CMS Performance Measurements

Dr Christopher Jones

RNTuple Workshop

02 12 2024

RNTuple Status in CMS

- We have prototype code which can read/write RNTuple files

- Can write any of our standard storage formats
  - RECO [*~1MB/ev*] , AOD [*~400kB/ev*], and MiniAOD [*~150kB/ev*]

- Only possible because of addition of streamer storage mechanism
  - allows storage of any class that TTree could handle
    - storage of bare pointers
    - storage of polymorphic types
    - storage of classes which directly/indirectly store themselves

MiniAOD Storage Comparison

- Storage Type
  - **select streamer**: apply streamer storage to data products if it reduces storage size
  - **select unsplit**: do not split *multi-byte basic type* fields if it reduces storage size
- Optimization
  - **zip cluster**: `ApproxZippedClusterSize` set to 25M down from 50M
  - **no buf write**: turn off `BufferedWrite`

Storage Type	Optimization	File Size Ratio	Max Allocation	Allocation Diff	Max RSS	RSS Difference
TTree		1.000	2,329,950,216	0	2,640,244,736	0
RNTuple default		**0.981**	2,975,349,944	**645,399,728**	3,188,834,959	**548,590,223**
RNTuple select streamer		**0.937**	3,001,023,440	**671,073,224**	3,285,766,144	**645,521,408**
	zip cluster	**0.933**	2,486,353,160	**156,402,944**	2,807,153,295	**166,908,559**
	no buf write	**0.951**	1,846,704,808	**-483,245,408**	2,240,806,912	**-399,437,824**
RNTuple select unsplit		**0.924**	3,036,038,160	**706,087,944**	3,308,414,566	**668,169,830**
	zip cluster	**0.918**	2,503,737,024	**173,786,808**	2,750,152,704	**109,907,968**
	no buf write	**0.939**	1,853,864,576	**-476,085,640**	2,143,961,108	**-496,283,628**

Fermilab

Thread Scaling for Writing

- No scaling without BufferedWrite

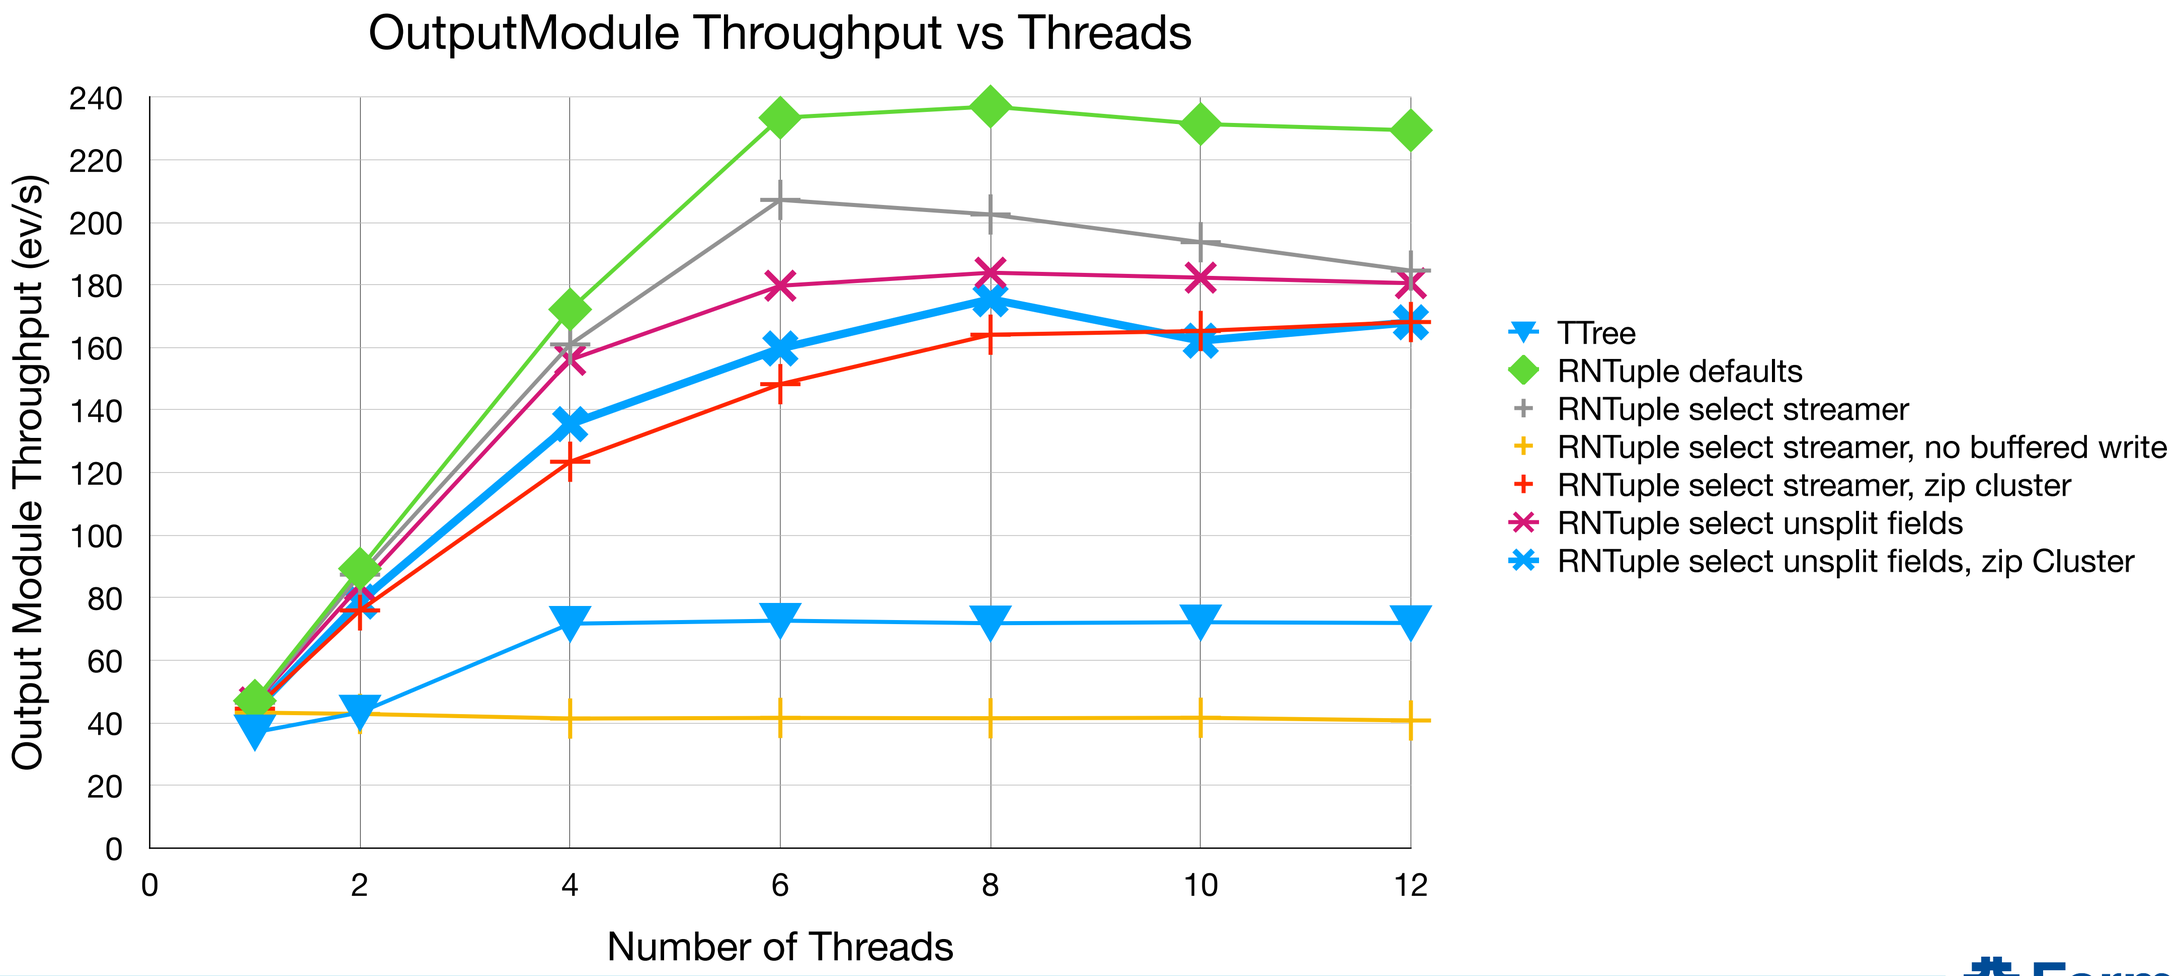

Fermilab

AOD Storage Comparison

- Additional Storage Type
  - **all Streamer**: Force all top level fields to be streamer type

AOD Storage Comparison

Storage Type	Optimization	File Size Ratio	Max Allocation	Allocation Difference	Max RSS	RSS Difference
TTree		1.000	1,365,851,992	0	2,213,576,704	0
RNTuple select streamer		**0.880**	2,314,973,832	**949,121,840**	2,732,180,111	**518,603,407**
RNTuple select unsplit		**0.859**	2,358,814,824	**992,962,832**	2,702,578,811	**489,002,107**
	zip cluster	**0.856**	2,018,192,176	**652,340,184**	2,443,108,680	**229,531,976**
	no buf write	**0.883**	1,513,042,232	**147,190,240**	1,848,754,831	**-364,821,873**
RNTuple all Streamer		**0.984**	1,766,328,840	**400,476,848**	2,199,472,046	**-14,104,658**
	no buf write	**1.001**	1,072,140,704	**-293,711,288**	1,629,864,591	**-583,712,113**

Fermilab

AOD Throughput

- RNTuple with buffered writing scales much better than TTree

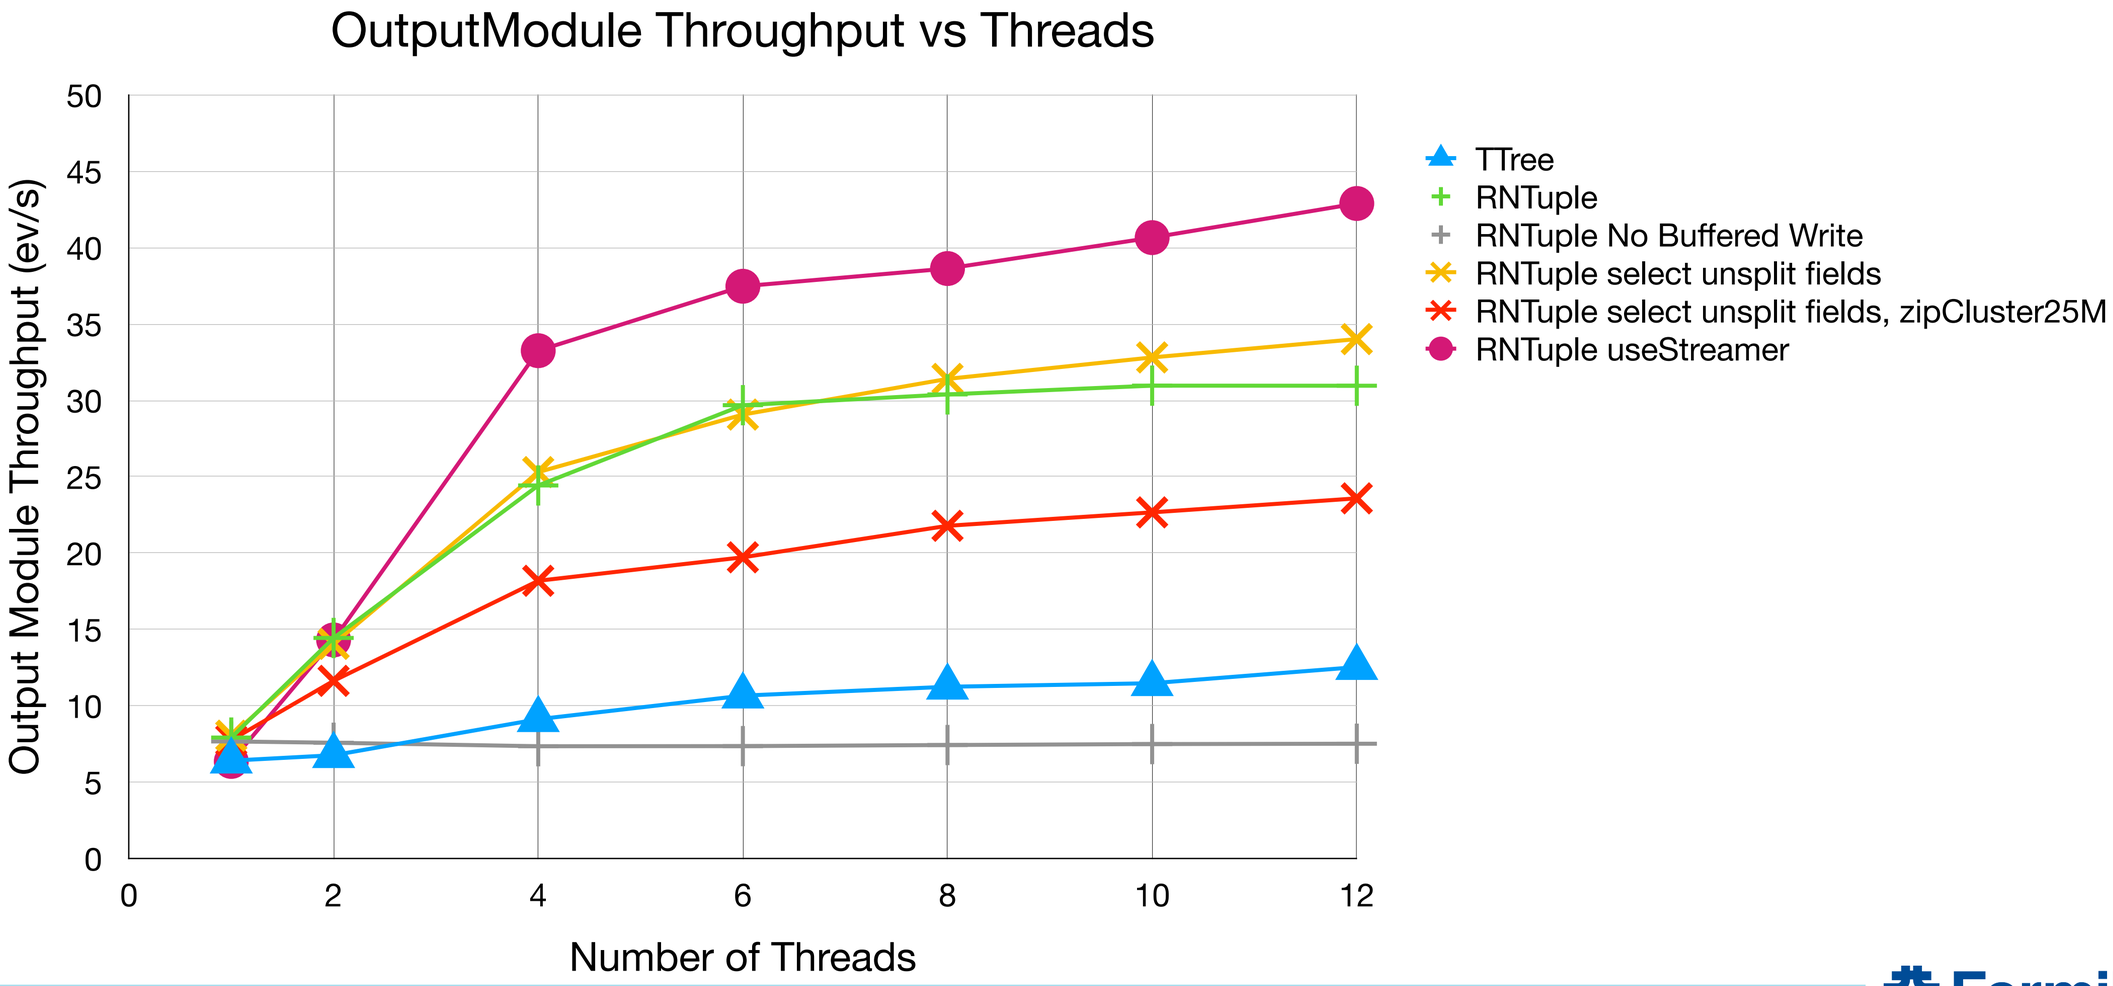

Fermilab

Conclusion

- Can store all CMS data products in RNTuple
- Thread scaling is better than TTree
  - assuming can use Buffered Writing
- Extra memory required is concerning