Xdas: a Python Framework for Distributed Acoustic Sensing

Alister Trabattoni[1], Marie Baillet[1], Martijn van den Ende[1], Diane Rivet[1], Eleonore Stutzmann[2], Claudio Strumia[3], Francesco Biagioli[2].

[1] *Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, IRD, Géoazur, France*
[2] *Institut de Physique du Globe de Paris, Université Paris Cité, Paris, France*
[3] *Università di Napoli Federico II, Complesso University Monte S. Angelo, Napoli, Italy*
Contact: alister.trabattoni@gmail.com

Abstract

Xdas is a Python library designed to manipulate Distributed Acoustic Sensing (DAS) data. It is capable of handling any dataset consisting of dense N-dimensional arrays. The software has the ability to read any DAS file format into a unified Python object abstraction, and to aggregate multi-file datasets produced by any number and kind of instruments, with any different acquisition parametrization into a unique virtual continuous object. It greatly facilitates the temporal and spatial selection and processing of the data while ensuring minimal overhead costs. Xdas utilizes a labeled N-dimensional arrays structure to provide a flexible, self-contained data model that encapsulates both data values and coordinate metadata. This data model adheres to the established formats provided by the NetCDF4/HDF5 file format and the Climate and Forecast (CF) conventions. Xdas has been designed to mirror the application programming interface (API) of the widely used NumPy/Scipy/Xarray Python libraries, with the objective of simplifying the learning process. Xdas is a flexible, extendible solution. The addition of a new file format support or the wrapping of a processing function typically requires less than ten lines of code. By default, Xdas functions are multithreaded, thereby leveraging the computational capacity of multicore machines. Xdas provides online processing routines, enabling continuous processing of long time series and real-time data streams.

Introduction

The advent of Distributed Acoustic Sensing (DAS) technology has firmly placed seismology into the realms of big data. While most DAS experiments last only several weeks, they produce data volumes comparable to what traditional seismological networks have accumulated over the span of years. Seismologists must now meticulously delineate data management plans and address the hardware necessities for storing and processing eventually hundreds of terabytes of data. DAS has introduced a novel observational paradigm by providing continuous sampling of the spatial dimension. In contrast to the previous practice of dealing with spatially sparse 1D time series, seismologists must now handle 2D time and space continuous representations of the wavefield. This necessitates the development of new standard analysis schemes, such as the frequency-wavenumber representation (FK analysis). It also requires the development of new tools that can effectively explore, process, and visualize the high volume of these higher dimensional datasets.

Even though DAS experiments have become commonplace in seismology, DAS remains a nascent technology. Each DAS interrogator model is unique, with each manufacturer utilizing their own set of parameterizations and file formats. At the time of writing this manuscript, the terminology, standard processing routines, and minimal set of metadata are still evolving. As a result, working with DAS data is inherently challenging for several reasons: (i) High data volume. The issue of processing DAS data is mostly constrained by input/output (I/O) and memory limitations. Most of the processing time is dedicated to reading and writing operations, and datasets are often too large to fit in memory. (ii) Chunked file storage. Datasets are partitioned into a multitude of files of reasonable size, ranging from seconds to hours,

depending on the manufacturer. This is a complication that can result in a slower temporal selection process compared to e.g. a single database. (iii) Absence of a standard format. Each manufacturer usually supplies a set of scripts for opening their preferred (and sometimes proprietary) file format. The specific metadata content is dependent on the choices made by the manufacturers, and sometimes depends on the acquisition technology.

Consequently, several software tools have recently been proposed to ease the burden of management and/or processing of DAS data. Some of these have been conceived to solve a specific part of the aforementioned challenges, e.g. as plugins for existing frameworks (Isken et al., 2022). Others have been developed as sets of processing routines and or I/O strategies (Nuwara, 2021; Hu & Li, 2024; Ni et al., 2024). Some are limited to a particular field of research (Bouffaut, 2023). Recently an initiative of providing a generic end-to-end solution was proposed by Chambers et al. (2024). Since DAS toolboxes are still in the early stages of development and usages are still evolving, it is unlikely that any single toolbox will satisfy all users' personal preferences and needs. Therefore, exploring multiple frameworks at the same time will help meet the diverse requirements of the community.

In this paper, we present Xdas, a Python framework for DAS, with the objective to propose a convenient and performant library designed to abstract away numerous logistical tasks, to deal with chunked, larger-than-memory datasets, and to provide performant implementations of the most common processing routines for DAS data. To achieve this objective, Xdas prioritizes the reuse of existing conventions and coding syntax, in addition to the provision of a complete documentation and extensive tutorials. We finally introduce Xpick, a manual picking web application for DAS which illustrates how Xdas can be used as the foundation of an ecosystem of DAS related toolboxes.

Xdas overview

Xdas is designed to support users through the entire data workflow: from preparation and exploration to deploying in production, whether for real-time or large-scale offline applications. Xdas foundations relies on several pillars: (i) Datasets virtualization and labeling: easy continuous data slicing in time and space, even when the underlying data are distributed over numerous smaller files. (ii) High-performance data processing: implementation of NumPy/SciPy functionalities with chunked-processing implementations for massive processing. (iii) Extensibility: great expressiveness and highly customizable routines that facilitate the development of external toolboxes.

Data structures

As DAS metadata harmonization is still an ongoing process, Xdas concentrates for now on a minimal subset of metadata that is applicable to all instruments. This subset consists of the correct labeling of the space and time dimensions. To represent data from DAS instruments, Xdas employs the N-dimensional (N-D) labeled array data structure. Xdas reimplements a subset of the well-established Xarray package (Hoyer & Hamman, 2017) and extends it with features required for DAS data handling that could not be implemented as is with the current state of Xarray (see Interpolated coordinates and Virtual datasets sections). Those versed in Xarray will find Xdas a familiar environment, with the conversion of Python objects between the two systems being a relatively straightforward process. Xdas is not constrained by the limitations of a specific file format. However, given that the majority of contemporary DAS interrogators generate NetCDF4/HDF5 files (it should be noted that the former is a subset of the latter), a number of crucial Xdas features have been developed with the specific intention of capitalizing on the capabilities of that file format (for further details, please refer to the section on virtual datasets). Xdas is compatible with numerous DAS file formats (please refer to the online documentation) and offers a custom NetCDF4 file structure for storing DAS data in accordance with the Climate and Forecast (CF) metadata convention. The addition of support for an additional format can be accomplished with a few lines of code (see Extensibility section).

Data arrays

An N-D labeled array is a data structure that extends the concept of a traditional N-D array by incorporating labels for each dimension, thereby facilitating axis/dimension selection. As opposed to the conventional practice of employing numerical indices, array elements may be referenced by means of labels or coordinate values. Xdas implements the Xarray DataArray structure (Fig. 1a). This structure is composed of an array-like object of any dimensionality, and a set of coordinate objects attached to each dimension. It is composed of three main components: (i) a *data* attribute that contains numerical values, (ii) a *dims* attribute which maps each numerical axis to a (physical) dimension label, and (iii) a *coords* attribute which maps each dimension label to one or several Coordinate objects. Coordinate objects are used for mapping numerical indices to physical values. Xdas allows the use of custom coordinates and data objects. In particular, it provides interpolated coordinates for quasi-uniformly spaced samples, as well as virtual datasets for the management of file-chunked datasets.
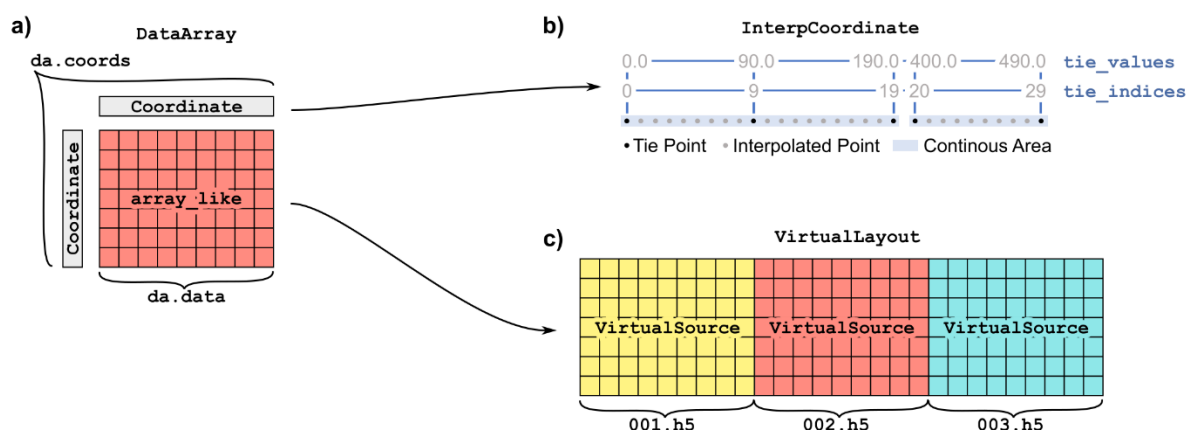

Figure 1. Overview of the data structure used in Xdas. (a) Xdas implements the N-D labeled array structure through a structure called data array. A data array can contain any array-like object and can attach to each of its dimensions one or more coordinate objects. Several types of array-like objects and coordinates can be used but Xdas introduces two objects that are particularly useful for DAS data handling: (b) Interpolated coordinates as defined by the CF conventions. They are a compact way to represent quasi evenly spaced coordinates with potential gaps/overlaps or tiny variations of the sampling rate. Tie points are stored for specific index/value pairs. To retrieve intermediate values interpolation is used. (c) Virtual datasets. HDF5 provides a powerful way to virtually access any number of subfile through a master file with almost no overhead. Here an illustration of the consolidation of three files into a virtually contiguous layout is shown.

Interpolated coordinates

In Xdas, the mapping between indices and related physical values or labels can be encoded in two ways: (i) DenseCoordinate requires a dense vector with the same length as the size of the related dimension. This is the labeling method utilized by Xarray. The main limitation of this approach is that when working with very long yearly time series, the time vector alone may already exceed the capacity of available memory (one year's worth of timestamps sampled at 100 Hz and stored in double-precision format exceeds 25 GB of data). (ii) InterpCoordinate implements a scheme adapted to quasi-evenly spaced samples (Fig. 1b). The strategy implemented in Xdas is in accordance with the conventions set forth by the CF conventions (Gregory, 2003) (see 8.3 Lossy Compression by Coordinate Subsampling – CF 1.11). In Xdas, temporal data is stored as a specified number of tie points. Each tie point is associated with a single index value. Values corresponding to intermediate indices are calculated through linear interpolation. In accordance with the established convention, the presence of two tie points within the same index interval indicates the existence of a discontinuity. This approach enables the handling of gaps and overlaps, as well as the accommodation of internal clock drifting for instruments that are not synchronized by GNSS. For additional flexibility and robustness, Xdas provides methods for the smoothing of minor artificial gaps and overlaps within a specified tolerance.

Virtual datasets

DAS datasets are comprised of a multitude of files. To make a temporal selection, it is necessary to traverse the file tree to identify the relevant files, and then to load them in memory and concatenate them. Xdas streamlines this process and enhances data accessibility through the virtual dataset feature (Fig. 1c) provided by HDF5 (Folk et al., 2011; Koziol et al., 2014). A virtual dataset is constructed based on one or more existing datasets, which are referred to as source datasets. It offers a logical representation of the data through references, thereby facilitating efficient data access without the necessity of rewriting or rearranging the data. In other words, an Xdas virtual dataset can be accessed as if all data were stored in a single file, without data duplication on-disk. Xdas offers a variety of functions, such as open_mfdataarray and open_mftree, which facilitate the conversion of a multifile dataset into a single virtual one. The master dataset comprises all the metadata, along with pointers to the specific data portions of the files. The process of cataloging and indexing the underlying data files must be completed only once to generate a master virtual file which can be subsequently accessed to allow seamless navigation of the entire dataset with minimal latency.

Combined with the data structure of interpolated coordinates, the user can query portions of the complete data set based on a range of datetime objects or timestamps and distances, without having to consider the corresponding N-D array indices in each data file, midnight transitions, leap years, or small inconsistencies in the data size from one file to the next. As any DAS data practitioner knows, the logistical burden that comes with each of these edge cases is non-trivial.

Data collections

The DataArray structure is optimized to handle data and metadata from a singular acquisition. It allows for the potential inclusion of gaps and overlaps but cannot handle change in the acquisition parameters. In particular, a dataset with changes in the number of samples along the spatial dimension cannot be encompassed into a unique DataArray which stores the data as a unique continuous array. Working with multiple acquisitions can happen when exploring different instrument parametrizations; or when multiple instruments are used simultaneously. To address such scenarios, Xdas introduces an additional layer of abstraction, referred to as a DataCollection.
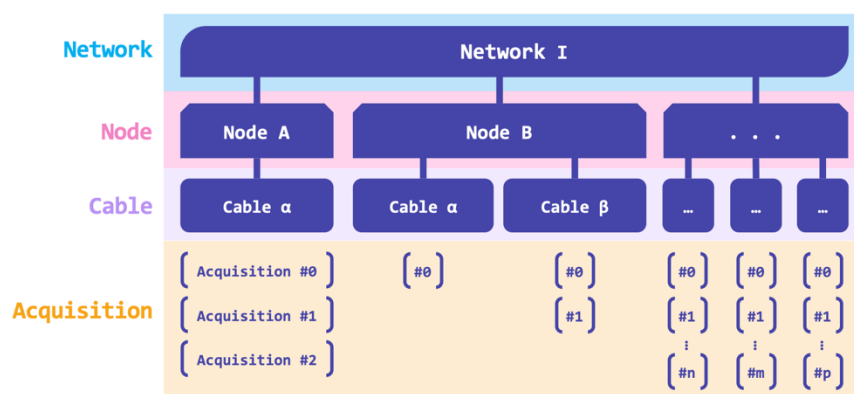

Figure 2. DataCollection structure. This flexible structure is a tree comprising nested dictionaries or lists where the leaves are DataArray objects (here illustrated by brackets). Here an example of a hierarchy with four levels (network, node, cable, acquisition). The first three levels are dictionary-like objects with labeled entries. The last level is a list-like object, containing a varying number of DataArray objects that correspond to different acquisition with potentially different parametrization.

A DataCollection object is a tree-like structure comprising a hierarchical nesting of either other DataCollections or, at the final leaves, DataArrays objects. Each node of the tree can either referenced by a label (e.g., the name of the instrument) or by position (e.g., the number of the acquisition). This flexible structure can represent any hierarchical organization, enabling the accommodation of a wide range of data structures. Fig. 2 depicts a structure that has been adapted for a pool of DAS instruments situated in

4

multiple telecommunication nodes and connected to disparate cables. The interrogator configuration has undergone modifications over time, resulting in the presence of multiple acquisitions. It is anticipated that this flexible structure will be able to adapt to emerging metadata schemes, such as the DAS metadata standard that has been proposed by the DAS-RCN group (Hui Lai et al., 2024).

A subset of the methods applicable to DataArray objects are likewise applicable to DataCollection objects. As a result, the tree is traversed, and the method is applied to every DataArray encountered within the DataCollection. This approach is beneficial when performing temporal selection over a complex multi-acquisition dataset.

Xdas native format

In the preceding sections, the principal structures used in Xdas are presented, along with their corresponding Python objects. Xdas can write any DataArray and DataCollection object to disk through the use of NetCDF4/HDF5 files, utilizing a specific format. The DataArray structure is written by simply reusing the proven NetCDF4 file structure in conjunction with the well-established CF conventions. If no changes have been made to the raw data other than slicing, both DataArray and DataCollection will be written without copying any data (unless the user requests it). The resulting file will rather include pointers to the original files using the high-performance virtual dataset feature of HDF5. In any other case the data will be written to disk. Due to their nature, raw DAS data are large, and it may be beneficial to save the data of interest in a compressed form. Xdas incorporates all standard NetCDF4 and HDF5 compression filters with the h5netcdf and the hdf5plugin libraries. In particular, Xdas includes the ZFP compression algorithm (Lindstrom, 2014), which is suited for conservative lossy compression of floating-point data (Issah & Martin, 2024). Compared to lossless compression, the ZFP compression ratio of DAS data is typically reaching a factor of four without the introduction of any discernible errors. The DataCollection structure is saved by employing the group feature of NetCDF4 to write multiple DataArrays in a single file. The Xdas native format is a valuable tool for the storage and transfer of processed data, as well as for the compression of raw data for long-term storage.

Typical workflow

The utilization of Xdas for data analysis is analogous to that of Xarray, which is, in turn, is analogous to NumPy (Harris et al., 2020) and/or SciPy (Virtanen et al., 2020). In contrast to other software that provides only large, opaque processing systems, Xdas offers a set of small, modular components that allow users to customize their workflow while maintaining flexibility. To facilitate simplicity, Xdas attempts to adhere to the function signatures observed in Xarray, NumPy, and SciPy, employing two conventions: (i) Rather than providing the axis number, the dimension label must be provided; (ii) a parallel keyword argument may be passed to parametrize multithreading processing (see Performance).

To illustrate a typical workflow, we process an example event (Fig. 3). In the case that a master virtual file has not been created prior to this step, it must first be created (step b in Fig. 3). This process may take a few minutes to complete for datasets containing millions of files. In subsequent instances, the master file can be accessed with minimal delay, usually less than a second (step c). Subsequently, we query a time and space selection of the data (step d). At this stage, the data can be loaded into memory. As an example, the FK representation can be plotted from scratch in a few lines of code (step e), applying a sequence of NumPy/SciPy-like operations to the data set. This example illustrates the versatility and expressiveness of Xdas and its ability to handle arbitrary coordinates, such as frequency and wavenumber in this case. Certain specific processing routines provided by Xdas can be applied, such as a conversion to velocity (Trabattoni et al., 2023) and time/space decimations (step g). The processed waveforms can then be plotted either as a 2D image or trace by trace in a single line of code (steps f, g, h), leveraging the encapsulated metadata to ensure accurate labeling of the axes. This style of processing and visualization is similar to that of ObsPy or Xarray, and therefore be familiar to most users.

a)
```python
import numpy as np
import xdas as xd
import xdas.signal as xs

# link virtual dataset (run it once)
da = xd.open_mfdataarray("abyss/*.hdf5", engine="asn")
da.to_netcdf("da.nc", virtual=True)

# open virtual dataset
da = xd.open_dataarray("da.nc")

# slice in time and space and load data
da = da.sel(
    time=slice("2024-01-22T22:38:30", "2024-01-22T22:39:00"),
    distance=slice(None, 120_000),
)
da = da.load()  # optional

# apply gain
da = 8.36e-11 * da

# compute FK representation and plot
fk = xs.taper(da, dim="distance")
fk = xs.taper(fk, dim="time")
fk = xd.fft.rfft(fk, dim={"time": "frequency"})
fk = xd.fft.fft(fk, dim={"distance": "wavenumber"})
fk = 20 * np.log10(np.abs(fk))  # numpy ufuncs
fk.plot(robust=True, interpolation="antialiased")

# convert to velocity
da = xs.integrate(da, "distance")
da = xs.sliding_mean_removal(da, wlen=1000.0)

# decimate in time and space and plot
da = xs.decimate(da, 16, ftype="fir", dim="distance")
da = xs.decimate(da, 4, ftype="iir", dim="time")
da.plot(yincrease=False, vmin=-1e-6, vmax=1e-6)

# extract one trace
tr = da.sel(distance=60_000, method="nearest")
tr.plot()
```

b)
```
Fetching metadata: 100%|████████| 8640/8640
[00:22, 376.42it/s]
Linking database: 100%|████████| 8640/8640
[00:02, 4243.04it/s]
```

c)
```
<xdas.DataArray (time: 5400000, distance: 10000)>
VirtualSource: 100.6G (int16)
Coordinates:
  * time (time): 2024-01-22T00:00:07.000 to 2024-
01-23T00:00:06.984
  * distance (distance): 0.0 to 153179.709
```

d)
```
<xdas.DataArray (time: 1875, distance: 7834)>
VirtualSource: 28.0M (int16)
Coordinates:
  * time (time): 2024-01-22T22:38:30.008 to
2024-01-23T22:38:30.992
  * distance (distance): 0.0 to 119997.6662
```

e)
```
<xdas.DataArray (frequency: 938, wavenumber: 7834)>
[[132.160 130.667 ... 129.657 130.667],
 ...,
 [ 99.072 116.218 ... 121.245 117.397]]
Coordinates:
  * frequency (frequency): 0.0 to 31.25
  * wavenumber (wavenumber): -0.033 to 0.033
```
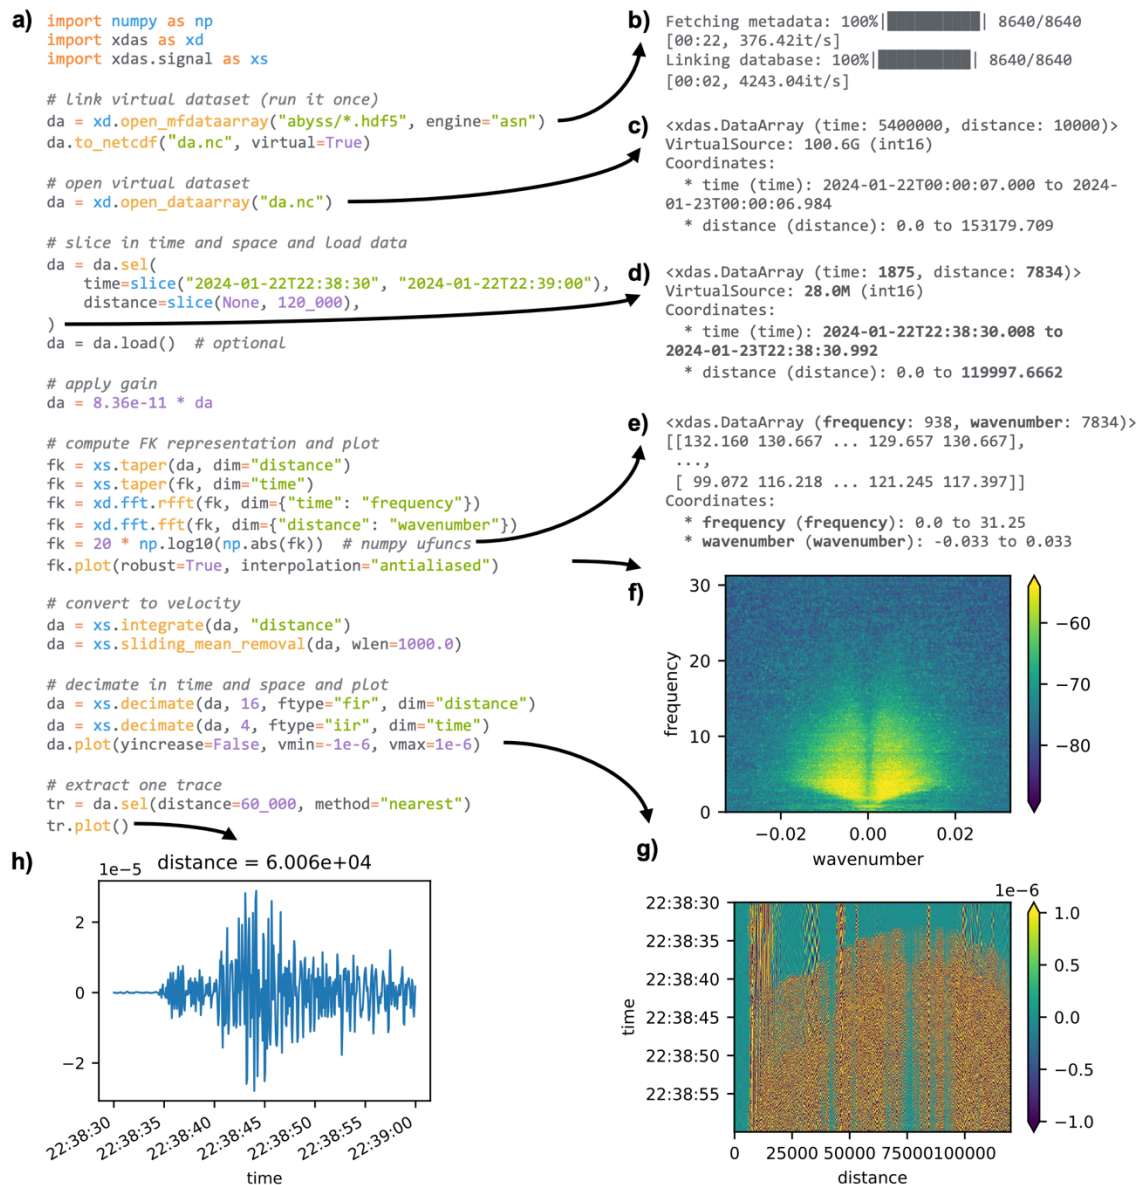
Figure 3. Typical workflow. (a) Example python scripts. (b) Liking of the virtual data array. Must be run only once. Here one day of ten-second files are linked, which takes about 20 s. (c) Virtual data array once all files have been linked. Opening it take less than one second (d) Data array after slicing. In bold, fields that have changed. Note that now the data array only contains 28MB (rather than 100 GB for the entire data set) and hence can be loaded in memory (note that Xdas will automatically load data when required). It is up to the user to apply any instrumental response, usually as a conversion factor. (e) The data array structure can be used with other coordinate systems than time-space dimensions. Here an FK result is stored. (f) The FK is plotted in a single line of code and metadata is used to produce a labeled plot. (g) Same but for the converted-to-velocity and decimated waveforms. (h) Xdas allows trace by trace analysis.

Performance

Given that data reading, transfer, and writing is frequently the limiting factor when processing DAS data, multithreaded optimization is relevant when the data are locally available, and not from cloud storage or Network-Attached Storage. In this section, we consider a local-access scenario.

Server-grade disk arrays allow data reading speeds of up to several gigabytes per second, and CPU core counts often exceed several dozen. In this context, the objective is to provide a computational pipeline that can optimally use the available compute resources to match the input/output (I/O) access speed. Xdas achieves this by avoiding unnecessary overhead by processing the data as a whole N-D array (and not trace

by trace) and by implementing a versatile multithreading solution. Under typical conditions, Xdas provides one order-of-magnitude speed gain compared to strategies that involve splitting the data into traces and then joining the processed traces (Fig. 4a) using standard tools such as ObsPy (Beyreuther et al., 2010). When multithreading is enabled, Xdas leverages the multithreaded acceleration of NumPy/Scipy operations scaling almost linearly with the number of threads (Fig. 4b).


Figure 4. Benchmark of Xdas. Each measure consists in the smaller elapsed time over 7 runs. (a) Single-threaded temporal decimation using a 12 order IIR filter for a 10 000 samples long signal with varying number of channels. Obspy scales linearly (ideal linear scalability in black dotted lines), meaning it has a fixed extra overhead per channel whereas Xdas take advantage of processing every channel in a single run, reducing the overhead per channel. (b) Multi-threaded decimation with Xdas. A spatial decimation of a factor 16 is applied with FIR filtering. The single-threaded version uses the base SciPy function and can be used as reference. Depending on the chunk size (different solid-colored lines, legend gives the number of samples) and number of threads, the overhead implied by multi-threading can limit the scalability (ideal scalability in black dotted lines). Note that throughput of several GB/s can be achieved.

Massive and real-time processing

A major challenge with DAS data lies in the processing of extensive time periods. In most cases, only small subsets of the data (up to few minutes) can be loaded into memory at once. To perform a posteriori analysis of the entire dataset, continuous chunked-based processing is necessary. A comparable challenge is encountered in applications requiring quasi-real time processing. Xdas provides a solution for the online processing of data while simultaneously optimizing input/output operations.

The chunk processing architecture is based on an overall first-in-first-out (FIFO) strategy, which is particularly well-suited for real-time applications. Furthermore, it allows for the use of infinite impulse response (IIR) processing, where the outcome of the previous iteration is needed to compute the next iteration. This contrasts with finite impulse response (FIR) processing, where the data can be processed in any order, provided that sufficient overlap is loaded.

To ensure the continuity of the processing over chunks, Xdas introduces "Atoms" which are elementary operations that are able to retain any final state from the previous processed chunked that is required as the initial state of the next iteration (Fig. 5a). Atoms are meant to be chained together and used as processing pipeline. To maximize throughput, each Atom is parallelized individually. In the case of straightforward filtering routines, this approach can achieve throughput of several gigabytes per seconds and to generally match the I/O bandwidth.

A common strategy to maximize processing bandwidth is to divide a pipeline into three concurrent operations: data reading (or awaiting new data in real-time), computation, and data writing. By staggering the output of each operation (as illustrated in Fig. 5b), the latency associated with I/O operations can be circumvented (hence the common term "latency hiding"). The user is thus required to create three objects:

a data loader (which could be a real-time data stream or an Xdas DataCollection), a signal processing chain (comprised of a sequence of Atoms), and a data writer.

Since the data is processed in discrete units, or "chunks", the optimal chunk size should be selected based on three key considerations: (i) Increasing the chunk size can lead to enhanced performance, while simultaneously reducing the overheads associated with the process of chunking. (ii) The size of the data chunk is constrained by the amount of available memory. Xdas maintains a maximum of three chunks in memory at any given time. However, the memory allocation required to store intermediate results must also be considered. (iii) In a real-time setup, an increase in the chunk size will result in a corresponding increase in delay.
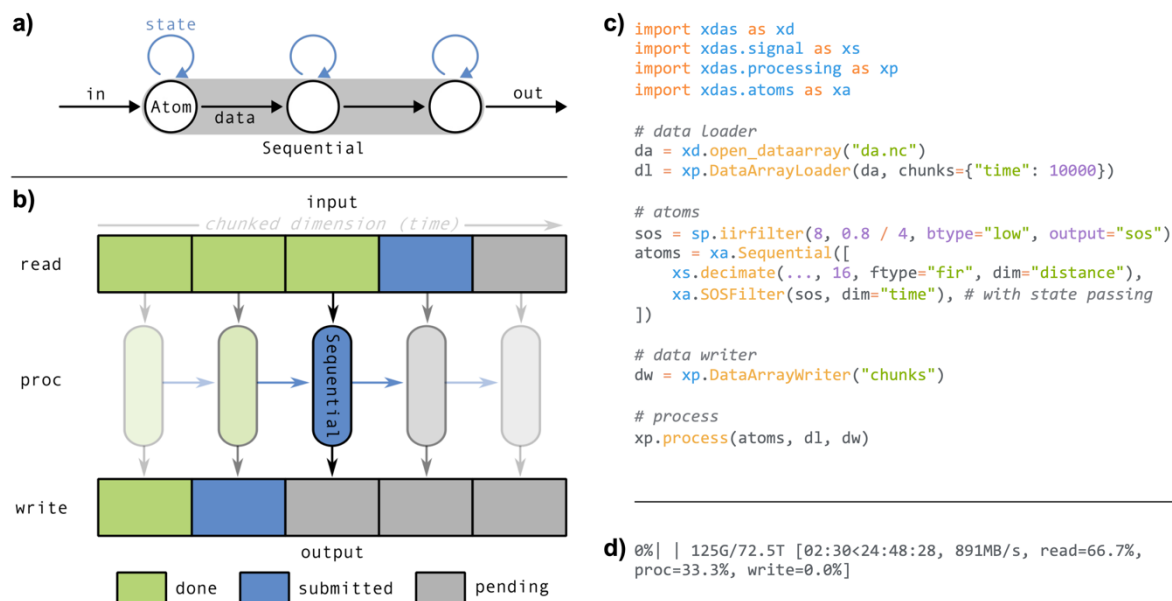

Figure 5. Xdas introduces the Atom object to enable chunk-based processing. (a) Each Atom takes one input and returns one output. To ensure continuity in a very generic way, Atom objects use state passing. The final state of the preceding iteration is used as the initial state of the next iteration. Atom objects can be concatenated in a processing chain through the Sequential Atom (in this example three atoms are shown). This latter also takes one input and returns one output hence is an Atom by definition. (b) Sequential graph of tasks. To perform continuous processing of DAS data, Xdas uses a chunk-based approach. Each input chunk (top rectangles) goes through a processing chain made of a sequence of Atom objects. Processed chunks are then written to disk chunk by chunk (bottom rectangles). While the input is generally DAS data, the output can be of any type (e.g., phase picks). Xdas decouples the I/O tasks (read, write) from the processing tasks (proc). As soon as a chunk is available in memory, the processing chain is applied, and the result is put in a writing queue. (c) Related code for a conventional space and time decimation routine to reduce data volume. First a data loader is defined specifying the chunk size in units of samples. The processing chain is composed of callable objects that takes and returns one chunk. Here: (i) a spatial decimation that is natively parallelized along time, (ii) a state aware filter (that uses cascaded second-order sections) followed by an integer decimation. Lastly, a data writer is defined and the processing chain is applied. (d) Monitoring of the processing.

Extensibility

Xdas is designed to be extensible. Two common use cases include: (i) adding a user-defined processing routine and (ii) supporting an additional file format. In most cases, extending Xdas entails the writing of a few lines of code (Fig. 6). The process involves addressing the two primary aspects of a Xdas data array: unpacking the data and coordinates objects, potentially performing operations on them, and then repacking them into a DataArray object. To add a new file format, the user is required to specify a function that must read one file and output a DataArray. This function can then be passed as an engine keyword argument to all Xdas reading functions. The custom reading function must fetch and parse the data and coordinates information. To wrap a Python function meant to work on NumPy arrays, minimal work is generally required. If the processing is done along an axis, the dimension label must be converted into its numerical

axis value. The function can be then applied to the data array values and, if necessary, the coordinates must be updated.

```python
import h5py
import numpy as np
from xdas import DataArray
from xdas.virtual import VirtualSource

def read(fname):
    with h5py.File(fname, "r") as file:
        t0 = np.datetime64(round(1e9 * file["t0"][()]), "ns")
        dt = np.timedelta64(round(1e9 * file["dt"][()]), "ns")
        dx = file["dx"][()]
        data = VirtualSource(file["data"])
    nt, nd = data.shape
    t = {"tie_indices": [0, nt - 1], "tie_values": [t0, t0 + (nt - 1) * dt]}
    d = {"tie_indices": [0, nd - 1], "tie_values": [0.0, (nd - 1) * dx]}
    return DataArray(data, {"time": t, "distance": d})
```

```python
import scipy.signal as sp
from xdas.parallel import parallelize
from xdas.atoms import atomized

@atomized
def decimate(da, q, *args, dim="last", parallel=None, **kwargs):
    axis = da.get_axis_num(dim)
    across = int(axis == 0)
    func = parallelize(across, across, parallel)(sp.decimate)
    data = func(da.values, q, *args, **kwargs)
    return da[{dim: slice(None, None, q)}].copy(data=data)
```

Figure 6. Extensibility of Xdas. (a) Adding the support for a new file format generally consists into providing the path to the data array, and parsing the start time and spatial and temporal spacing. Here a hypothetical simple file format is used as example. For real world application, the four lines into the with statement block must be adapted. (b) Wrapping a signal processing function that is applied along one axis require to convert the dimension name in its related axis number, processing the data (optionally in a parallelized manner), and handling coordinates changes if any. Here we reimplement the decimate function almost as it is implemented in Xdas.

Xpick: a toolbox for manual picks edition

In the long term, rather than developing a large and complex monolithic library, we propose the creation of a constellation of toolboxes. This approach will facilitate the development of an ecosystem of tools for addressing diverse advanced aspects of data analysis in a more flexible and community-oriented manner. By offering a comprehensive solution for DAS data management and processing, Xdas serves as the foundation for the development of other DAS toolboxes.

One illustrative example of such a toolbox is Xpick (Fig. 7). Xpick is a web-based application designed for the editing of manual picks on DAS data. The software is designed to perform free-hand delineations of seismic phase arrivals on two-dimensional data, enabling the manual identification of any phase along thousands of DAS traces. Xpick can be operated on a remote machine and accessed via a local browser, which makes it suitable for deployment as a web service by any observatory.
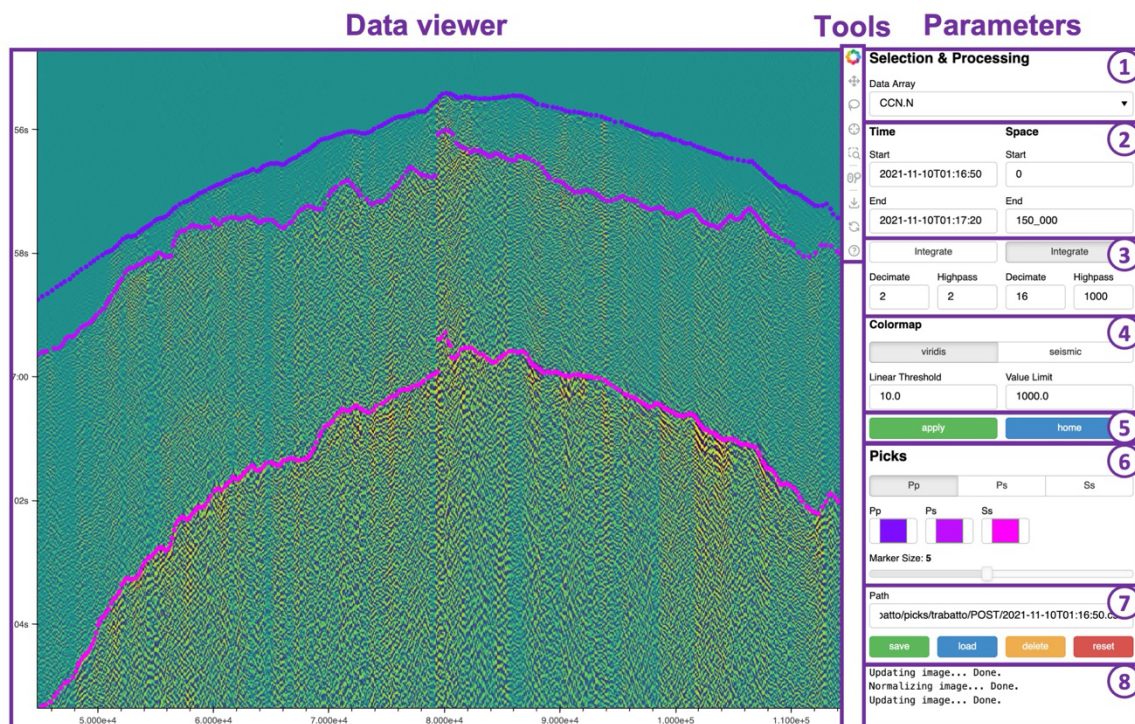
Figure 7. Xpick, the Xdas toolbox for editing manual picks on DAS data. Users explore any dataset both in time and in space and draw freely on the 2D plot of the waveform to delineate arrivals. This toolbox is an example of external tools that will be included in the Xdas ecosystem.

Conclusion

Xdas is an open-source Python framework that provides the fundamental functionalities for DAS data management and processing. The software is capable of handling large multi-file datasets of any format. It provides its own data format based on existing conventions and implements a state-of-the-art compression algorithm. It facilitates the time-based querying across multiple data files, provides multithreaded common processing routines, enables online processing for massive or real-time data analysis, and allows for straightforward extension. Xdas is intended to serve as the basis for other Python toolboxes that focus on advanced DAS data analysis and processing. We encourage anyone interested to contribute to the improvement of Xdas' main functionalities or to develop their own Xdas-based toolboxes within the GitHub xdas-dev organization (https://github.com/xdas-dev).

Data and Resources

Xdas is available at https://github.com/xdas-dev/xdas. A complete documentation can be accessed at https://xdas.readthedocs.io. Tutorials are provided at: https://github.com/xdas-dev/tutorials. Xdas can easily be installed via pip (https://pypi.org/project/xdas). The scripts used to produce the figures of this study can be found at (https://github.com/atrabattoni/xdas_framework).

Declaration of Competing Interests

The authors acknowledge there are no conflicts of interest recorded.

Acknowledgement

References

Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010). ObsPy: A Python Toolbox for Seismology. *Seismological Research Letters*, *81*(3), 530–533. https://doi.org/10.1785/gssrl.81.3.530

Bouffaut, L. (2023). *DAS4Whales: A Python package to analyze Distributed Acoustic Sensing (DAS) data for marine bioacoustics* (v0.1.0) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.7760187

Chambers, D., Jin, G., Tourei, A., Issah, A. H. S., Lellouch, A., Martin, E., Zhu, D., Girard, A., Yuan, S., Cullison, T., Snyder, T., Kim, S., Danes, N., Punithan, N., Boltz, M. S., & Mendoza, M. M. (2024). DASCore: A Python Library for Distributed Fiber Optic Sensing. *Seismica*, *3*(2), Article 2. https://doi.org/10.26443/seismica.v3i2.1184

Folk, M., Heber, G., Koziol, Q., Pourmal, E., & Robinson, D. (2011). An overview of the HDF5 technology suite and its applications. *Proceedings of the EDBT/ICDT 2011 Workshop on Array Databases*, 36–47. https://doi.org/10.1145/1966895.1966900

Gregory, J. (2003). The CF metadata standard. *CLIVAR Exchanges*, *8*(4), 4.

Harris, C. R., Millman, K. J., Van Der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., Van Kerkwijk, M. H., Brett, M., Haldane, A., Del Río, J. F., Wiebe, M., Peterson, P., … Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357–362. https://doi.org/10.1038/s41586-020-2649-2

Hoyer, S., & Hamman, J. (2017). xarray: N-D labeled Arrays and Datasets in Python. *Journal of Open Research Software*, *5*(1), 10. https://doi.org/10.5334/jors.148

Hu, M., & Li, Z. (2024). DASPy: A Python Toolbox for DAS Seismology. *Seismological Research Letters*. https://doi.org/10.1785/0220240124

Hui Lai, V., Hodgkinson, K. M., Porritt, R. W., & Mellors, R. (2024). Toward a Metadata Standard for Distributed Acoustic Sensing (DAS) Data Collection. *Seismological Research Letters*, *95*(3), 1986–1999. https://doi.org/10.1785/0220230325

Isken, M. P., Heimann, S., Wollin, C., Bathke, H., & Dahm, T. (2022). *Lightguide—Seismological Tools for DAS data* (v0.1.0-beta) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.6580579

Issah, A. H. S., & Martin, E. R. (2024). Impact of Lossy Compression Errors on Passive Seismic Data Analyses. *Seismological Research Letters*, *95*(3), 1675–1686. https://doi.org/10.1785/0220230314

Koziol, Q., Pourmal, E., & Fortner, N. (2014). RFC: HDF5 virtual dataset. *Documentation, The HDF5 Group*. https://docs.hdfgroup.org/archive/support/HDF5/docNewFeatures/VDS/HDF5-VDS-requirements-use-cases-2014-12-10.pdf

Lindstrom, P. (2014). Fixed-Rate Compressed Floating-Point Arrays. *IEEE Transactions on Visualization and Computer Graphics*, *20*(12), 2674–2683. https://doi.org/10.1109/TVCG.2014.2346458

Ni, Y., Ragland, J., & Fatland, R. (2024). *niyiyu/DASstore: V0.2.1* (v0.2.1) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.10714765

Nuwara, Y. (2021). *DASLab: Distributed Acoustic Sensing Lab* (1.0.0) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.5797215

Trabattoni, A., Biagioli, F., Strumia, C., van den Ende, M., Scotto di Uccio, F., Festa, G., Rivet, D., Sladen, A., Ampuero, J. P., Métaxian, J.-P., & Stutzmann, É. (2023). From strain to displacement: Using deformation to enhance distributed acoustic sensing applications. *Geophysical Journal International*, *235*(3), 2372–2384. https://doi.org/10.1093/gji/ggad365

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., Van Der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., … Vázquez-Baeza, Y. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2