

Speeding HEP Analysis with ROOT Bulk I/O

B Bockelman¹, Z Zhang² and O Shadura²

¹ Morgridge Institute for Research, Madison, WI 53715, USA

² Holland Computer Center, University Nebraska – Lincoln, Lincoln, NE 68588, USA

E-mail: bbockelman@morgridge.org

Abstract. Distinct HEP workflows have distinct I/O needs; while ROOT I/O excels at serializing complex C++ objects common to reconstruction, analysis workflows typically have simpler objects and can sustain higher event rates. To meet these workflows, we have developed a “bulk I/O” interface, allowing multiple events’ data to be returned per library call. This reduces ROOT-related overheads and increases event rates – orders-of-magnitude improvements are shown in microbenchmarks.

Unfortunately, this bulk interface is difficult to use as it requires users to identify when it is applicable and they still “think” in terms of events, not arrays of data. We have integrated the bulk I/O interface into the new RDataFrame analysis framework inside ROOT. As RDataFrame’s interface can provide improved type information, the framework itself can determine what data is readable via the bulk IO and automatically switch between interfaces. We demonstrate how this can improve event rates when reading analysis data formats, such as CMS’s NanoAOD.

1. Introduction

LHC experiment event data models are very complex and slow to read. The problem is that experiments do not care because input I/O time is minimal compared to the reconstruction process. Another critical factor is that experiments care about volume because they have lots of expensive disks.

For analysis case, the situation is different, since the data model is often more straightforward. The same case is about data volume used during the analysis phase, and there will be generated smaller a data volume and often used from SSD (NVMe). It causes minimal CPU costs and allows to iterate over events many times quickly.

ROOT IO is an incredibly flexible format. It can easily store the complex objects that correspond to the experiment’s data. In the same time, ROOT has high overheads for the serialization of simple objects.

2. Bulk IO

The typical mechanism for iterating through data in a `TTree` is a handwritten for-loop. ROOT uses a API shown in Listing 2 to read objects from a branch (`TTree` is a structure that contains one or multiple `TBranches`). This function runs in two steps. First, it searches the underlying storage medium for the basket where the event is located and then read the basket into a memory buffer. The `TBasket` is the data structure that represents the in-memory buffer. ROOT decompresses the buffer and put the uncompressed buffer in so-called “kernel” space. In the

second step, once the basket appears in memory, `GetEntry` deserializes the requested event from the kernel-space buffer and copy it to user-space buffer.

```
Int_t TBranch::GetEntry(Long64_t entry)
```

Listing 1: GetEntry in TBranch

When user application is computationally expensive, the cost of library calls, frequently deserializing objects and copying data between memory buffers are amortized to effectively nothing. To overcome such overheads, we introduce a new interface for ROOT to copy all events in an on-disk `TBasket` directly to a user-provided memory buffer. For the simplest cases - primitives and C-style arrays of primitives, the serialization can be done without a separate buffer or “fixing up” pointer contents. The user can request the serialized data or deserialized data to be delivered to the user buffer. By requesting the serialized data directly and deserializing directly in the event loop, the user can avoid an expensive scan from main memory.

Pragmatically, the user will not implement code for deserializing data themselves: rather, we have provided a header-only C++ facade around the data, allowing the user to work with a proxy object. This allows the compiler to inline the deserialization code in the correct place.

3. Implementation

Bulk IO interface is a set of APIs that are built in the existing ROOT IO framework. The user can choose between regular APIs and Bulk IO APIs. We implement Bulk IO in three common use cases: `TBranch`, `TTreeReader` and `RDataFrame`. We discuss about our interface design and integration in this section.

3.1. Bulk IO in TBranch

Listing 3.1 shows Bulk IO API in `TBranch` in which two input arguments need to be parsed into the function: `entry` and `user_buf`. The `entry` defines an event index number indicating which event the function is going to read. The `user_buf` parses an user-space `TBuffer` structure as a reference into the function. In the end of the function call, the `user_buf` should contain the whole basket of data that contains the inquiry event.

```
Int_t TBranch::GetBulkEntries(Long64_t entry, TBuffer &user_buf)
```

Listing 2: Bulk API in TBranch

It is worthwhile to mention that `GetBulkEntries` deserializes events on-the-fly when the data read into the `user_buf`. Thus no further manipulation is required for user applications. An user can later on access an event in the basket using the code snippet shown in Listing 2 where T is the object type and idx is the event index in the `user_buf`.

```
*reinterpret_cast<T*>(user_buf.GetCurrent())[idx]
```

Listing 3: Loop over user_buf

3.2. Bulk IO in TTreeReader

`TTreeReader` is an interface for an user to access simple object (primitives, arrays, etc.) in a ROOT file. `TTreeReaderValue` is the interface to access primitives and `TTreeReaderArray` is the interface to access arrays (each event is either an fixed-size or varialbe-size array). Listing 3.2 shows a code sample that uses `TTreeReaderValue` to read events (floats) from a file. `TTreeReader` internally relies on `GetEntry` to access events.

We design a Bulk API - `GetEntriesSerialized` in `TBranch` shown in Listing 4. We introduce a new interface - `TTreeReaderFast` that uses `GetEntriesSerialized` to function as `TTreeReader`. Unlike `GetBulkEntries`, `GetEntriesSerialized` does not deserialize events while reading basket

into `user_buf`. Instead, it waits until the user calls `*myF`. Dereference operator invokes the appropriate deserialization code.

```
TTreeReader myReader("T", hfile);
TTreeReaderValue<float> myF(myReader, "myFloat");
Long64_t idx = 0;
Float_t sum = 1;
while (myReader.Next()) {
    sum += *myF;
}
```

Listing 4: Access to Events using TTreeReader

```
Int_t TBranch::GetEntriesSerialized(Long64_t entry, TBuffer &user_buf)
```

Listing 5: Bulk API in TTreeReaderFast

3.3. Bulk IO in RDataFrame

During our work, Bulk IO is also integrated into RDataFrame [2] which is a python Pandas [4] like data analysis framework for ROOT users. RDataFrame provides a proxy interface - RDataSource (RDS). It allows RDF to read arbitrary data formats such as TTree, CSV, etc..

```
Int_t TBranch::GetEntriesSerialized(Long64_t entry,
                                    TBuffer &user_buf, TBuffer *count_buf)
```

Listing 6: Bulk API in RDataFrame

We define a new Bulk API `GetEntriesSerialized` shown in Listing 3.3. The only difference from Listing 4 is that there is one more argument `count_buf`. Listing 4 actually calls Listing 3.3 and set the `count_buf` as `nullptr`. `count_buf` is used to store array length information when events in RDataFrame are arrays. Variable-size arrays need such information to deserialize the buffer into multiple individual arrays.

4. Evaluation

4.1. Experiments

All tests are conducted on a desktop Intel i5 4-Core @ 3.2GHz. A TTree with 100 million float values is read with different APIs. We tested three different use cases: GetBulkEntries, TTreeReaderFast and RDataSource.

4.2. Results

Figure 1 shows the time spent on iterating all events in the `TTree` with `GetEntry` and `GetBulkEntries`. Figure 2 shows the read time between `TTreeReader` and `TTreeReaderFast`. As shown in the figures, Bulk IO spends 10+ times less than `GetEntry` and `TTreeReader`. Bulk IO in both use cases spends similar time on reading events. `TTreeReader` interface spends more than 3 times reading events than `GetEntry` due to the overheads of TTreeReader itself (TTreeReader internally calls `GetEntry`).

Figure 3 shows the results of Bulk IO in RDataFrame. In the figure, the standard RDF shows the performance by using regular RDataFrame function calls. Bulk RDF and Bulk RDS show the result of Bulk APIs. The difference is that, Bulk RDS test detaches RDataSource from RDataFrame stack and run the test directly through RDS function calls. As shown in Figure 3, Bulk RDS outperforms standard RDF by more than 2 times. In addition, RDataFrame has extra overheads compared to RDataSource (RDataFrame internally relies RDataSource), therefore Bulk RDF runs slower than Bulk RDS, but still outperforms standard RDF.

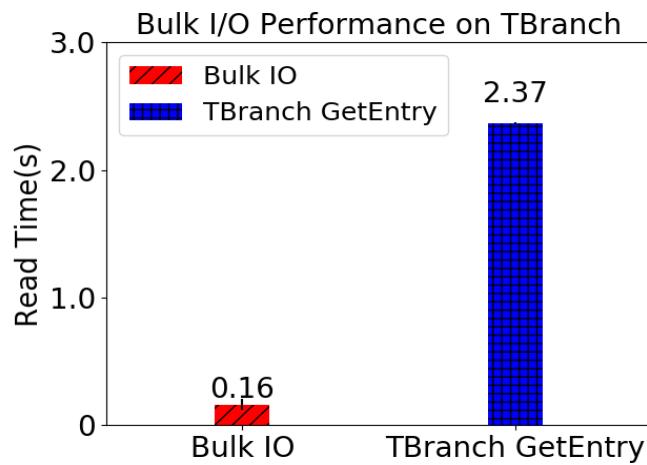


Figure 1: Performance between *GetEntry* and *GetBulkEntries*.

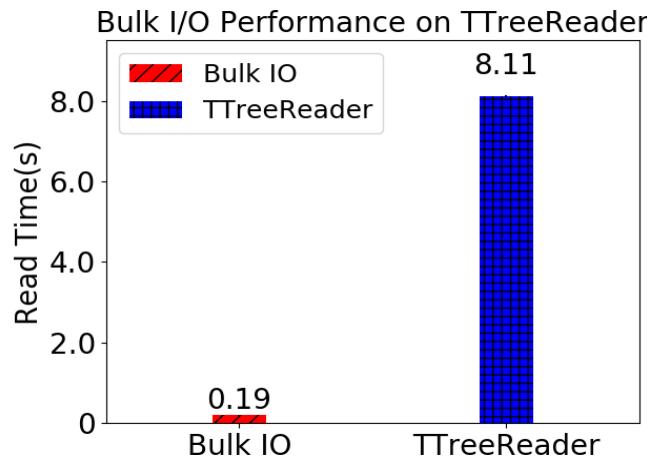


Figure 2: Performance between *TTreeReader* and *TTreeReaderFast*.

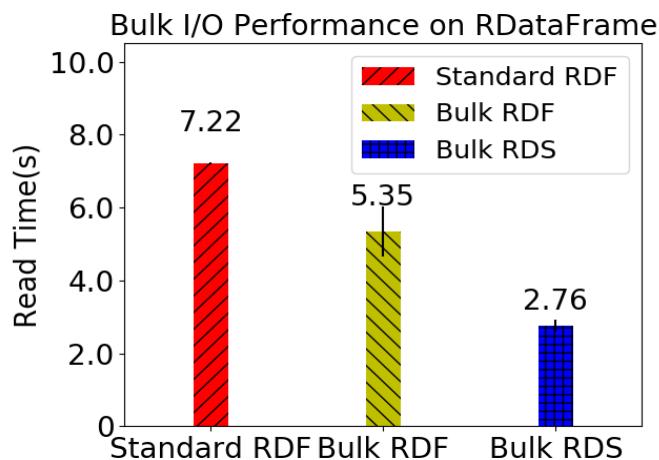


Figure 3: Performance improvements on RDataFrame with Bulk IO.

Acknowledgments

This work was supported by the National Science Foundation under Grant ACI-1450323. This research was done using resources provided by the Holland Computing Center of the University of Nebraska.

References

- [1] Brun R and Rademakers F “ROOT - An object oriented data analysis framework”, *Nucl. Instr. Meth. Phys. Res.* **389** (1997) 81-86
- [2] Guiraud E, Naumann A and Piparo D “RDataFrame: functional chains for ROOT data analyses”, (2017) doi: 10.5281/zenodo.260230. url: <https://doi.org/10.5281/zenodo.260230>.
- [3] Bockelman B, Zhang z and Pivarski J “Optimizing ROOT IO For Analysis”, *J. Phys.: Conf. Ser.*, **1085** (2018) 032012
- [4] Mckinney W “pandas: a Foundational Python Library for Data Analysis and Statistics”, *PyHPC 2011 : Python for High Performance and Scientific Computing*, (2011)