An HDF5-WRF module

-A performance report

MuQun Yang, Robert E. McGrath, Mike Folk

National Center for Supercomputing Applications
University of Illinois, Urbana-Champaign

ymuqun@ncsa.uiuc.edu

URL: **http://hdf.ncsa.uiuc.edu/apps/WRF-ROMS/**

The uniqueness about the study

- The HDF5 is not used to store satellite data, but the output from a sophisticated numerical weather model
- Explore the performance of parallel HDF5 in parallel computing environments
- Investigate the performance of the compression feature inside HDF5 when applying to the numerical model
-

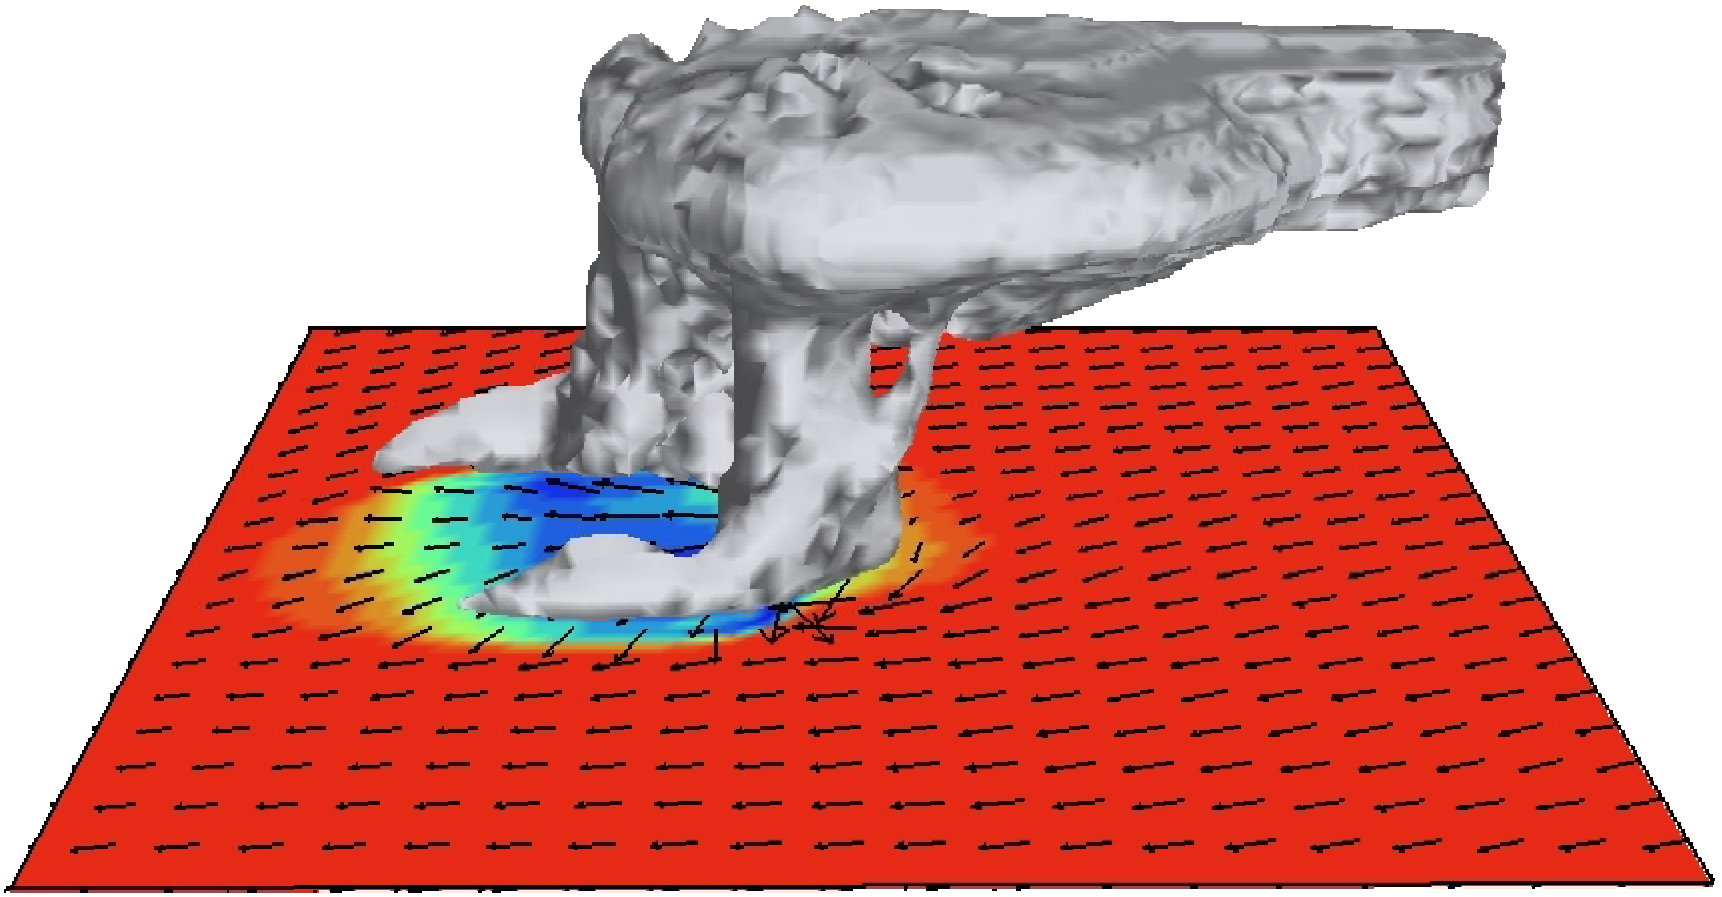

From WRF tutorial

Schematic HDF5 File Structure of Sequential HDF5-WRF Output

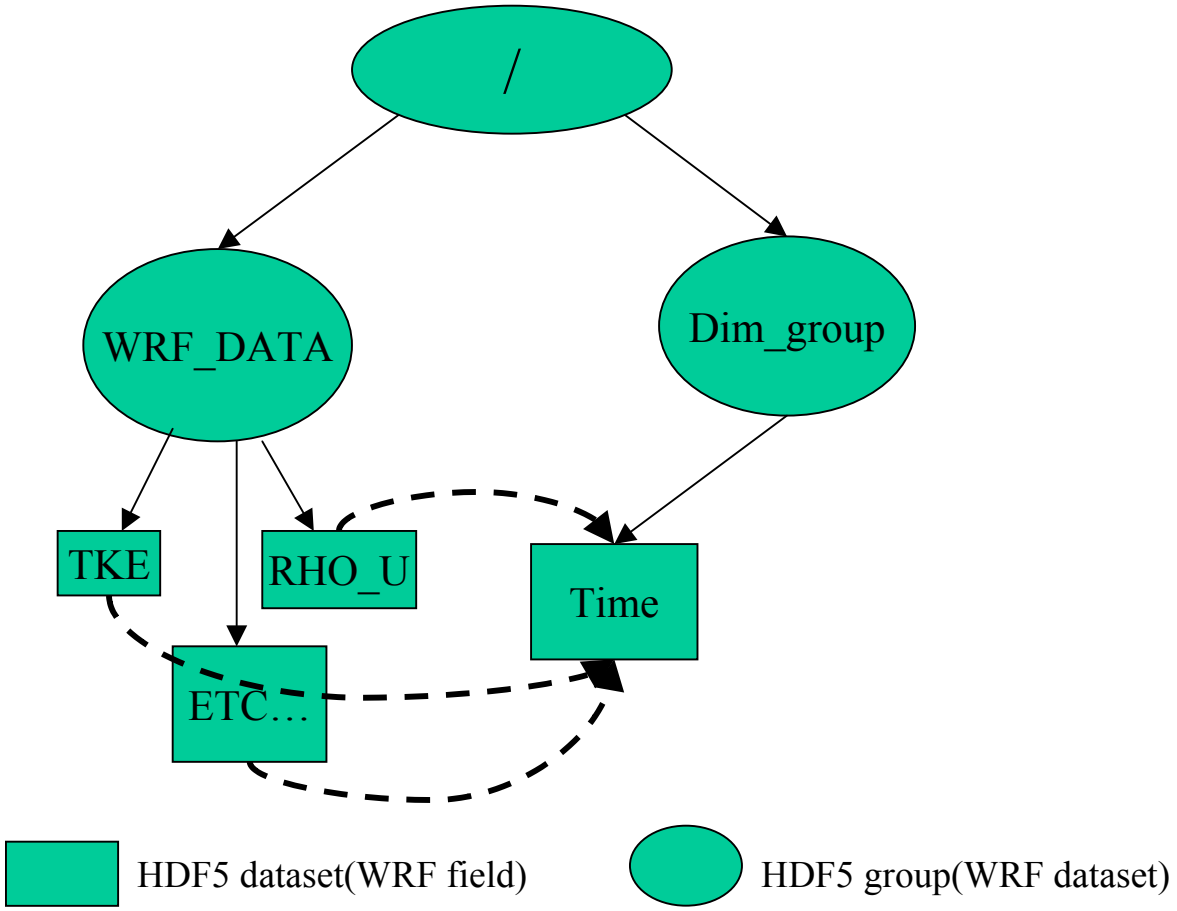

Solid line: HDF5 datasets or sub-groups (the arrow points to) that are members of the HDF5 parent group.

Dash line: The association of one HDF5 object to another HDF5 object through dimensional scale table.

Schematic HDF5 File Structure of Parallel HDF5-WRF Output

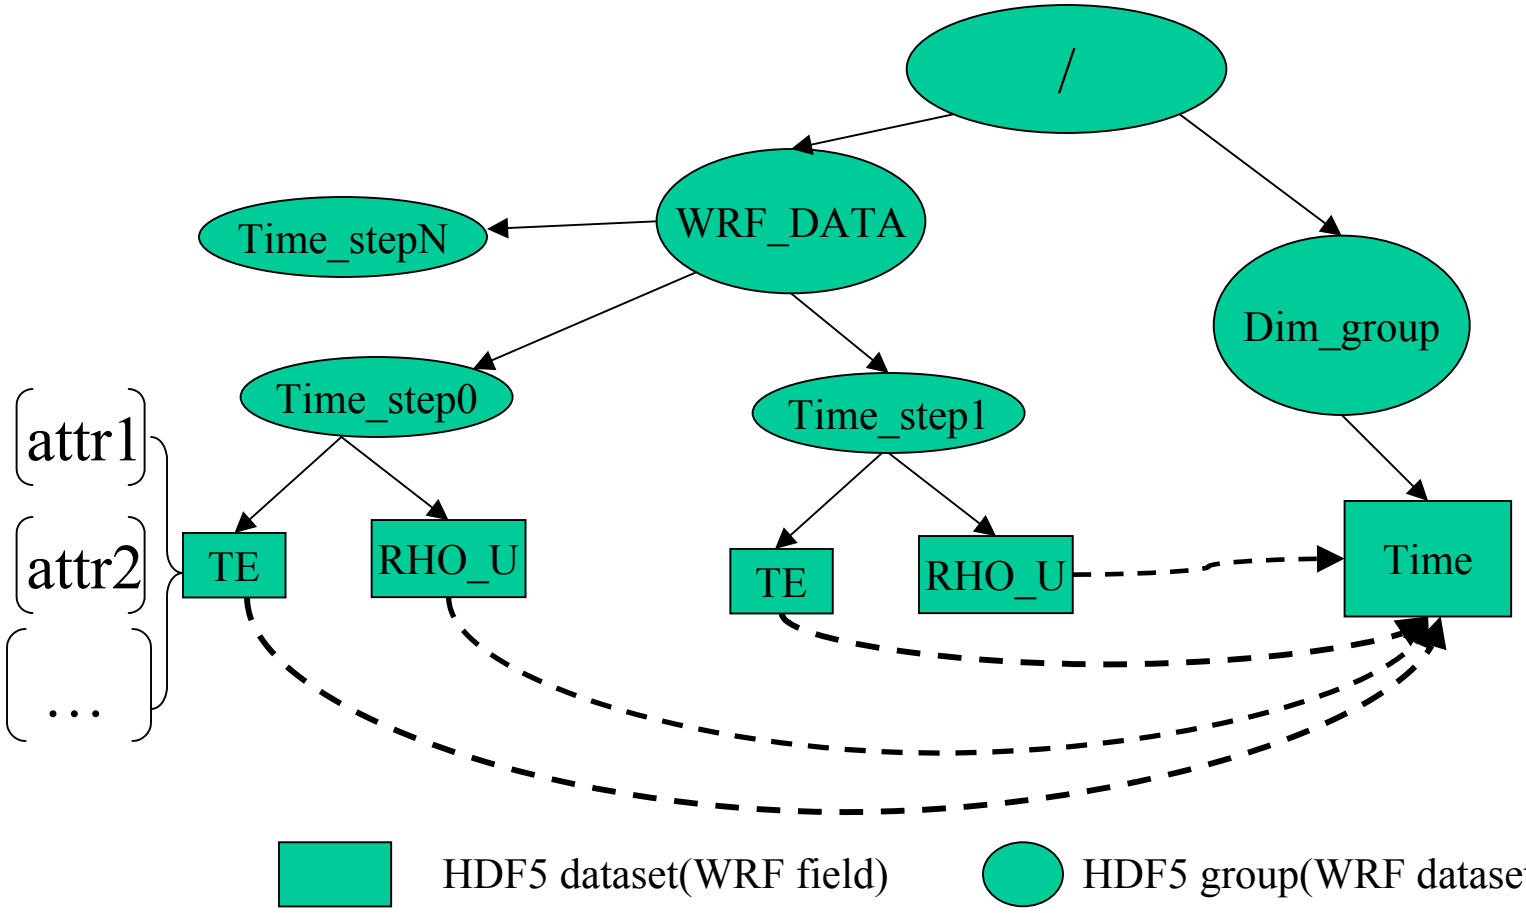

Solid line: HDF5 datasets or sub-groups (the arrow points to) that are members of the HDF5 parent group.

Dash line: the association of one HDF5 object to another HDF5 object through dimensional table.

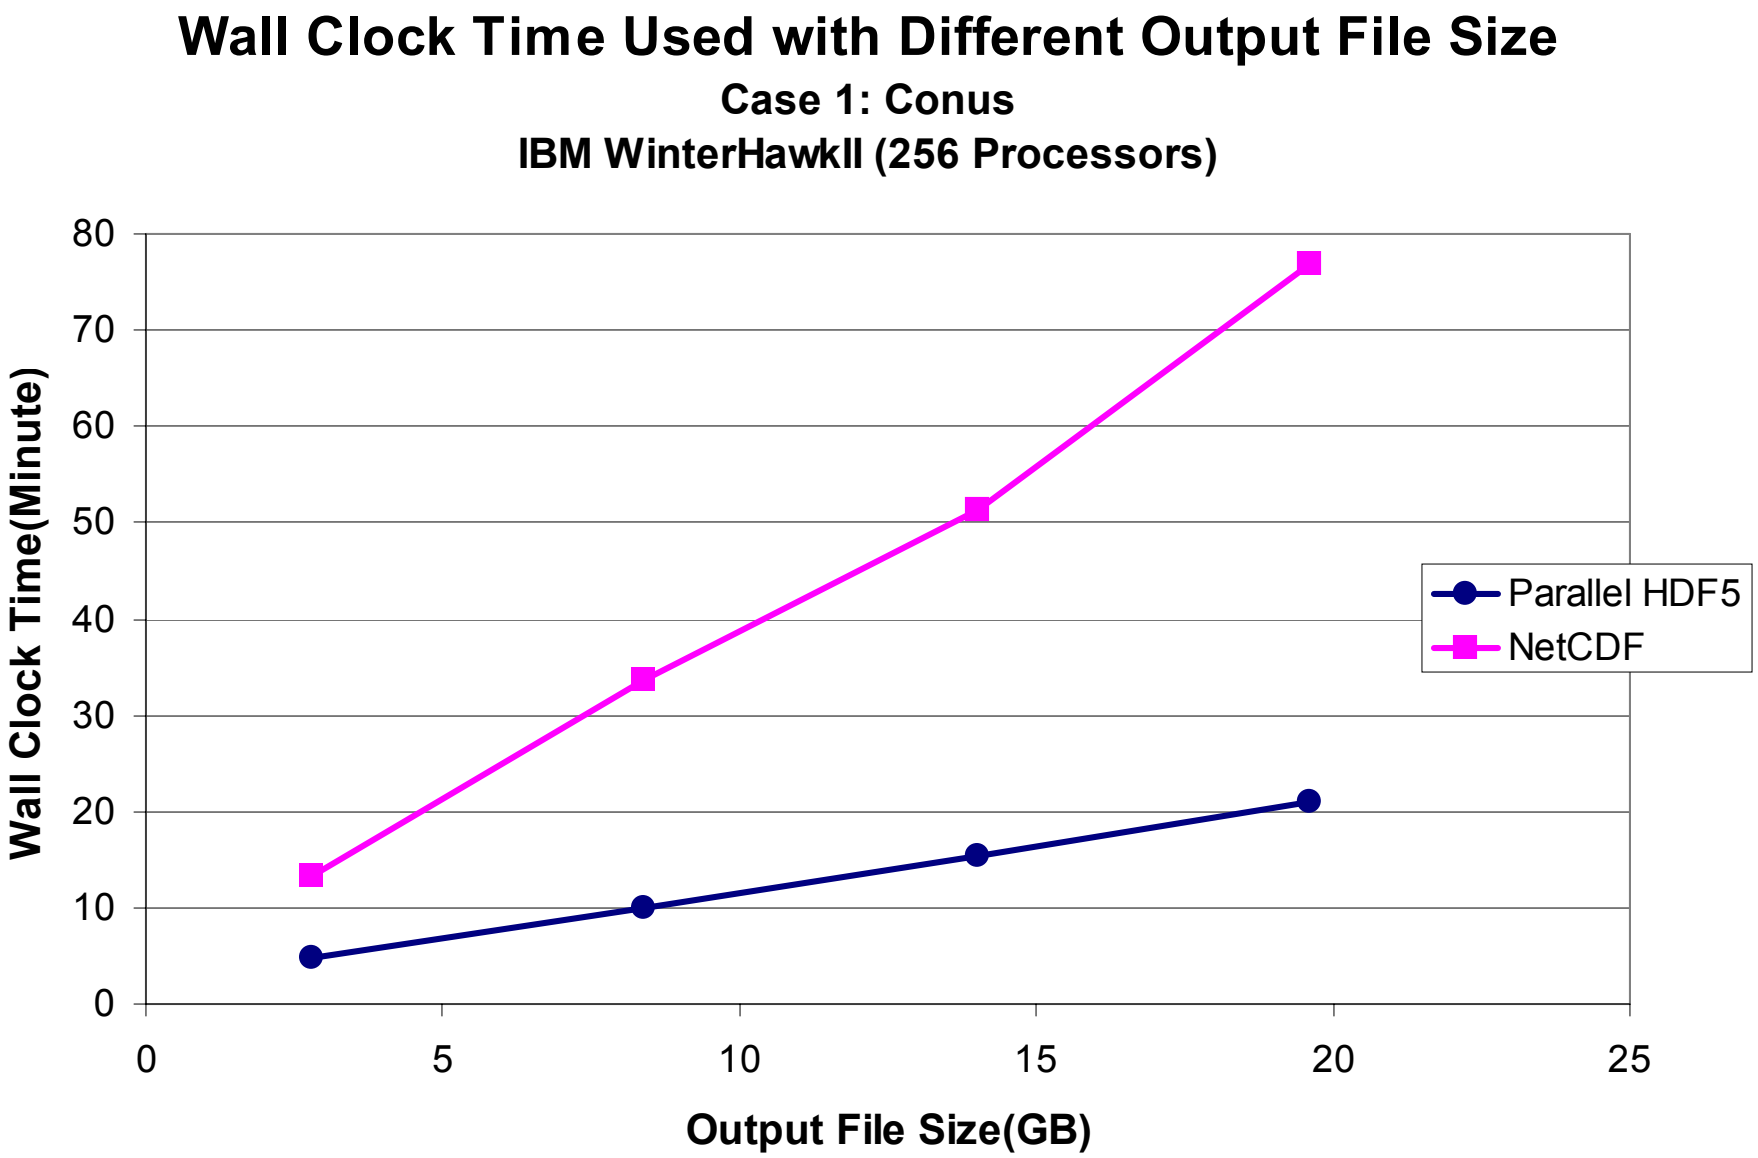
Wall Clock Time Used with Different Output File Size
Case 1: Conus
IBM WinterHawkII (256 Processors)
Wall Clock Time(Minute)
80
70
60
50
40
30
20
10
0
0
5
10
15
20
25
Output File Size(GB)
Parallel HDF5
NetCDF

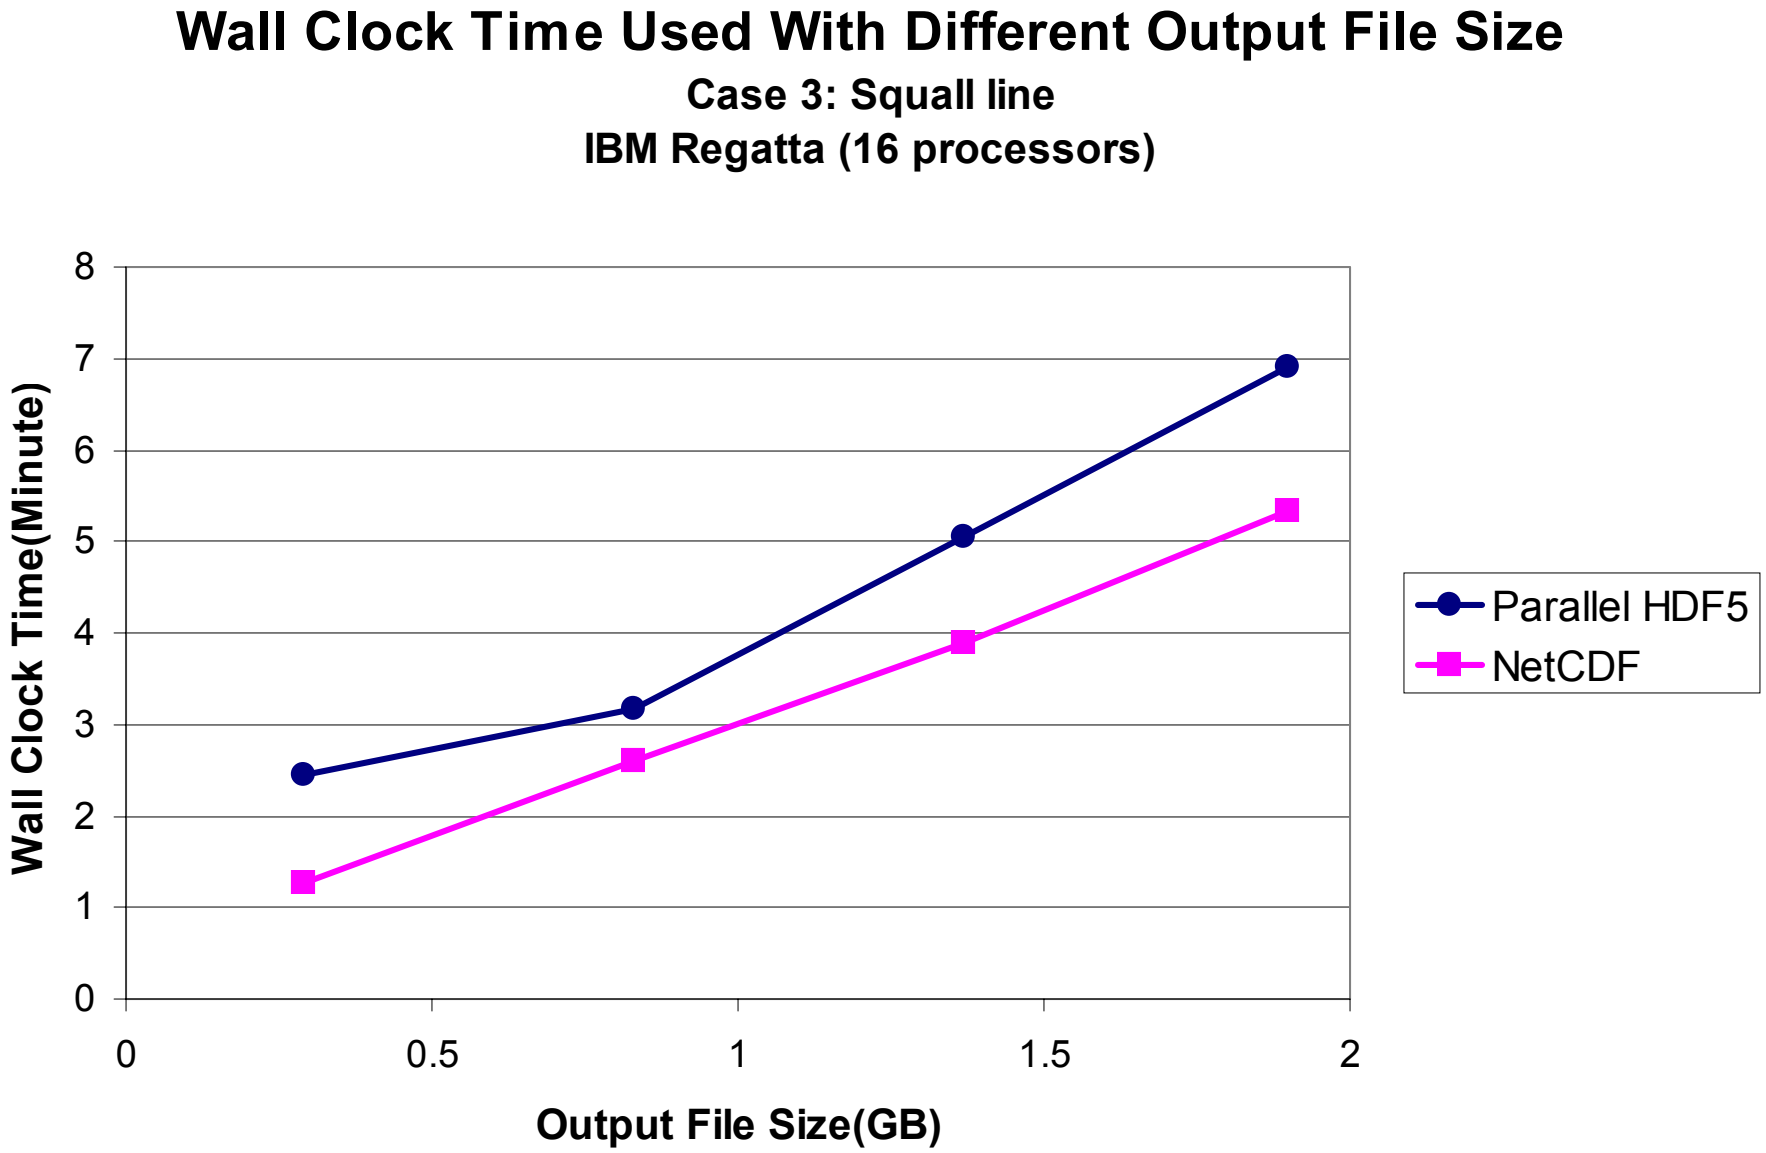
Wall Clock Time Used With Different Output File Size
Case 3: Squall line
IBM Regatta (16 processors)
Wall Clock Time(Minute)
0
1
2
3
4
5
6
7
8
0
0.5
1
1.5
2
Output File Size(GB)
Parallel HDF5
NetCDF

Model Output File Size With Different Compressions

Case 3: Squall Line
IBM Regatta(16 processors)

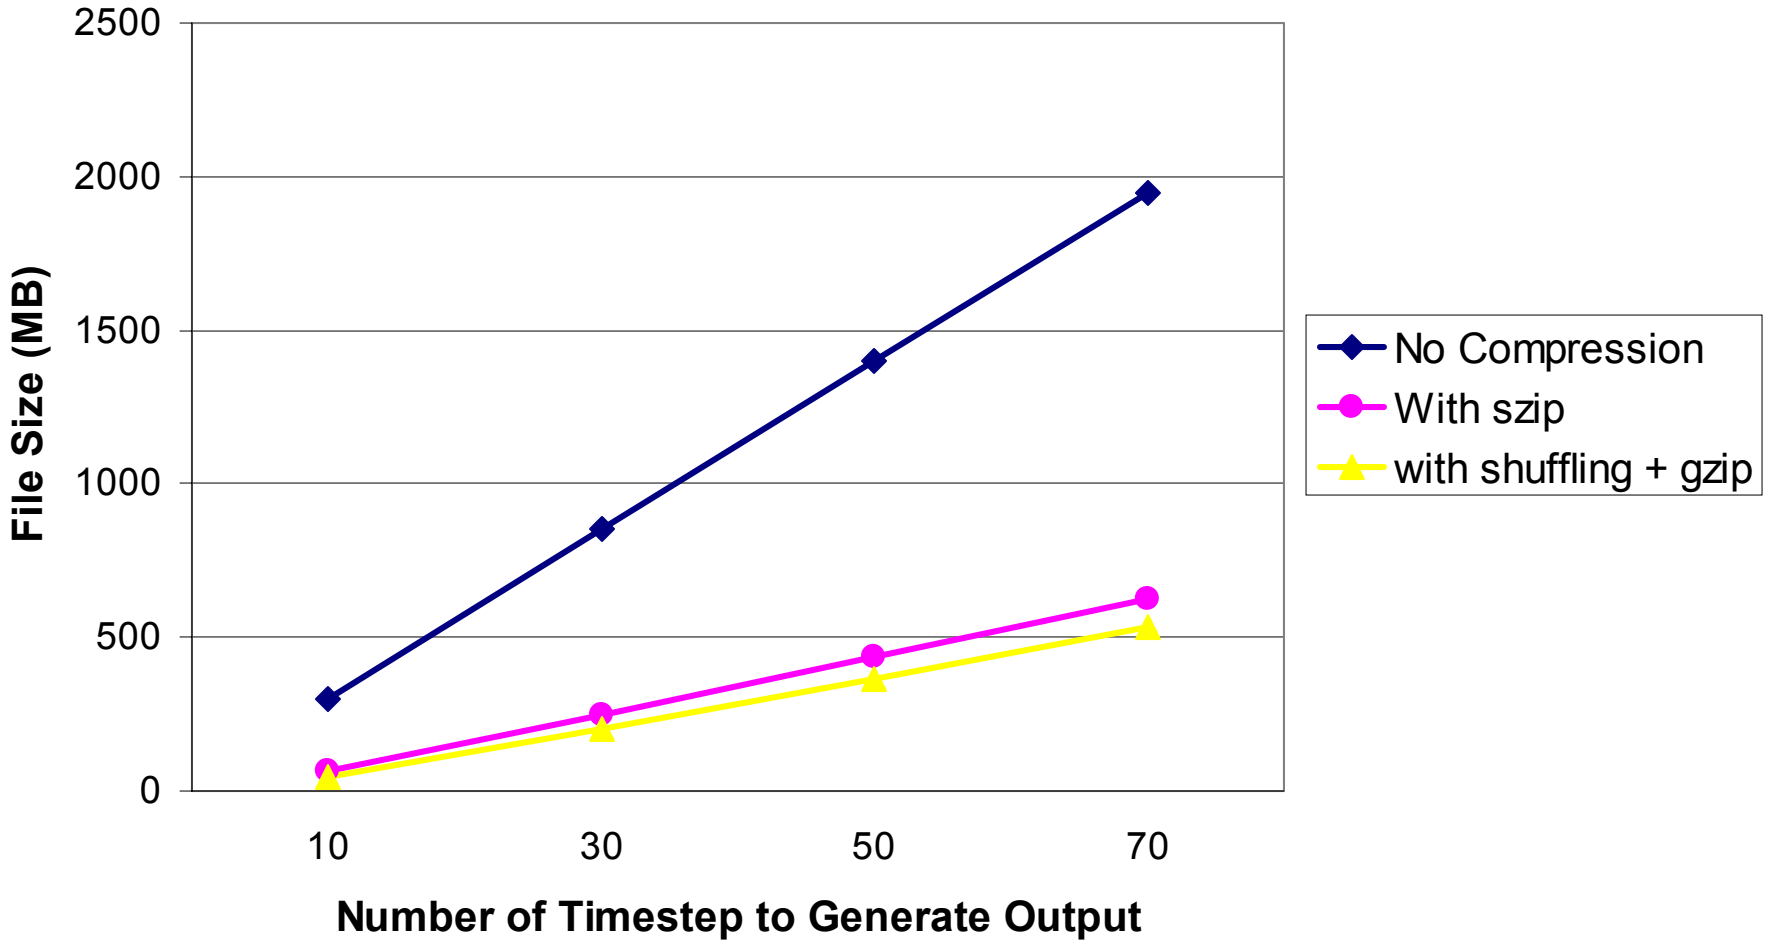

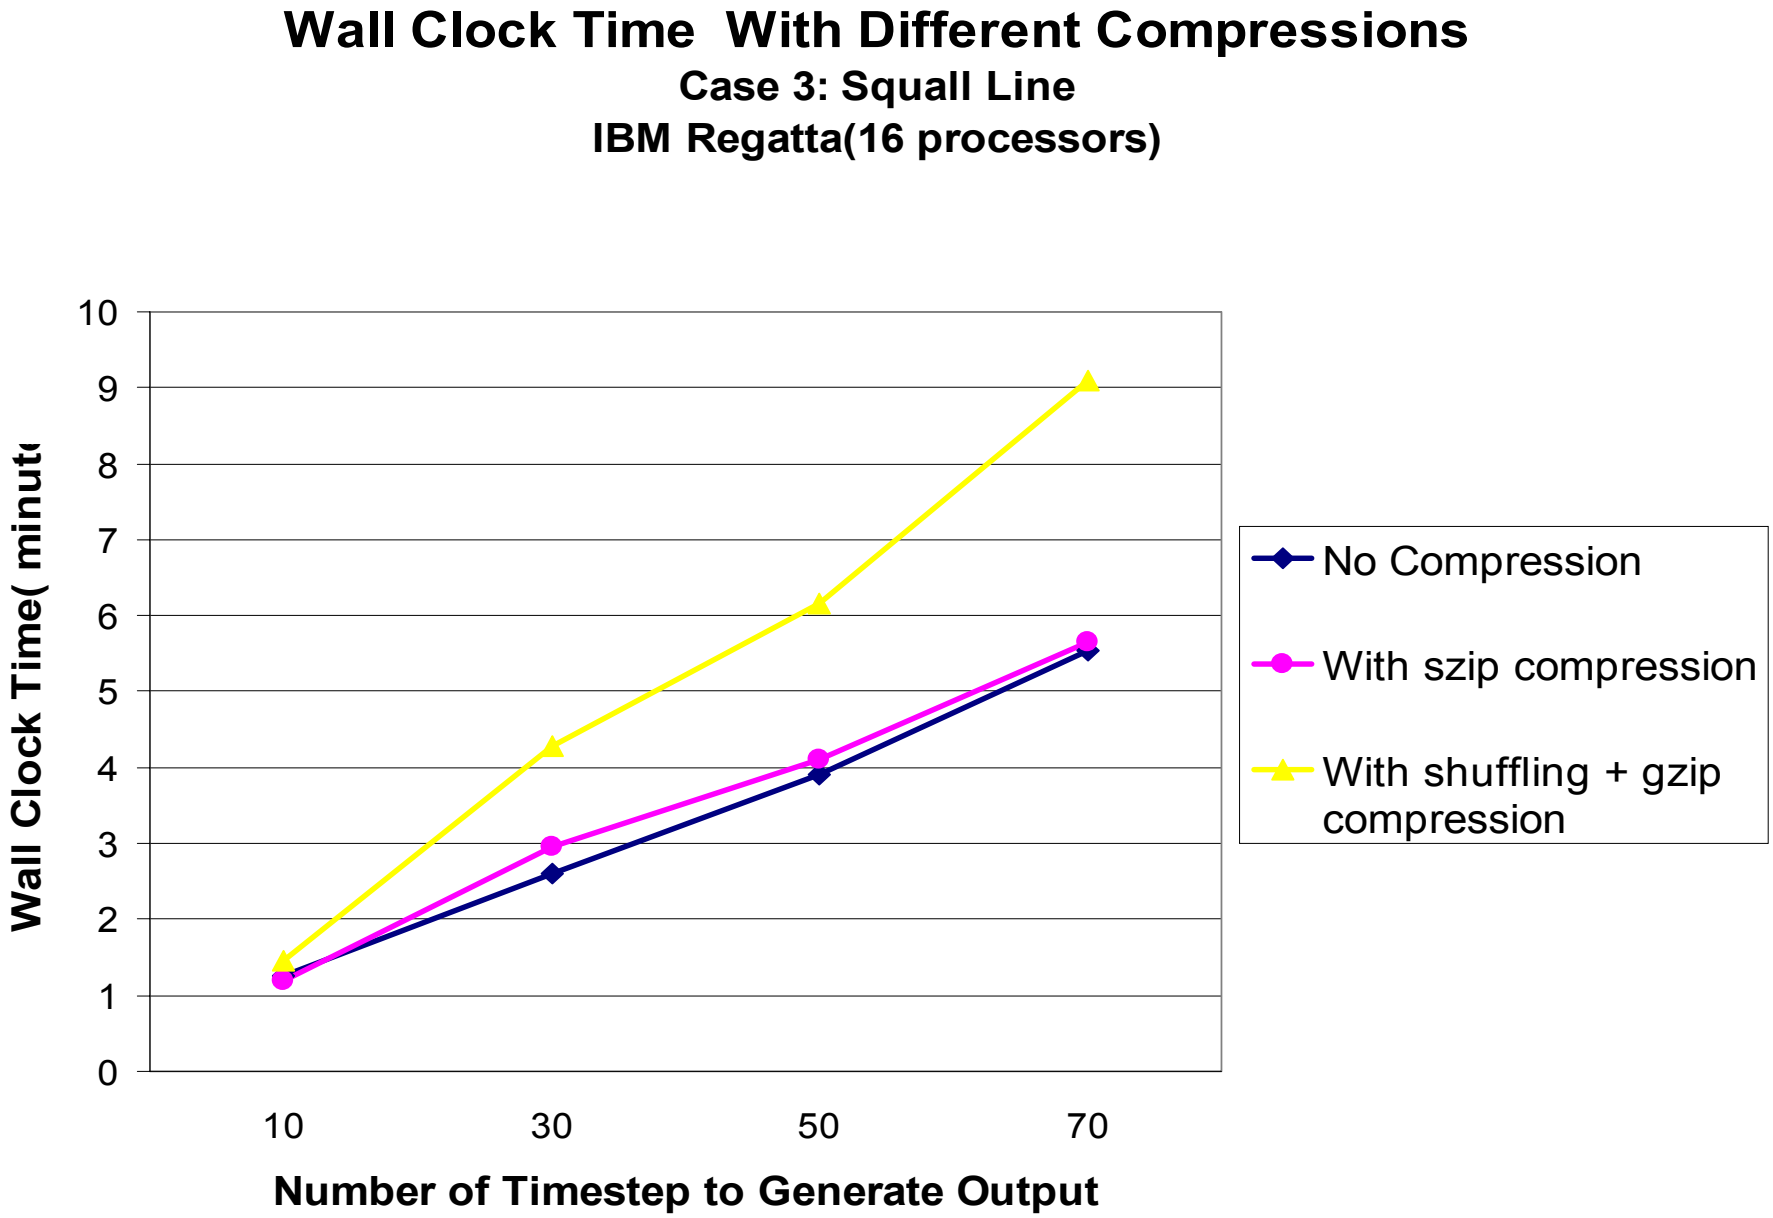
Wall Clock Time With Different Compressions
Case 3: Squall Line
IBM Regatta(16 processors)
Wall Clock Time(minut
10
9
8
7
6
5
4
3
2
1
0
10
30
50
70
Number of Timestep to Generate Output
No Compression
With szip compression
With shuffling + gzip compression

Summary

- Parallel IO is not trivial
- Effective chunking with MPI-IO inside HDF5 library is the key to improve parallel HDF5 performance
- Szip compression can improve performance for WRF application
- Shuffling algorithm with gzip compression can further improve compression ratio