



UPPSALA
UNIVERSITET

Prototype of an online track and event reconstruction scheme for the $\overline{\text{PANDA}}$ experiment at FAIR

Project in Computational Science (10 weeks)

Björn Andersson Johan Nordström

Supervisors: Michael Papenbrock, Karin Schönning

March 6, 2017

Uppsala University

Table of contents

1. Project introduction
2. Event reconstruction algorithms
3. Implementation details
4. Performance analysis
5. Summary and outlook

Project introduction

- \bar{P} ANDA experiment: event rates up to 20MHz, event size 10kB

- \bar{P} ANDA experiment: event rates up to 20MHz, event size 10kB
- Data rate: 200GB/s

- \bar{P} ANDA experiment: event rates up to 20MHz, event size 10kB
- Data rate: 200GB/s
- Reduce data rate with a completely software based trigger

- **Project goal:** Prototype of an online track and event reconstruction scheme for the $\bar{\text{P}}\text{ANDA}$ experiment.

- **Project goal:** Prototype of an online track and event reconstruction scheme for the $\bar{\text{P}}\text{ANDA}$ experiment.
- **Requirements:**

- **Project goal:** Prototype of an online track and event reconstruction scheme for the $\bar{\text{P}}\text{ANDA}$ experiment.
- Requirements:
 - Fast and efficient

- **Project goal:** Prototype of an online track and event reconstruction scheme for the $\bar{\text{P}}\text{ANDA}$ experiment.
- Requirements:
 - Fast and efficient
 - Scalable

- Generation of semi-realistic data

- Generation of semi-realistic data
- Prototype event reconstruction algorithms

- Generation of semi-realistic data
- Prototype event reconstruction algorithms
- Benchmarking of algorithms

Relevant sub-detectors

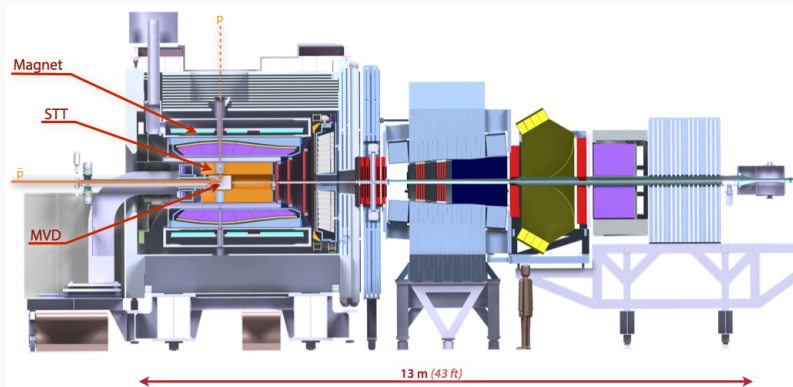


Figure 1: The \bar{P} ANDA Detector

Straw Tube Tracker

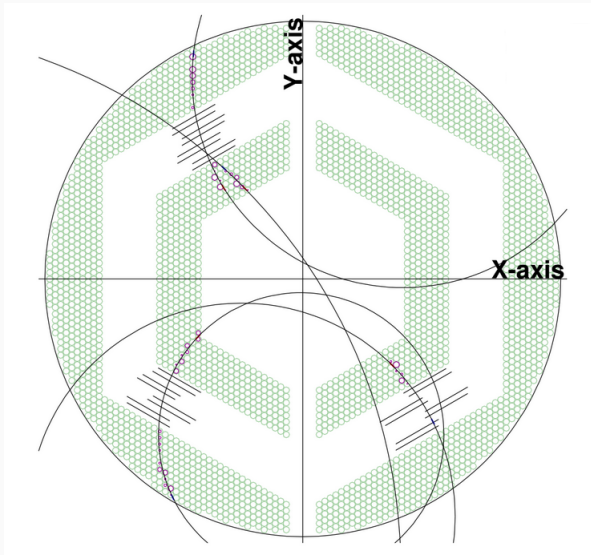


Figure 2: Cross section of STT with trajectories

Generated data structure

- Sequence of STT hits

```
class Hit {  
    int strawID;  
    float timeStamp;  
};
```


Generated data structure

- Sequence of STT hits

```
class Hit {  
    int strawID;  
    float timeStamp;  
};
```

- Hyperons → displaced vertices

Generated data structure

- Sequence of STT hits

```
class Hit {  
    int strawID;  
    float timeStamp;  
};
```

- Hyperons → displaced vertices
- Hits generated in lines and V-shapes

Generated data structure

- Sequence of STT hits

```
class Hit {  
    int strawID;  
    float timeStamp;  
};
```

- Hyperons → displaced vertices
- Hits generated in lines and V-shapes
- Additional noise hits.

Event intermixing

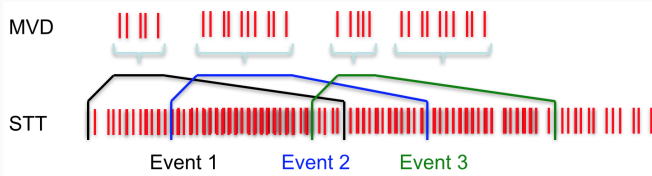


Figure 3: Illustration by T. Stockmanns.

Event intermixing is simulated by incrementing the timestamps with:

$$t \rightarrow t + \Delta t + \xi, \quad \xi \sim \mathcal{N}(0, \sigma^2).$$

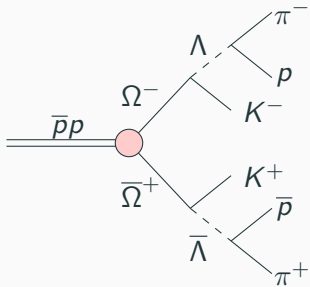


Figure 4: Proton Anti-proton reaction.

Event reconstruction algorithms

- Three different algorithms
 - Clustering of STT hits based on position in space and time

- Three different algorithms
 - Clustering of STT hits based on position in space and time
 - Track reconstruction

- Three different algorithms
 - Clustering of STT hits based on position in space and time
 - Track reconstruction
 - Displaced vertex detection (at clustering stage)

- Proposed solution: Split hit stream into bins and cluster separately

- Proposed solution: Split hit stream into bins and cluster separately
- Assumption (*):

- Proposed solution: Split hit stream into bins and cluster separately
- Assumption (*):
 - Time difference between two STT hits in the same event can not exceed STT response time (≈ 250 ns)

- Proposed solution: Split hit stream into bins and cluster separately
- Assumption (*):
 - Time difference between two STT hits in the same event can not exceed STT response time (≈ 250 ns)
- Place STT hits in multiple different bins so that (*) holds.

- Track reconstruction through modified version of STTCellTrackFinder (Schumann, FZ Jülich)

- Track reconstruction through modified version of STTCellTrackFinder (Schumann, FZ Jülich)
- Simplification: Straight particle trajectories (Line fitting)

Cellular automaton

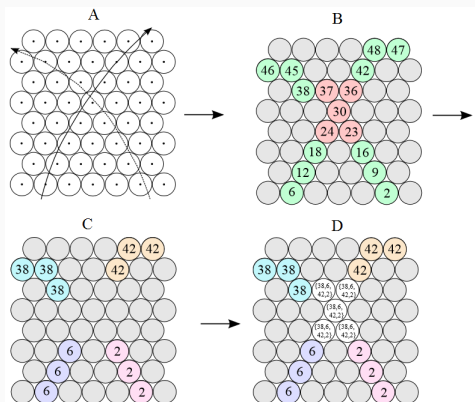
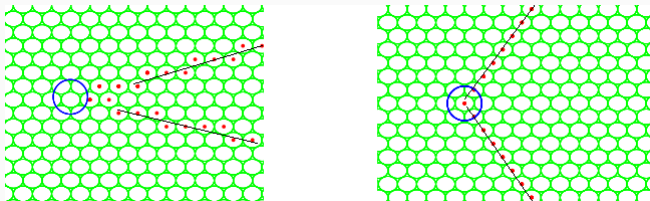


Figure 5: Illustration showing four steps in the track finding cellular automaton.

Displaced vertex detection

Two cases:

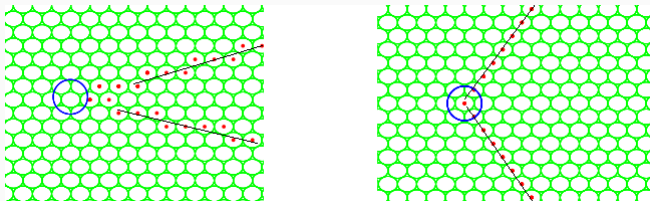
- An ambiguous node with exactly two possible IDs



Displaced vertex detection

Two cases:

- An ambiguous node with exactly two possible IDs
- A tracklet with a high mean square error to the curve fit



Implementation details

- Various computing hardware is in consideration

- Various computing hardware is in consideration
- Prototype uses CPU cluster

- Various computing hardware is in consideration
- Prototype uses CPU cluster
 - Suitable for developed algorithms

- Various computing hardware is in consideration
- Prototype uses CPU cluster
 - Suitable for developed algorithms
 - Efficient to prototype in

- Various computing hardware is in consideration
- Prototype uses CPU cluster
 - Suitable for developed algorithms
 - Efficient to prototype in
 - **Benchmarking hardware available**

- Independent C++ project

- Independent C++ project
- Hybrid parallelization

- Independent C++ project
- Hybrid parallelization
 - Message Passing Interface (MPI) (non-shared memory)

- Independent C++ project
- Hybrid parallelization
 - Message Passing Interface (MPI) (non-shared memory)
 - OpenMP (shared memory)

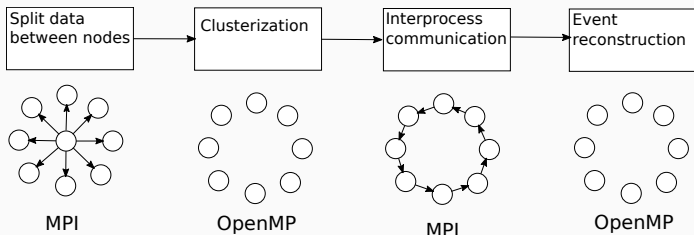
- Independent C++ project
- Hybrid parallelization
 - Message Passing Interface (MPI) (non-shared memory)
 - OpenMP (shared memory)
- Applicable in multi-core CPU environments

- Independent C++ project
- Hybrid parallelization
 - Message Passing Interface (MPI) (non-shared memory)
 - OpenMP (shared memory)
- Applicable in multi-core CPU environments
- Modular

- Hit data sequence is divided into stacks

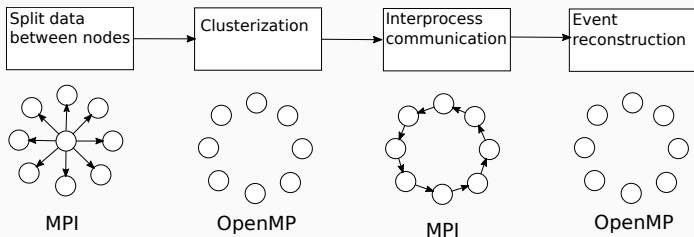
Parallel structure

- Hit data sequence is divided into stacks
- Stacks processed in four main stages



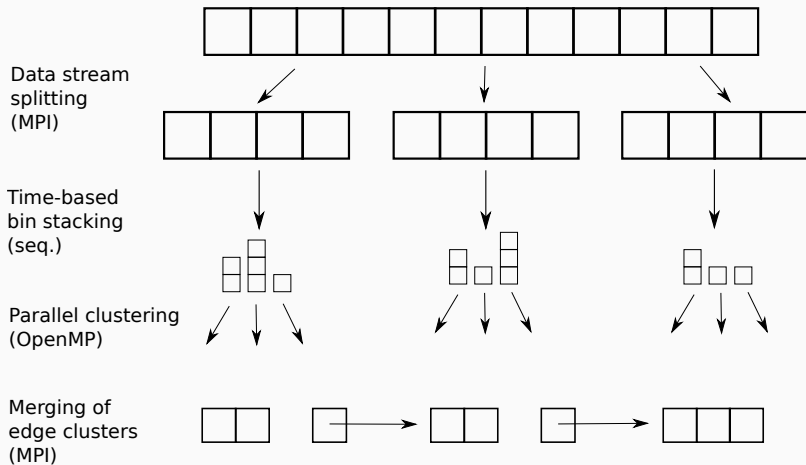
Parallel structure

- Hit data sequence is divided into stacks
- Stacks processed in four main stages



- Load balancing by master/slave model

Parallel clustering pipeline



Performance analysis

- MPI and OpenMP components are analyzed separately

- MPI and OpenMP components are analyzed separately
- Performance metrics:

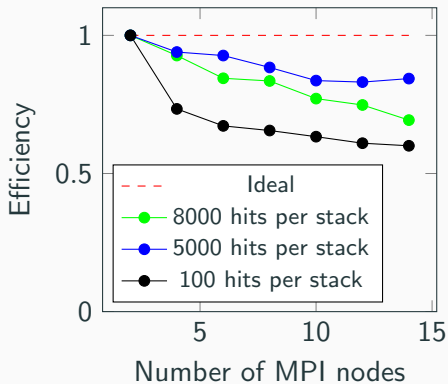
- MPI and OpenMP components are analyzed separately
- Performance metrics:
 - Efficiency: $E(N) = \frac{\tau_1}{N\tau_N}$

- MPI and OpenMP components are analyzed separately
- Performance metrics:
 - Efficiency: $E(N) = \frac{\tau_1}{N\tau_N}$
 - Speedup: $S(N) = \frac{\tau_1}{\tau_N}$.

- MPI and OpenMP components are analyzed separately
- Performance metrics:
 - Efficiency: $E(N) = \frac{\tau_1}{N\tau_N}$
 - Speedup: $S(N) = \frac{\tau_1}{\tau_N}$.
 - τ_1 and τ_N are the execution times using 1 and N computing nodes respectively for a fixed problem size.

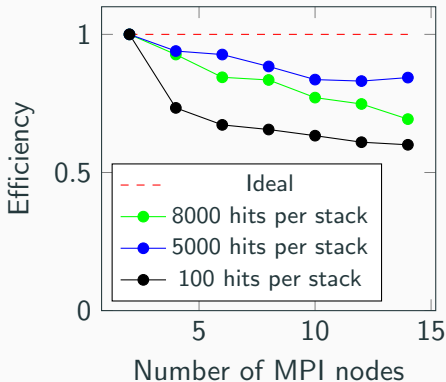
Non-shared memory efficiency (MPI)

- 5000 STT hits per stack
→ optimal efficiency

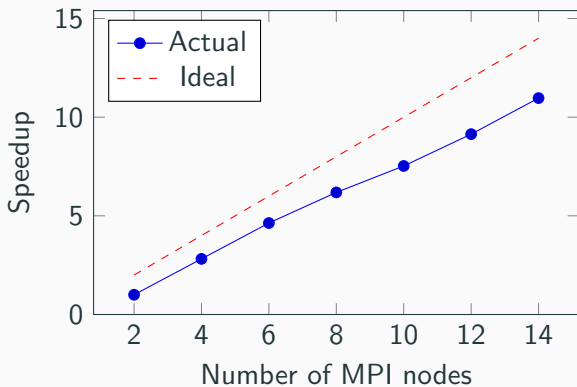


Non-shared memory efficiency (MPI)

- 5000 STT hits per stack
→ optimal efficiency
- Good efficiency!

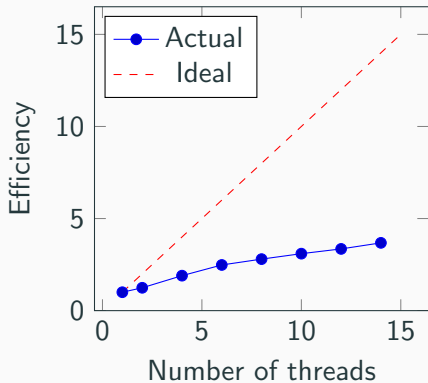


Non-shared memory speedup (MPI)



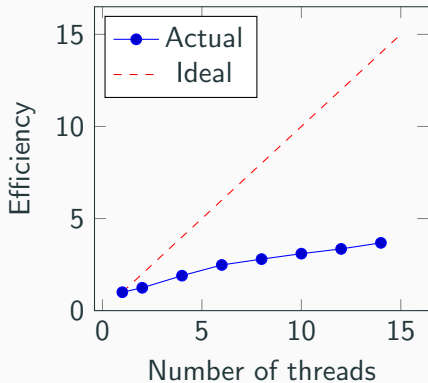
Shared memory speedup (OpenMP)

- Quite low speedup



Shared memory speedup (OpenMP)

- Quite low speedup
- Will improve with more realistic algorithms.



Parallel clustering run time

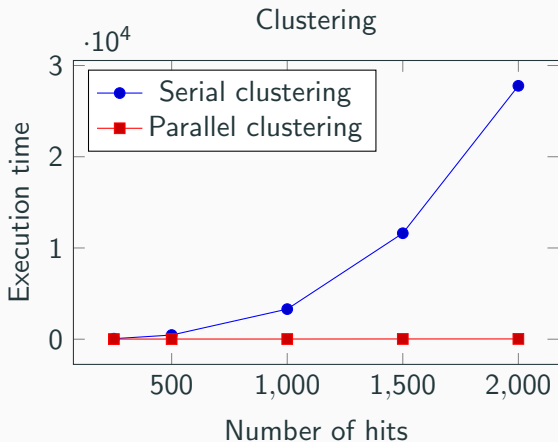


Figure 6: Execution time of the clustering algorithms.

Summary and outlook

- Prototype of an online track and event reconstruction scheme.

- Prototype of an online track and event reconstruction scheme.
- Hybrid parallelization model (MPI/OpenMP)

- Prototype of an online track and event reconstruction scheme.
- Hybrid parallelization model (MPI/OpenMP)
- Possible extension of STTCellTrackFinder

- Prototype of an online track and event reconstruction scheme.
- Hybrid parallelization model (MPI/OpenMP)
- Possible extension of STTCellTrackFinder
- Promising scaling.

- Modify the event reconstruction algorithms to handle realistic data.

- Modify the event reconstruction algorithms to handle realistic data.
- Integrate the system with the $\bar{\text{P}}\text{ANDA}$ simulation framework (PandaRoot).

- Modify the event reconstruction algorithms to handle realistic data.
- Integrate the system with the $\bar{\text{P}}\text{ANDA}$ simulation framework (PandaRoot).
- Investigate the use of a dynamic load balancing scheme.

- Modify the event reconstruction algorithms to handle realistic data.
- Integrate the system with the $\bar{\text{P}}\text{ANDA}$ simulation framework (PandaRoot).
- Investigate the use of a dynamic load balancing scheme.
- More thorough performance analysis on larger scale systems

- Modify the event reconstruction algorithms to handle realistic data.
- Integrate the system with the $\bar{\text{P}}\text{ANDA}$ simulation framework (PandaRoot).
- Investigate the use of a dynamic load balancing scheme.
- More thorough performance analysis on larger scale systems
- Look into frameworks that allow for streaming data processing.

- Modify the event reconstruction algorithms to handle realistic data.
- Integrate the system with the $\bar{\text{P}}\text{ANDA}$ simulation framework (PandaRoot).
- Investigate the use of a dynamic load balancing scheme.
- More thorough performance analysis on larger scale systems
- Look into frameworks that allow for streaming data processing.
- The report is available as an internal $\bar{\text{P}}\text{ANDA}$ document.

Questions?

Inter-process communication

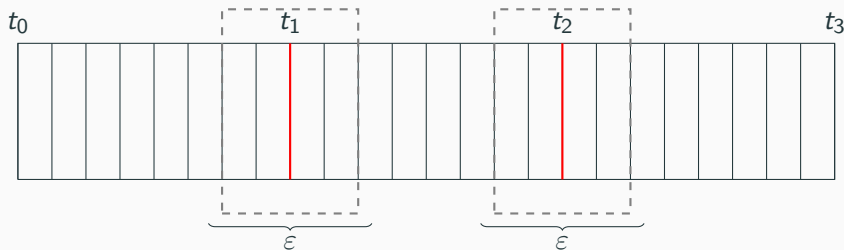


Figure 7: Ordered hit sequence by time.

- Stacks in different nodes \rightarrow need interprocess-communication.
- Can be ignored if one allows to throw away a proportion of events.