



Fake it 'till you make it

**Zarr-like Access of Existing NetCDF4
Datasets**

Lucas Sterzinger

*Mentors: Chelle Gentemann (Farallon Institute), Julia Kent (NCAR), Kevin
Paul (NCAR)*

July 28, 2021



The Problem





Image Source - Pangeo/Ryan Abernathy

- Peta(Exa?)bytes of geophysical data now available in the cloud
- Much of the data are in NetCDF4/HDF5 format
 - Not efficient for cloud computing


NetCDF vs Zarr

NetCDF



Single File Object

Zarr



Multiple File Objects
(One per chunk)



Plaintext
Metadata
File

ReferenceMaker/ReferenceFileSystem:

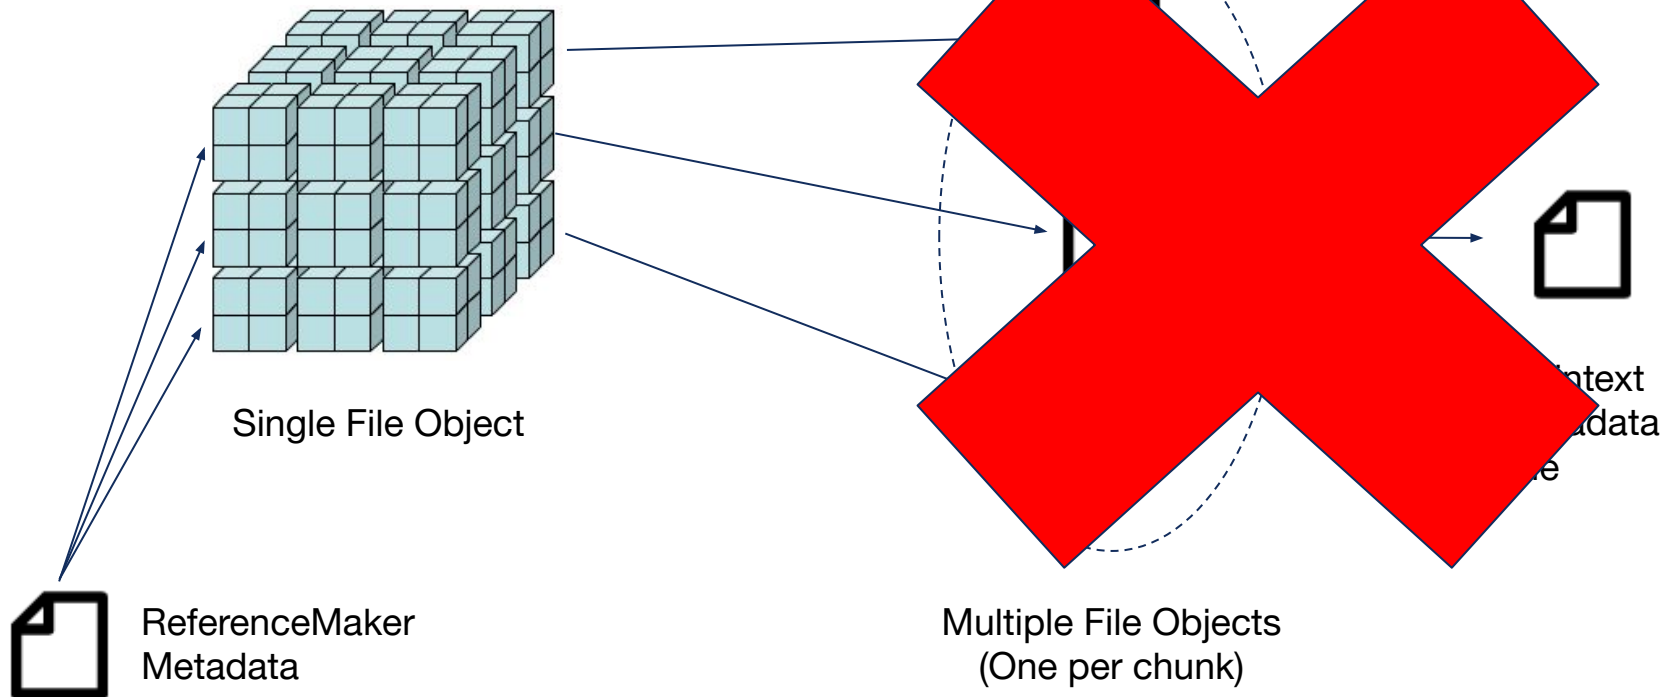
- Part of Intake group's `fsspec` project
- Access existing NetCDF/HDF files
as if they are Zarr
- Zarr-like Metadata files (few MB each)
 - Remote file address (Azure Blob, AWS S3, etc)
 - Data variable and chunking info
- Extracts byte-ranges for individual variables/chunks
 - Instead of individual chunk files in Zarr



ReferenceMaker

NetCDF

Zarr



Creating JSONS

@dask.delayed

Easily distributed with Dask!

```
def gen_json(url):
```

```
    so = dict(
```

```
        mode="rb", anon=True, default_fill_cache=False, default_cache_type="none"
```

```
)
```

```
    outpath = './jsons/' + url.split('/')[-1].split('.nc')[0] + '.json'
```

```
    with fsspec.open(url, **so, account_name='goesewest') as inf:
        h5chunks = SingleHdf5ToZarr(inf, u, inline_threshold=300)
```

Read remote file and process chunks

```
    with open("./reference.json", 'wb') as outf:
        outf.write(ujson.dumps(h5chunks.translate()).encode())
```

Write metadata to file

Reference JSON Example

OR_ABI-L2-MCMIPF-M6_G' X

Filter...

▼ root:

version: 1

▼ u: "az://noaa-goes16/ABI-L2-MCMIPF/2020/002/00/OR_ABI-L2-MCMIPF-M6_G16_s20200020000216_e20200020009524_c20200020010031.nc"

▼ attrs:

.zgroup: "{ "zarr_format": 2, "filters": [{ "name": "GOES R Series Advanced Baseline Imager", "id": "GOES R Series Advanced Baseline Imager", "parent": "GOES R Series Advanced Baseline Imager", "children": [] }] }

.zattrs: "{ "Conventions": "GOES R Series Advanced Baseline Imager Dataset Discovery v1.0", "cdm_data_type": "GOES R Series Advanced Baseline Imager", "date_created": "2020-01-02T00:10:03.1Z", "id": "087244ef-58ee-4c60-a351-2d8b63086250", "institution": "DOC/NOAA/NESDIS > U.S. Department of Commerce, National Oceanic and Atmospheric Administration, National Environmental Satellite, Data, and Information Services", "instrument_ID": "FM1", "instrument_type": "GOES R Series Advanced Baseline Imager", "iso_series_metadata_id": "8c9e8150-3692-11e3-aa6e-0800200c9a66", "keywords": "ATMOSPHERE > ATMOSPHERIC RADIATION > REFLECTANCE, SPECTRAL/ENGINEERING > VISIBLE WAVELENGTHS > REFLECTANCE, SPECTRAL/ENGINEERING > INFRARED WAVELENGTHS > BRIGHTNESS TEMPERATURE", "keywords_vocabulary": "NASA Global Change Master Directory (GCMD) Earth Science Keywords, Version 7.0.0.0.0", "license": "Unclassified data. Access is restricted to approved users only.", "naming_authority": "gov.nesdis.noaa", "orbital_slot": "GOES-East", "platform_ID": "G16", "processing_level": "National Aeronautics and Space Administration (NASA) L2", "production_data_source": "n/a", "production_environment": "OE", "production_site": "NSOF", "project": "GOES", "scene_id": "Full Disk", "spatial_resolution": "2km at nadir", "standard_name_vocabulary": "CF Standard Name Table (v25 - 20 July 2014)" }

File URL (Azure Storage for this example)

Reference JSON Example

```
ministration (NASA) L2 , production_data_source : n/a , production_environment : OE ,  
"production_site": "NSOF", "project": "GOES", "scene_id": "Full Disk", "spatial_resolution":  
n": "2km at nadir", "st", "name": "ABI L2 Cloud and Moisture Imagery Product Table (v35, 20 July 201  
6)", "summary": "Multiple products are digital maps of the Earth's surface in visible, near-IR,  
and IR bands.", "time_coverage_start": "2020-01-02T00:00:21.6Z", "time_coverage_end": "2020-01-02T00:00:21.6Z", "timeline_id": "ABI Mode 6", "title": "ABI L2 Cloud and Moisture Imagery" }
```

Variable (CMI_C01) chunk info,
compression, shape, etc

```
CMI_C01/.zarray: "{ \"chunks\": [ 226, 226 ], \"compressor\": { \"id\": \"zlib\", \"level\": 1 }, \"dtype\": \"<i2\", \"fill_value\": -1, \"filters\": null, \"order\": \"C\", \"shape\": [ 5424, 5424 ], \"zarr format\": 2 }"
```

```
CMI_C01/.zattrs: "{ \"_ARRAY_DIMENSIONS\": [ \"y\", \"x\" ], \"_FillValue\": -1, \"_Netcdf4Dimid\": 0, \"_Unsigned\": \"true\", \"add_offset\": 0.0, \"ancillary_variables\": \"DQF_C01\", \"cell_methods\": \"t: point area: sum (interval: 0.000028 rad)\", \"coordinates\": \"band_id_C01 band_wavelength_C01 t y x\", \"downsampling_method\": \"average\", \"grid_mapping\": \"goes_imager_projection\", \"long_name\": \"ABI Cloud and Moisture Imagery reflectance factor\", \"resolution\": \"y: 0.000056 rad x: 0.000056 rad\", \"scale_factor\": 0.00031746001332066953, \"sensor_band_bit_depth\": 10, \"standard_name\": \"toa_lambertian_equivalent_albedo_multiplied_by_cosine_solar_zenith_angle\", \"units\": \"1\", \"valid_range\": [ 0, 4095 ] }"
```

```
▼ CMI_C01/0.4: [] 3 items  
0: "{u}"  
1: 253643  
2: 468
```

Single Chunk Byte Range

Reference JSON Example

Set “reference” filesystem
and specify JSON file

```
fs = fsspec.filesystem('reference',  
                        fo='reference.json',  
                        remote_protocol='az',  
                        remote_options={'account_name': 'goeseuwest'},  
                        simple_templates=True)
```

Specify remote file
protocol (Azure here) and
authentication

```
m = fs.get_mapper("")  
ds = xr.open_dataset(m, engine='zarr')
```

Pass mapper to xarray and
specify Zarr engine

GOES Workflow



Data stored in same region as compute




4 CPU, 32 GB Memory,
8 Dask Workers

- Created a test workflow on Microsoft Planetary Computer
- Analysis of 24 hours of GOES data
- Test access to data stored in Azure in 3 formats:
 - NetCDF4
 - NetCDF4 + ReferenceMaker
 - Zarr

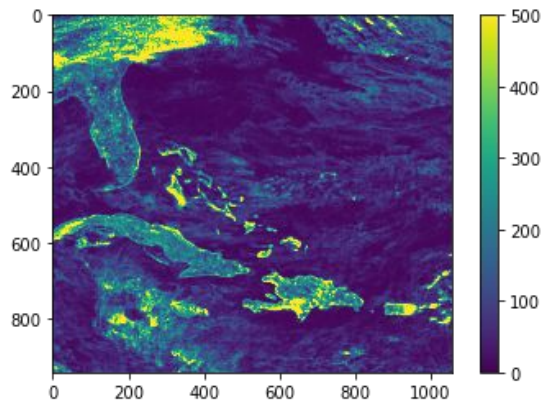
GOES Workflow



Data stored in same region as compute



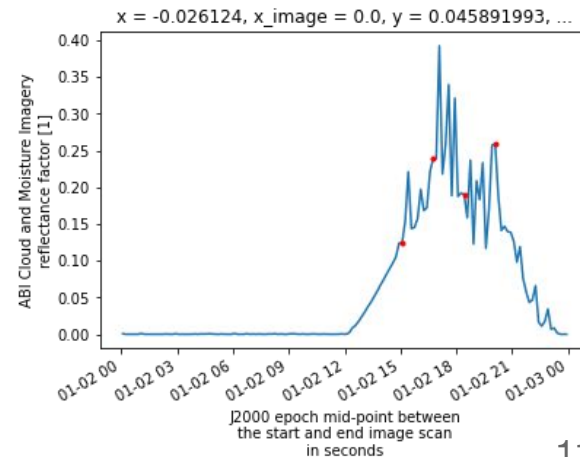
4 CPU, 32 GB Memory,
8 Disk Workers



True Color 2020-01-02 0005 UTC



Veggie Color 2020-01-02 0005 UTC



GOES Processing Times

For 24 hours of GOES-16 “ABI-L2-MCMIPF” product data

<u>Format</u>	<u>Preprocess Time</u>	<u>Data Open Time</u>	<u>Workflow Time</u>	<u>Extra Storage</u>
Remote NetCDF4	0 min	10 minutes	40 min	0 GB
Zarr	1 h 38 min	30 seconds	4 min	52 GB
ReferenceFileSystem	1 h 25 min	35 seconds	5 min 30 s	<u>416 MB</u>

- NetCDF has no upfront preprocess time cost
 - But very slow open and workflow times
- Zarr comes at cost of high storage costs (full data duplication)
- ReferenceFileSystem has speed of Zarr but **≤1% storage cost**
 - Reference file can be compressed further

The Implications

- ReferenceMaker/FileSystem harnesses the cloud-optimization of Zarr without converting any data
- Allows current institutions to cheaply make existing cloud-hosted data more optimized
- Reference files can be also created, hosted, and shared by third parties
- Automatable
 - e.g. Pangeo Forge

Future Work Needed

- Project is still in early development
- Test compatibility across different datasets
- Scale to larger datasets and compare performance to Zarr
- Test performance on a local parallel filesystem (like GLADE)

Acknowledgements

Thank you to the ReferenceFileSystem team!

- Special thanks to Martin Durant (Anaconda) and Rich Signell (USGS) for their help learning and contributing to this project
- <https://github.com/intake/fsspec-reference-maker>

Thank you to Kevin Paul (NCAR), Julia Kent (NCAR), and Chelle Gentemann (Farallon Institute) for their amazing mentorship through the SIParCS program

Thanks to the SIParCS team for working so hard to make a fully-online internship experience so successful

Thanks to Microsoft for providing credits for Planetary Computer



Link to site with interactive code examples, tutorials, resources, and contact info

<https://lucassterzinger.com/2021-siparcs-poster/>


Please feel free to reach out with questions
lsterzinger@ucdavis.edu or twitter [@lucassterzinger](https://twitter.com/lucassterzinger)

If you are a scientist that works with these tools

- Join the Pangeo community
- Talk to project maintainers about your workflow needs
- Open issues/bug reports on GitHub
- Submit pull requests, even if your code isn't 100% ready

PANGEO

Microsoft Planetary
Computer



xarray



Iris

