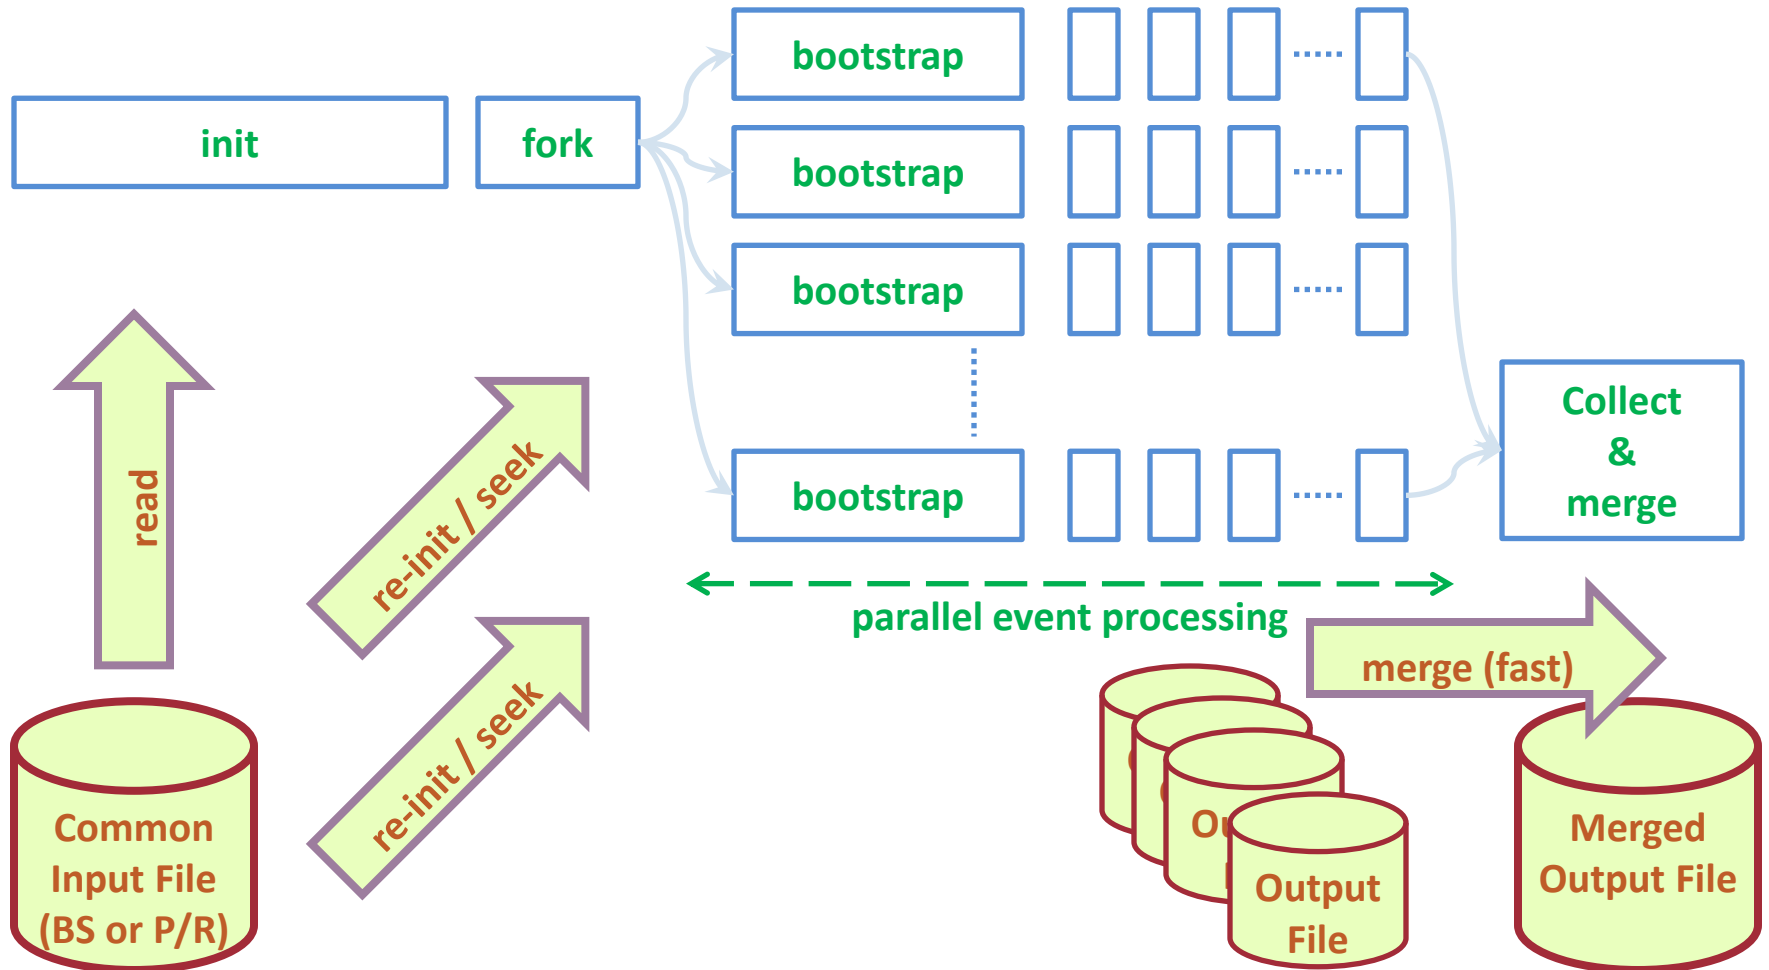


I/O extensions for athenaMP

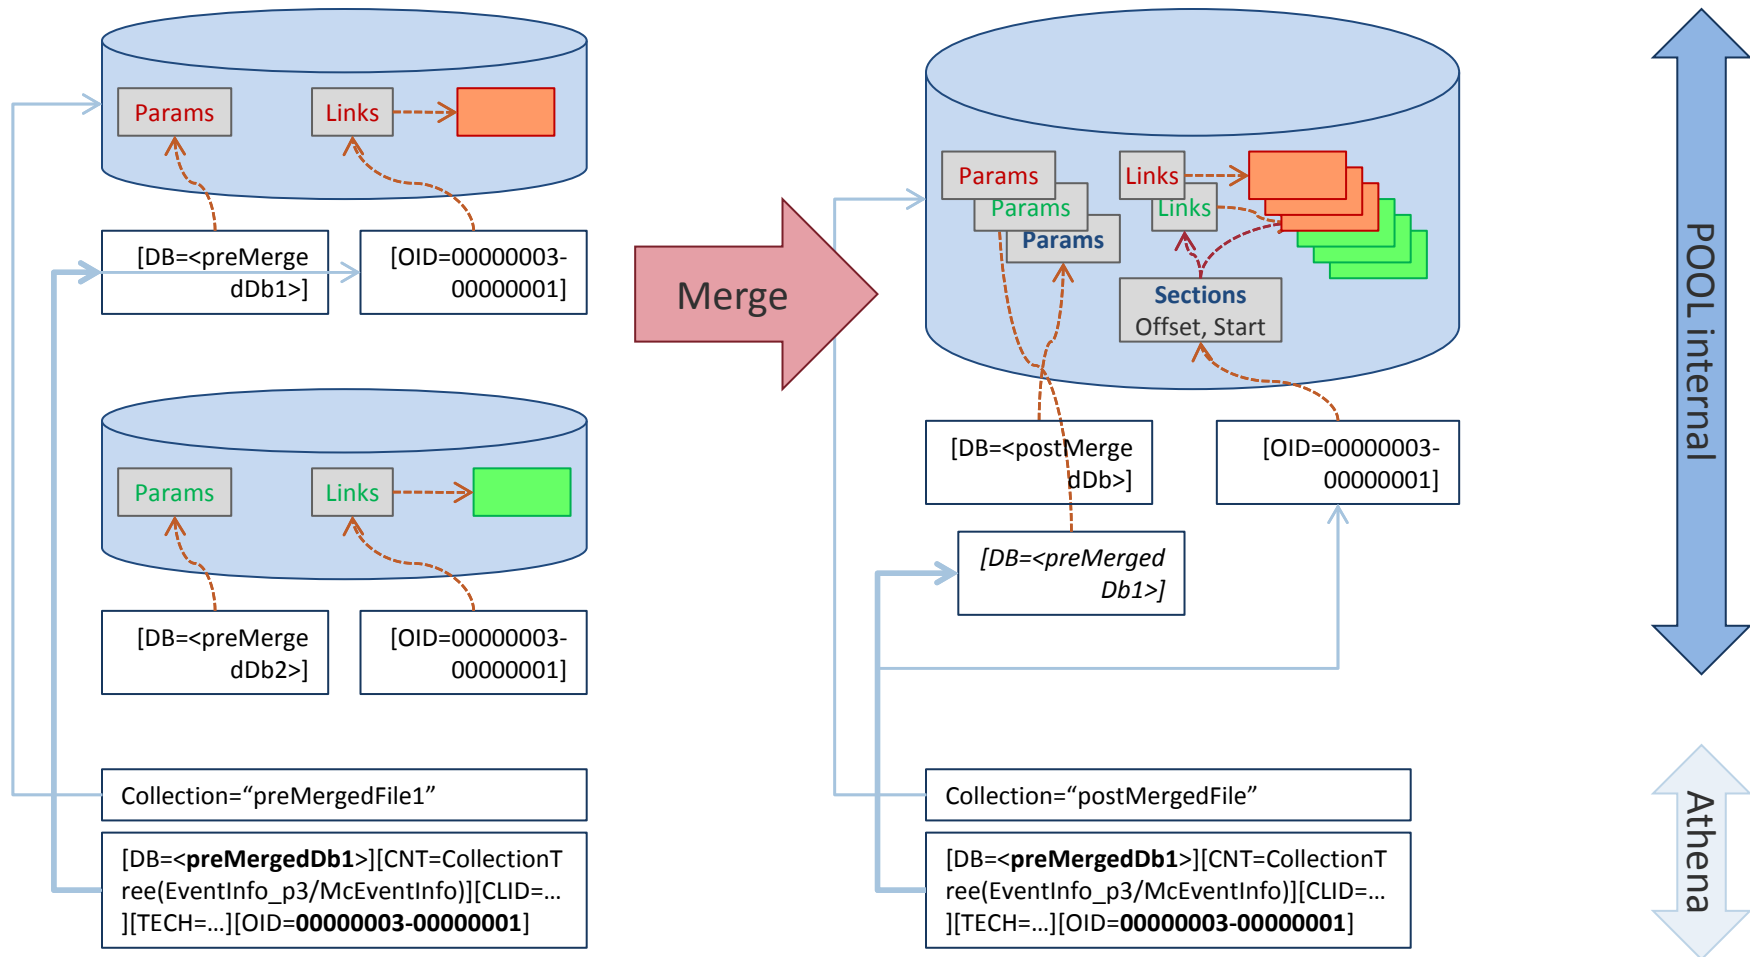
- POOL / ROOT fast merge
- ByteStreamEventSelector
 - ISeekable
 - IIoComponent

Peter van Gemmeren (ANL)

AthenaMP



POOL / ROOT fast merge



POOL / ROOT fast merge

- With help of Markus Frank, we got POOL/ROOT fast merge to work for ATLAS event data some time ago:
 - In POOL 2.9.9 -> **Athena 15.9.0**
 - With new redirection of preMerged GUIDs and utilization of POOL ##Sections table, POOL/ROOT fast merge works even for external tokens.
 - POOL/ROOT fast merge will not (can not) summarize metadata
 - Instead appends multiple entries in the POOL container for metadata
 - Usually, there was only one such entry
 - Most metadata clients only retrieve one entry per file
 - I/O framework supports multiple metadata records
 - Using StoreGate versioning
- Toy Test in AthenaPoolExample works.
- Merging event ESD and AOD appears to work
 - Not tested in detail yet
- Metadata clients have been updated.
 - Not yet tested (other than by developer)

POOL / ROOT fast merge

- There is a simple toy test in **Database/AthenaPOOL/AthenaPoolExample/AthenaPoolExampleAlgorithms** testing the fast merge:
 - Produces two files SimplePoolFileA.root and SimplePoolFileB.root and merges them:
 - **poolMerge -o SimplePoolFile2.root -i SimplePoolFileA.root -i SimplePoolFileB.root**
 - Delete the original files:
 - **rm SimplePoolFileA.root SimplePoolFileB.root**
 - Create POOL file catalog entry for merged file:
 - **pool_insertFileToCatalog -u xmlcatalog_file:Catalog1.xml SimplePoolFile2.root**
 - Read the merged file and produce a down stream file
 - Read downstream file and backNavigate to merged file.
- Some Test were done with ESD / AOD.
 - Catalog entries biggest cause for confusion

ByteStreamEventSelector

- For athenaMP to do data reconstruction, it needs to be able to read RAW in ByteStream format.
- Mother process will initialize **ByteStreamEventSelector** and **ByteStreamEventStorageInputSvc**.
 - Optionally process some events
- After Worker processes are forked, their **ByteStreamEventSelector** and **ByteStreamEventStorageInputSvc** need to be reinitialized:
 - Need to be **IloComponent** and register (themselves and their files) with **IloComponentMgr**
 - Need to implement **io_reinit()** to update file handles and **reinit** to reinitialize services.
- Each Worker processes events by (forward) seeking to them.
 - Both, **ByteStreamEventSelector** and **ByteStreamEventStorageInputSvc** need to be **ISeekable** and implement a **seek()** function
- Requested in October, most recent bug fix yesterday...

Outlook / Conclusion

- Fast POOL merge works in initial tests and promises large improvements for athenaMP.
 - However, utilizing fast merge in athenaMP poses its very own challenges
 - Most are related to the POOL File Catalog handling
 - No showstopper so far, but work is not yet complete
- Using athenaMP for ByteStream RAW data reconstruction appears to work, but it is still too early to call done.
- Both features are only for short-term, longer term (maybe as part of the computing upgrade) we will need to implement parallel I/O:
 - Upgrades of the I/O framework, such as Scatter/Gather architecture for event reading, do not only help with I/O, but also:
 - Memory, (only one copy of TTree in memory, rather than one per Worker)
 - Storage, (better output file compression / optimization)
 - CPU, (avoid Worker uncompressing identical baskets (which contain several events))