Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Exploring the SZ lossy compressor use for the XIOS I/O server

Xavier Yepes-Arbós

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

The research leading to these results has received funding from the EU H2020 Framework Programme under grant agreement no. 823988
This material reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains

25/2/2021

12th JLESC Workshop

Introduction

- Exascale supercomputers will allow Earth System Models (ESMs) to make simulations at an unprecedented level of horizontal resolution.
- But this has implications:
  - A huge amount of data will be generated that must be efficiently written into the storage system.
  - A high cost of storage systems due to the huge data size.

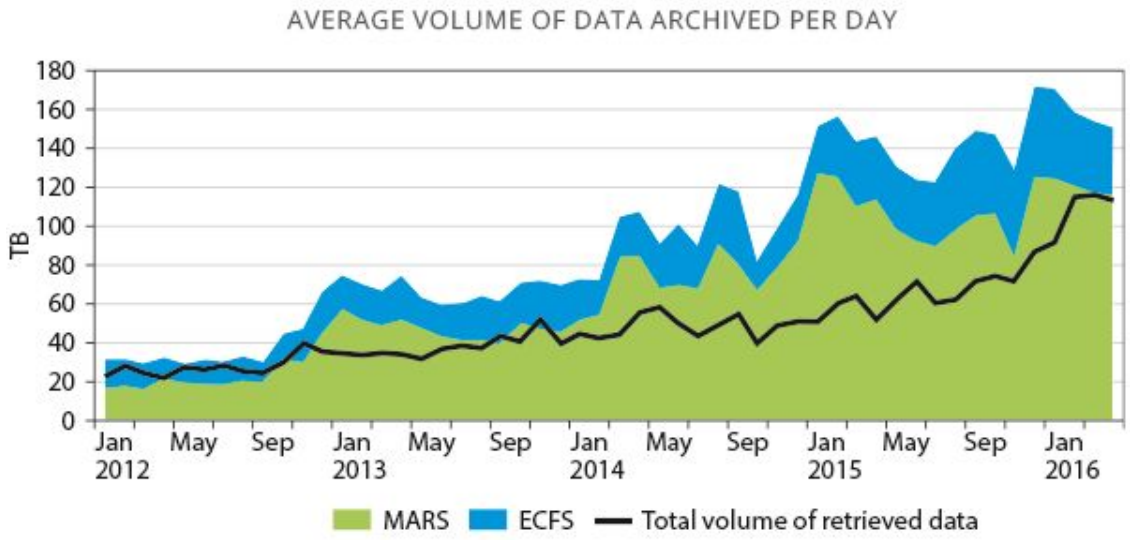

The XIOS I/O server

- The I/O issue is typically addressed by adopting scalable parallel I/O solutions.

- In the climate community, a widely I/O tool used is XIOS.

- The XML Input/Output Server (XIOS) is an asynchronous MPI parallel I/O server developed by the Institut Pierre-Simon Laplace (IPSL).

- XIOS has the following features needed for climate modelling:
  - Output files are in netCDF format.
  - Written data is CMIP-compliant (CMORized).
  - It is able to post-process data inline to generate diagnostics.

XIOS: Some technical features

- From a computational point of view, XIOS is thought to address:
  - The inefficient legacy read/write process.
  - The unmanageable size of “raw” data.

- By implementing:
  - Scalable parallel I/O.
  - Inline post-processing.

- In addition, it offers **lossless data compression** using gzip through HDF5.

Test case: OpenIFS and XIOS integration

- OpenIFS is a global meteorological forecasting model developed and maintained by the European Centre for Medium-Range Weather Forecasts (ECMWF).
- In the past we integrated XIOS into OpenIFS to address the former inefficient sequential I/O scheme.
- Although the **overhead** of outputting data through XIOS is really **small** for current resolutions, in the **future** this may well become a **bottleneck** because of the exponential growth of the output volume.
- The default lossless compression filter of HDF5 does not fit our needs:
  - If compression ratio is high, it takes too much time.
  - If it takes a reasonable amount of time, compression ratio is not enough.

OpenIFS-XIOS integration scheme

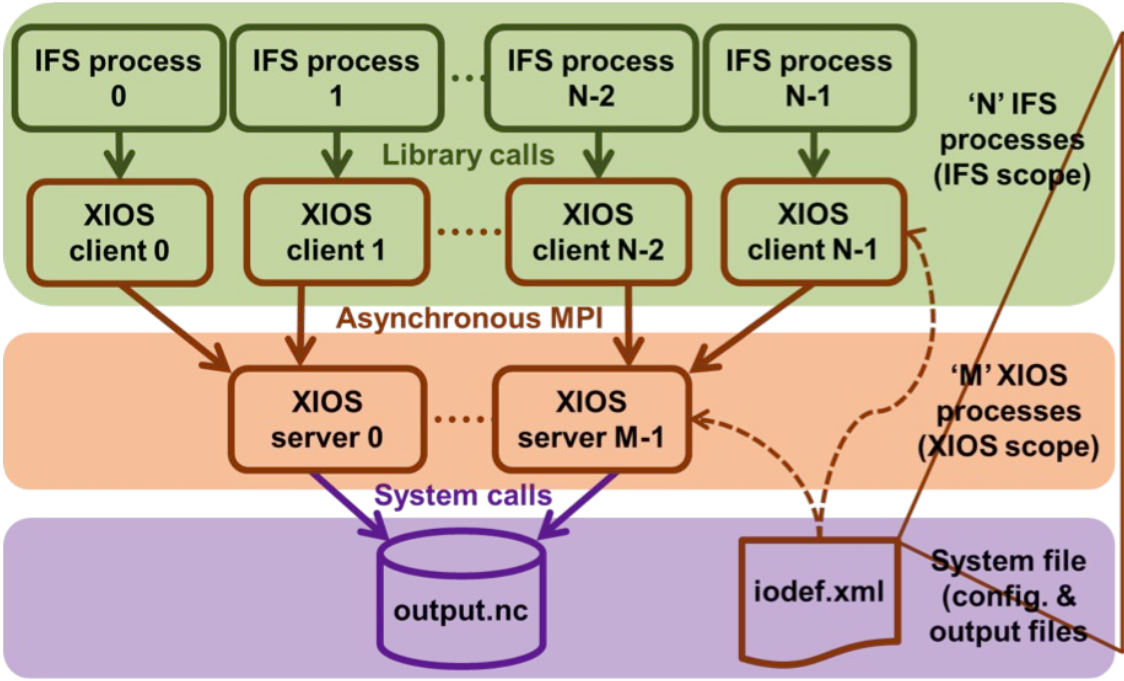


```
<grid id="grid_3d">
   <domain id="domain_2d"
   type="rectilinear" ni_glo=… />
   <axis id="axis_1d" n_glo=… />
</grid>

<field_definition>
   <field id="temp" grid_ref="grid_3d" />
   <field id="pressure" grid_ref="grid_3d" />
</field_definition>

<file_definition>
   <file name="output" freq_output="1mo">
      <field field_ref="temp" />
      <field field_ref="pressure" />
   </file>
</file_definition>
```

What about XIOS compression?

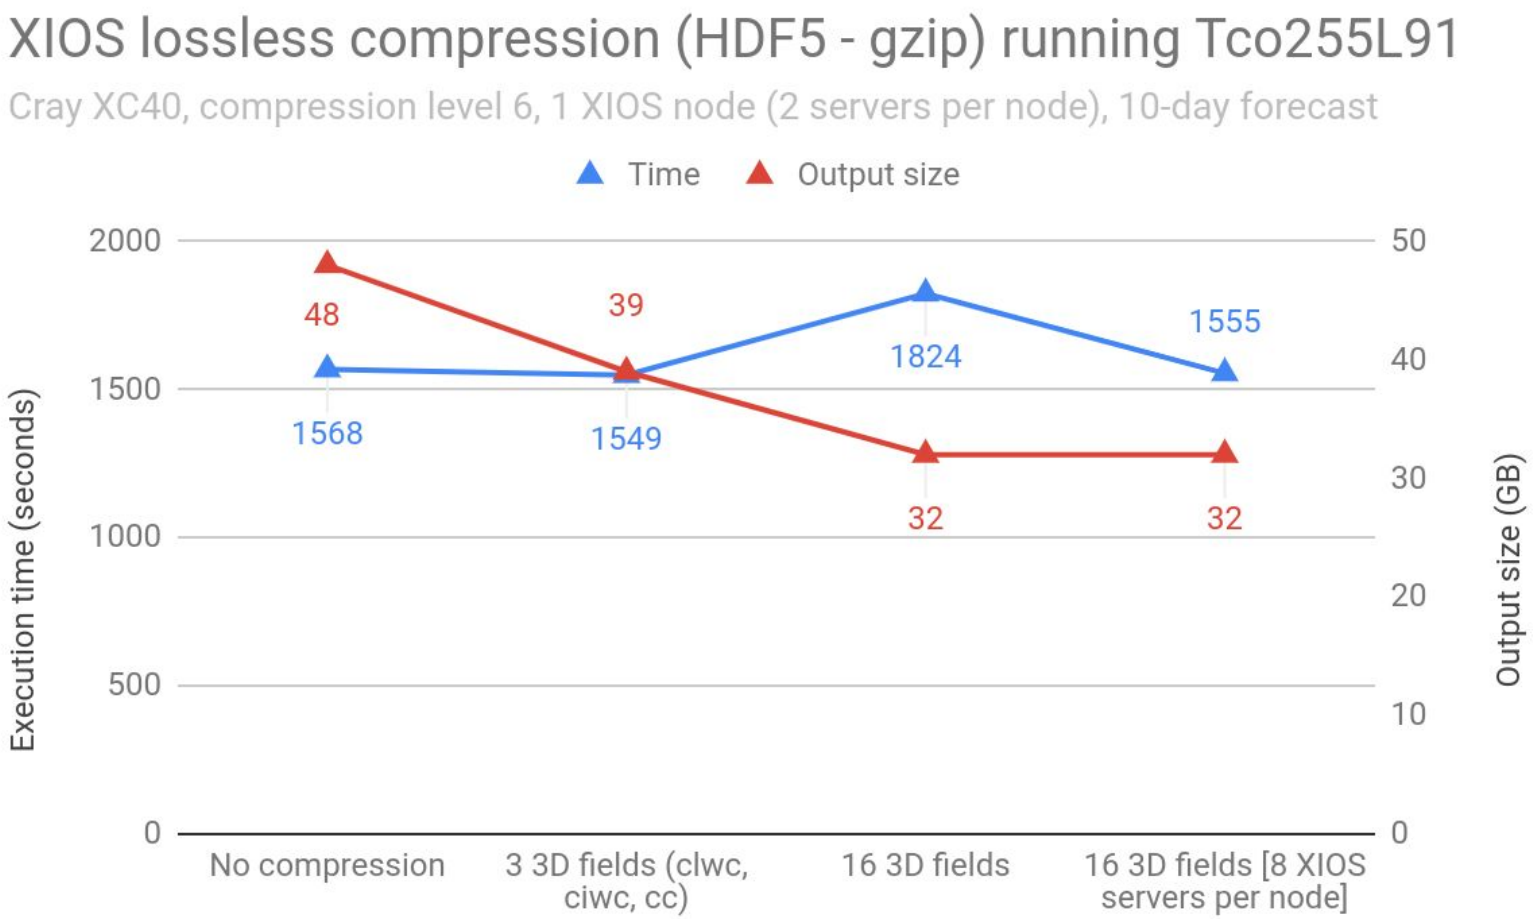

What about XIOS compression?

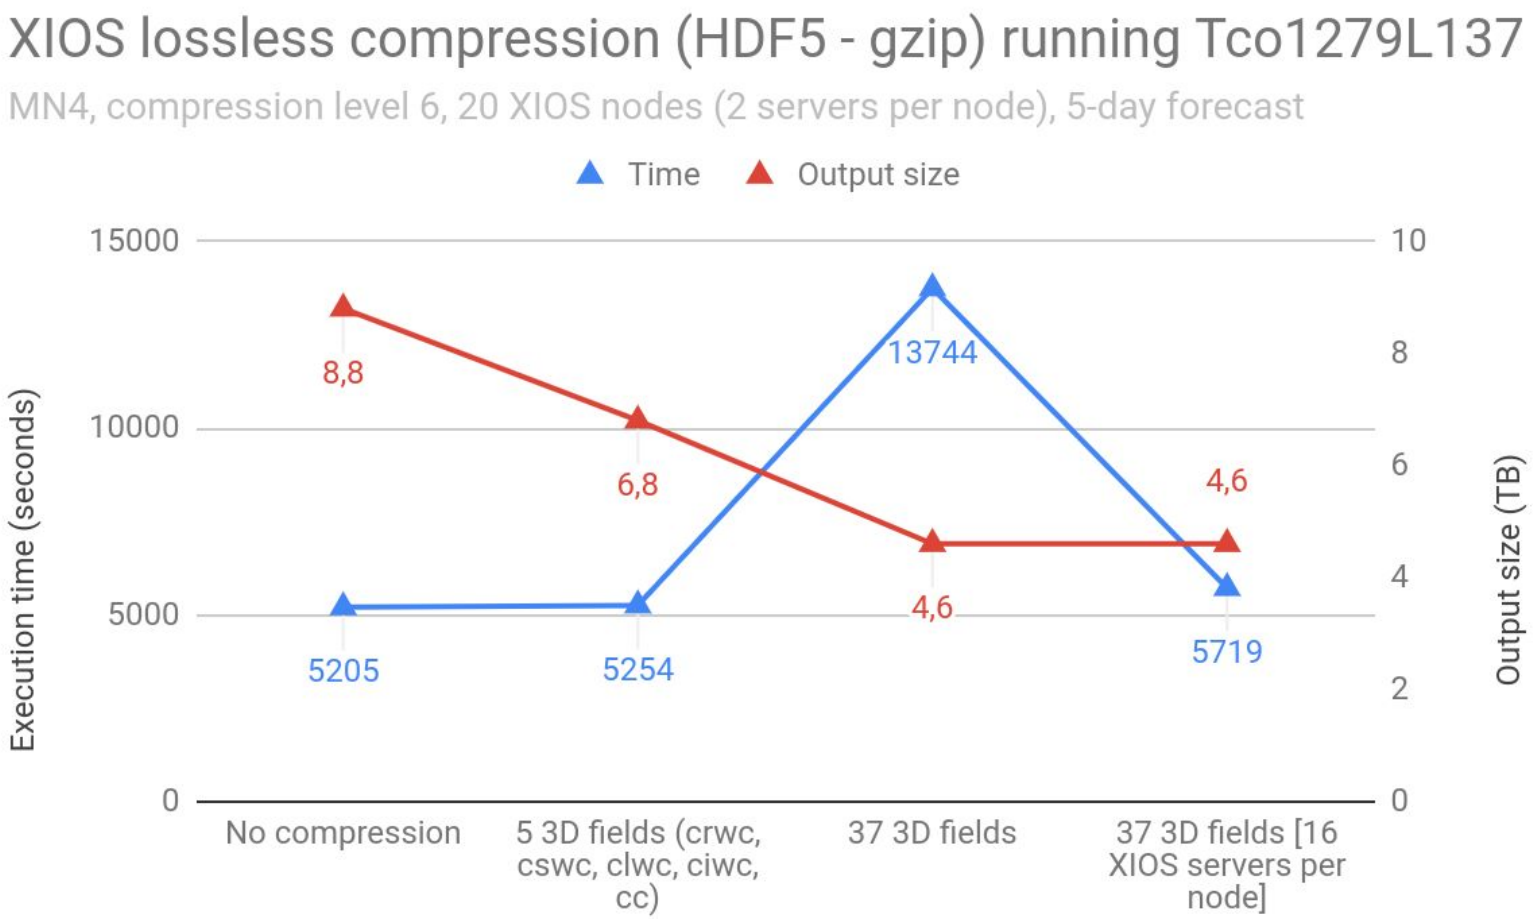

SZ lossy compression filter in XIOS

- Alternatively to lossless compression we wanted to test lossy compression and in particular, the SZ compressor from ANL.
- We started a collaboration in the 11th JLESC workshop to explore if SZ is suitable for XIOS regarding these points:
  - Reach high compression ratios.
  - Enough compression speed to considerably mitigate the I/O overhead.
  - Keep high accuracy.
- The SZ compressor is registered as a third-party filter of HDF5 which facilitates the integration in XIOS:

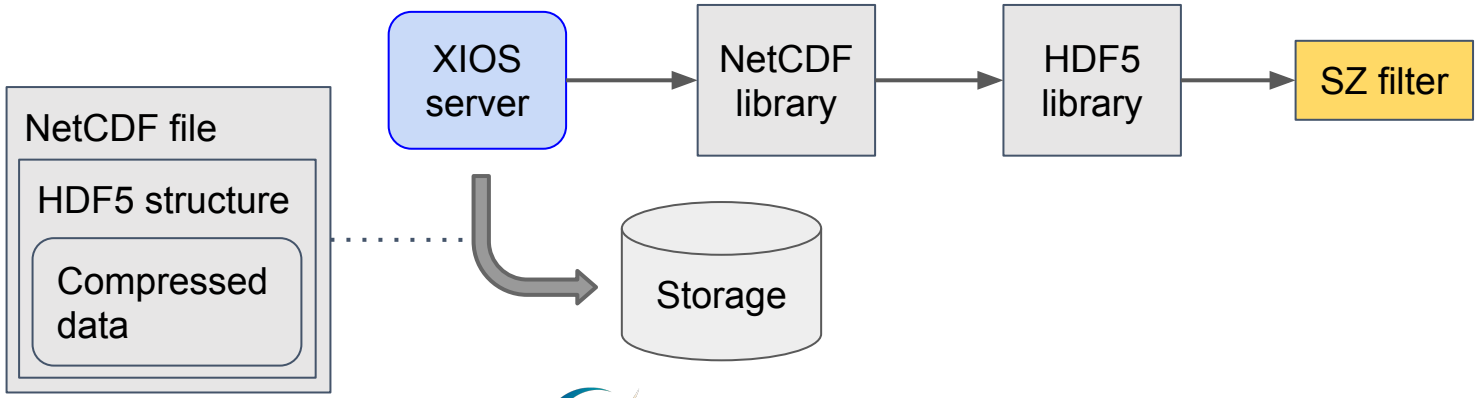

gzip vs. SZ compression: preliminary results

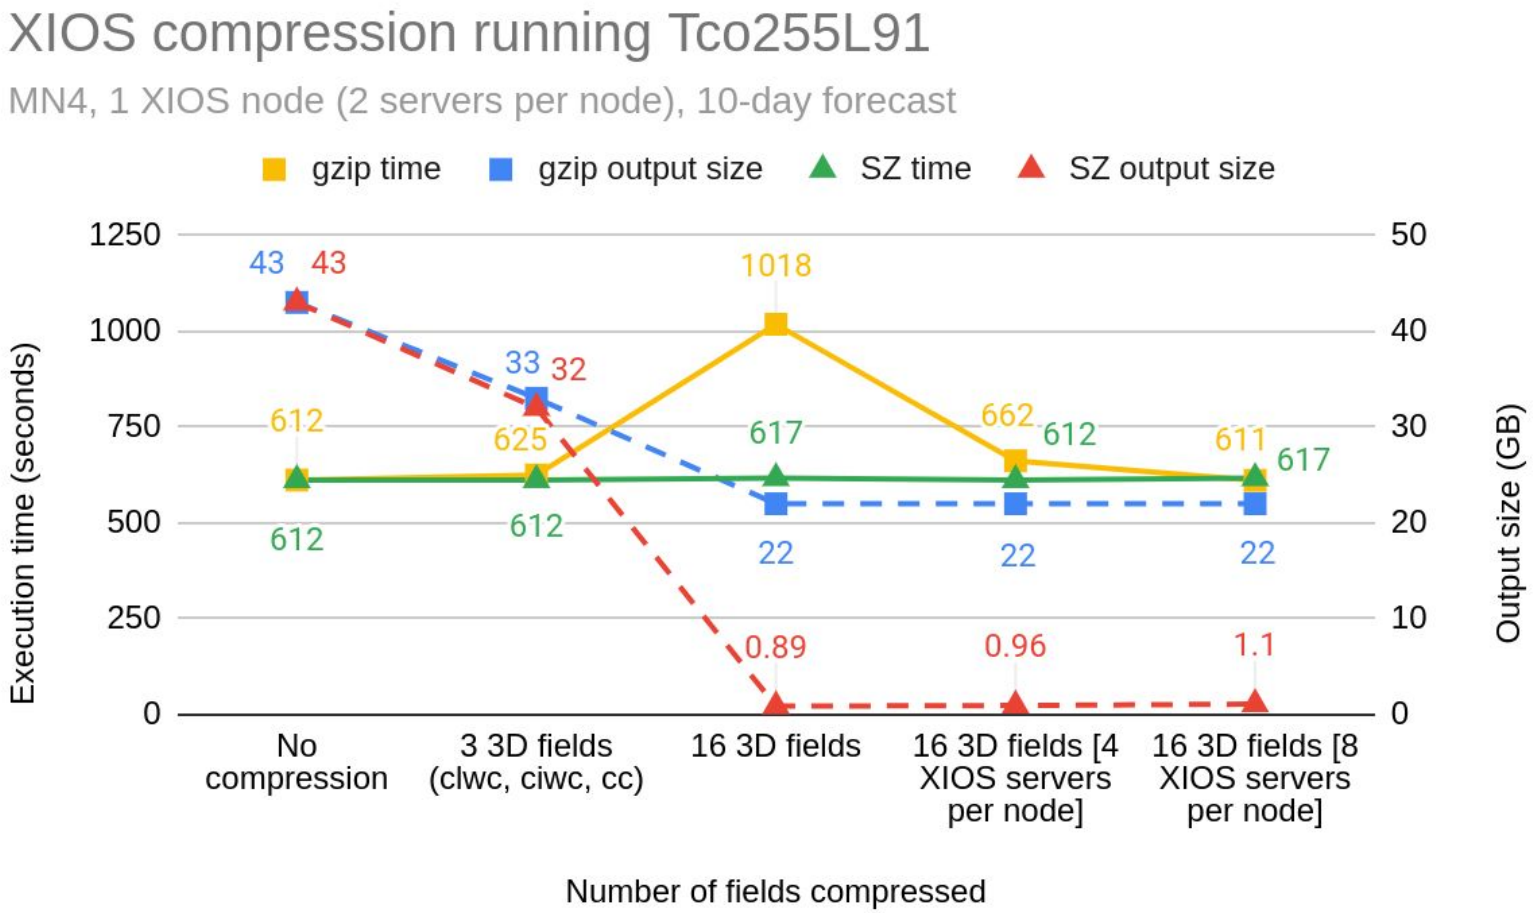

- gzip: compression level 6
- SZ: relative error bound 0.01

Validation

- Specific humidity (q):

```
This is little-endian system.
reading data from q_reduced_pl_0000.dat
Min = 9.999999939225290 2908E-09, Max = 0.027128605172038078308, range = 0.027128595172038139083
Max absolute error = 0.0002728566
Max relative error = 0.010058
Max pw relative error = 26571.363721
PSNR = 48.814234, NRMSE = 0.0036248356857680394394
normErr = 1.430658, normErr_norm = 0.032378
pearson coeff = 0.999425
```

- Temperature (t):

```
This is little-endian system.
reading data from t_reduced_ml_0000.dat
Min = 178.822265625, Max = 312.271209716796875, range = 133.448944091796875
Max absolute error = 1.4429931641
Max relative error = 0.010813
Max pw relative error = 0.007225
PSNR = 45.400245, NRMSE = 0.0053701664860087792303
normErr = 15926.175641, normErr_norm = 0.002937
pearson coeff = 0.999734
```

Ongoing and future work

- ANL has added support to set different compression parameters per field.
  - Change XIOS code to read individual compression parameters and set the SZ filter via the NetCDF API.
  - This will allow us to tune the accuracy of each field depending on scientific needs.

- Discuss with climate scientists the adequate compression parameters for each field, such as specific humidity or temperature.

- Test the SZ filter with the experimental HDF5 parallel I/O.

Open questions and collaboration opportunities

- Continuing the collaboration between ANL and BSC.
- Any feedback will be welcome.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Thank you

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

xavier.yepes@bsc.es