Physics Institute III A
RWTH Aachen University
Contact: dennis.noll@rwth-aachen.de

SPONSORED BY THE
Federal Ministry of Education and Research

Funded by
DFG Deutsche Forschungsgemeinschaft

RWTH AACHEN UNIVERSITY