



*International  
Virtual  
Observatory  
Alliance*

## **Time Series Cube Data Model Version 1.1**

**IVOA Note 2017-02-05**

Working group

Time domain interest group

This version

<http://www.ivoa.net/documents/cubeDM/20170205>

Latest version

<http://www.ivoa.net/documents/cubeDM>

Previous versions

Author(s)

Jiří Nádvořník, Petr Škoda, Dave Morris, Pavel Tvrđík

Editor(s)

Jiří Nádvořník

## Abstract

The purpose of this document is to create a data model and its serialization for expressing generic time series data. Already existing IVOA data models are reused as much as possible. The model is also made as generic as possible to be open to new extensions but at the same time closed for modifications. This enables maintaining interoperability throughout different versions of the data model.

We define the necessary building blocks for metadata discovery, serialization of time series data and understanding it by clients.

We present several categories of time series science cases with examples of implementation. We also take into account the most pressing topics for time series providers like tracking original images for every individual point of a light curve or time-derived axes like frequency for gravitational wave analysis.

The main motivation for the creation of a new model is to provide a unified time series data publishing standard - not only for light curves but also more generic time series data, e.g., radial velocity curves, power spectra, hardness ratio, provenance linkage, etc.

The flexibility is the most crucial part of our model - we are not dependent on any physical domain or frame models. While images or spectra are already stable and standardized products, the time series related domains are still not completely evolved and new ones will likely emerge in near future. That is why we need to keep models like *Time Series Cube DM* independent of any underlying physical models. In our opinion, this is the only correct and sustainable way for future development of IVOA standards.

## Status of This Document

This is an IVOA Note expressing suggestions from and opinions of the authors. It is intended to share best practices, possible approaches, or other perspectives on interoperability with the Virtual Observatory. It should not be referenced or otherwise interpreted as a standard specification.

A list of current IVOA Recommendations and other technical documents can be found at <http://www.ivoa.net/Documents/>.

## Contents

<b>1</b>	<b>Term Definition</b>	<b>5</b>
<b>2</b>	<b>Introduction</b>	<b>5</b>

<b>3</b>	<b>Dependent data models</b>	<b>6</b>
3.1	Imported Data models . . . . .	7
3.1.1	Dataset DM . . . . .	8
3.1.2	ObsCore DM . . . . .	8
3.1.3	VO-DML . . . . .	8
3.2	Extended Data models . . . . .	8
3.2.1	Cube DM . . . . .	8
3.2.2	Pixelated Cube DM . . . . .	9
3.2.3	Image DM . . . . .	9
3.2.4	Sparse Cube DM . . . . .	9
3.2.5	Time Series Cube DM . . . . .	9
<b>4</b>	<b>Time Series Cube DM</b>	<b>9</b>
4.1	DataSet DM . . . . .	10
4.2	Cube DM . . . . .	12
4.2.1	Sparse Cube class . . . . .	12
4.2.2	Time Series Cube class . . . . .	12
4.2.3	CubeAxis class . . . . .	12
4.3	Axis Domain DMs . . . . .	13
4.4	Time Series Cube DM summary . . . . .	13
<b>5</b>	<b>Discovery - ObsCore</b>	<b>13</b>
<b>6</b>	<b>Serialization</b>	<b>13</b>
<b>7</b>	<b>Supported science cases</b>	<b>14</b>
7.1	Requirements . . . . .	14
7.2	Science case definitions . . . . .	15
7.3	Simple Light curves . . . . .	16
7.3.1	Examples . . . . .	16
7.4	Groups of light curves . . . . .	16
7.4.1	Examples . . . . .	18
7.5	Generic time series . . . . .	18
7.5.1	Examples . . . . .	20
7.6	Complex light curves . . . . .	20
7.6.1	Examples . . . . .	20
7.7	Gravitational waves . . . . .	21
7.7.1	Examples . . . . .	22
<b>8</b>	<b>Reference implementation</b>	<b>25</b>
<b>9</b>	<b>Summary</b>	<b>25</b>

## Acknowledgements

This work was supported by grant LD-15113 of Ministry of Education Youth and Sports of the Czech Republic and by grant No. SGS16/123/OHK3/1T/18 from Czech Technical University in Prague.

## 1 Term Definition

- Data cube - When we are speaking about a data cube, we have in mind the pure database point of view. In astronomy a term data cube is often implicitly understood as an image cube or a spectral cube and we do not want to be restricted only to these domains. In this document, a data cube is much closer to an OLAP data cube (Kimball (1997a)) (while again not restricting us to business-related domains). This data cube is usually implemented in the database as a star-schema with its dimensions being our axis domains (Kimball (1997b)).
- Time Series Cube DM - The *Time Series Cube data model (DM)* defines data structures for holding metadata and data of a time series data cube.
- Axis domain - By axis domain we mean the physical interpretation of an axis of the data cube. This is usually axis metadata, e.g., spectral limits of that axis, resolution, etc. From the IT perspective, while *Time Series Cube DM* holds the *data*, axis domain model stores the relationships and context which enables us to get *information* from the data.
- Dataset - We mean an IVOA Dataset as specified in *Observation Data Model* (Mireille Louys (2011)), also known as the *ObsCore DM*, i.e. a collection in which can be stored any kind of observation. As such, a dataset is holding information common for every element of its collection, e.g., facility, instrument or creator.

## 2 Introduction

This is a working draft of a new approach for representing generic time series data. It was designed to be modular and extensible, new changes will extend the model, not modify it. Extending a model means reusing what was in the extended model, while adding new structures to it. That said, any part of this design is open to discussion. If you have any ideas how it can be improved, please contact us on Time Domain Interest Group mailing list [voevent@ivoa.net](mailto:voevent@ivoa.net).

We can incorporate any kind of axis in the data cube, where all other up-to-now approaches tried to embed domain-specific metadata inside the model. They described specialized axes relevant only for light curves - SimpleTimeSeries (Matthew Graham (2014)) and SSAP-based approaches (Doug Tody (2012)). For these approaches, making any enhancement to the model required re-work of what was already in the model, not an extension.

The description of axis domain, i.e. the metadata which enables us to extract information and knowledge from the data, is kept outside of this model. The relationships towards other models can be found in Chap. 3.

Compared to previous approaches, we are only referencing the physical domain specific metadata, instead of incorporating them inside the *Time Series Cube DM*. The detailed description of the *Time Series Cube DM* can be found in Chap. 4

The position of *Time Series Cube DM* inside IVOA Architecture (Christophe Arviset (2010)) can be seen on Fig. 1.




Figure 1: IVOA architecture

### 3 Dependent data models

In this chapter we describe *Time Series Cube DM* position and relationships to other models.

On Fig. 2 we can see other models that we import without modifications, i.e. copy them inside our model, and without extending them.

On Fig. 3 we show the *Time Series Cube DM* relationships to other data cube models.




Figure 2: This UML describes models that we import AS-IS, without any modifications to them or extending them.




Figure 3: UML model specifying the position of the *Time Series Cube DM* compared to the *Cube DM* and the *ImageCube DM*.

### 3.1 Imported Data models

We import these data models without modification. We basically copy the model definition into *Time Series Cube DM*, making us completely dependent on what is defined within these - *Time Series Cube DM* wouldn't work without them.

However, because we use them as black boxes, it does not matter if these data models change - *Time Series Cube DM* has no common parts shared with them.

### 3.1.1 Dataset DM

The *Dataset DM* is used for describing common information for the whole set of time series cubes. The individual cubes, i.e. elements of the dataset collection described by *Dataset DM*, are then described within the *Time Series Cube DM*.

### 3.1.2 ObsCore DM

*ObsCore DM* in combination with TAP (Mireille Louys (2011)) is used for discovery of the time series cube metadata and we use it without modification. The *ObsCore DM* has already support for time series data with `dataProductType=timeseries` parameter. We are now working with a reference implementation, using GAVO DaCHS for server side and modified Splat-VO (Petr Skoda (2014)) for consumption of this ObsCore TAP service and discovering time series. More about discovery can be found in Chap. 5.

### 3.1.3 VO-DML

The *Time Series Cube DM* uses the VO-DML (Gerard Lemson (2016)) for annotating parts of the model and referencing individual axis domain models from inside the time series cube. The serialization of *Time Series Cube DM* model into VO-DML is not the purpose of this document.

We use VO-DML because it is the only reasonable way how to reference an actual model, not only entities already defined by that model. This is useful for referencing axis domains from within the axis definitions inside the *Time Series Cube DM*.

We are not importing the whole VO-DML, the only purpose it serves us is brining a simple way to reference other models as a whole, not only their structures.

## 3.2 Extended Data models

The following models are either base classes or siblings with relationships to the *Time Series Cube DM*. We reuse parts of them or reference them and in some cases, we override what has been defined in them if it was not generic enough for the purposes of *Time Series Cube DM*.

### 3.2.1 Cube DM

The *Cube DM* describes generic data cubes and is defined in the *N-Dimensional Cube/Image Model* (et al. (2015)) document. This document is also describing the *Sparse Cube DM* and the *Pixelated Cube DM*, however, it is not describing *Image Cube DM* (as could be erroneously understood from the title). We will refer to this standard as *N-Dimensional Cube DM*.



### 3.2.2 Pixelated Cube DM

This data model describes an abstraction of pixelated data cubes, e.g., image cubes. Essentially this is still the same data cube as the generic one, we only put restrictions on the main axes. The pixelated cube must have pixelated axes, i.e., regular intervals between coordinates.

### 3.2.3 Image DM

The *Image Cube DM* is currently being worked upon in the *IVOA Image DM* (Douglas Tody (2013)) working draft but is not being aligned completely with the *Cube DM* effort. It is an extension of *Pixelated Cube DM* and the image axes have naturally regular intervals between their coordinates.

### 3.2.4 Sparse Cube DM

This data model extends the *Cube DM* while adding restrictions on the main axes. The main axes must hold sparse data, i.e., we can have irregular intervals between the coordinates.

### 3.2.5 Time Series Cube DM

The *Time Series Cube DM* is describing time series data cubes, which are by definition sparse. We describe them in more detail in Chap. 4.

## 4 Time Series Cube DM

In this chapter we work with the time series data cubes defined within the *Time Series Cube DM*. These data cubes are instances of the **Time Series Cube** class seen on Fig. 4.

Note that the purpose of the *Time Series Cube DM* metadata is to enlist all the axes of the cube and provide linkage towards their values and errors, not to completely describe axis domains.

The axis domain descriptions will be held within other specialized data models, e.g., spatial and time domain in STC (Rots (2007)), spectral domain in Photometry DM (Jesus Salgado (2013)), etc. These external models will only be referenced by the *Time Series Cube DM*. This is where we differ from the original *Cube DM* that already focuses on supporting SIAv2 (Patrick Dowler (2015a)) image-based protocol and we do not want to restrict ourselves only to pixelated types of data cubes.

Withdrawing axis domain definitions as a dependency from the *Time Series Cube DM* helps us in the future when new data models will be created, e.g., for probability distributions or gravitational wave domains. If such a new data model appears, we can just add it as a reference to axis metadata

within the data cube and the *Time Series Cube DM* does not need to be changed at all. It also solves the main drawback of previous approaches and that is extensibility - SSAP or SimpleTimeSeries standards would need to be changed if a new domain was to be supported. Moreover, such a domain would not be logically part of spectral or time aspects of those models and therefore it does not belong there.

The only difference between axis coordinate and the value for that coordinate, from the data point of view, is that the axis coordinates are independent while the data values are dependent on these coordinates.

Therefore, we can afford (from the data point of view) to treat them equally, which is one of the biggest advantages of a data cube. We can browse the data cube by filtering the axis coordinates and the data values in the same way. This enables a client for example to slice the time axis of a light curve (coordinates) within an interval and the magnitude axis (values) within another interval using the same algorithm.

The time series data cube serialized accordingly to the *Time Series Cube DM* describes only the data, not the physical interpretation. From the IT perspective, it holds *data*, not *information* or *knowledge*. We are building on the fact that any information that we would like to store within the time series data cube, even the provenance, can be represented by some value and can be treated as an axis from the data point of view.

The **Time Series Cube** class as seen on Fig. 4 is the main class of *Time Series Cube DM*. It extends the **Sparse Cube** class taken from *N-Dimensional Cube DM*. This means that we are reusing what was already defined within **Sparse Cube** class coming from *Sparse Cube DM* and add new parameters in the **Time Series Cube** class.

The UML representation of our model can be seen on Fig. 4. On this figure we describe two data models, i.e., *Dataset DM* and *Cube DM*. The *Axis Domain DMs* package is collecting external models that we reference. *Time Series Cube DM* imports classes from the *DataSet DM* model.

A time series data cube can be distinguished from a more generic sparse data cube by having the `dataProductType` attribute with value "timeseries".

The *DataSet DM*, *Cube DM* and *Axis Domain* packages seen on Fig. 4 are described below.

## 4.1 DataSet DM

Everything you see in the *Dataset DM* package on Fig. 4 is imported, i.e., not changed within the scope of this document.

The *DataSet DM* holds common metadata for all **Time Series Cube** class instances collected within this dataset and stores them in the **ObsDataset** class. Typically that would be **DataID** class attributes like `creator`,




Figure 4: The UML diagram of Time Series Cube DM ecosystem. The *DataSet DM* classes (red) collect sparse cubes and time series cubes (cyan). The *Axis Domain DMs* referenced by the *CubeAxis* class are split between already existing (green) and potentially coming (grey) models.

Curation class metadata like `publisherDID` or `releaseDate` and Provenance class metadata like `facility` or `instrument` identifying the origin of the data. A `Target` class can also have its metadata specified, if the dataset has a common observation target.

The *SparseCube Dataset* class extends *ObsDataset* class and puts restriction on what can be stored within this collection. An instance of the *SparseCube Dataset* class has an attribute `sparseCubeData` which holds the collection of *SparseCube* class instances, i.e. `Time Series Cube` instances.

## 4.2 Cube DM

In this section we speak about relevant classes from the *Cube DM* package, i.e., `Sparse Cube` class, `Time Series Cube` class and `CubeAxis` class as described on Fig. 4.

### 4.2.1 Sparse Cube class

On Fig. 4 we show the `Sparse Cube` class that we are re-using from *N-Dimensional Cube DM*. The `Time Series Cube` class is extended from this one.

### 4.2.2 Time Series Cube class

A client can identify the `Time Series Cube` instance by having the keyword `timeseries` in the `dataProductType` attribute of its instance. By defining *CubeAxis* class and referencing the axis domain models from it, we override the *Cube DM* dependency on axis domains, e.g., *STC v2.0 DM*.

A client reading the data cube might want to know what should it treat as independent (coordinates) and dependent (data) values - they are collected within the `independent_axes` and in `dependent_axes` attributes of the `Time Series Cube` instance.

### 4.2.3 CubeAxis class

The `CubeAxis` class annotates the most important information about the axis. Where do I find the data (`field` attribute) and its simple statistical error (`error` attribute). If a specialized client wants to find more about specific metadata and data of an axis (its frame, domain, all related parameters and fields), it can go to the axis domain model through the `model` parameter of that axis and find the information in there.

And because this client is specialized towards that domain, it understands metadata in the referenced domain-specific data model (e.g., gravitational wave domain, spectral domain, etc.) and we don't need to make it part of *Time Series Cube DM*.

### 4.3 Axis Domain DMs

An answer to what all can be described with an axis (the axis domain) is not within the scope of this document. If we want to describe our spectral coordinates (e.g., list of filters) by the *Photometry DM* or in *STC DM* or add a completely new model for my axis domain (e.g., Gravitational Wave data), it does not affect the *Time Series Cube DM*.

### 4.4 Time Series Cube DM summary

With the growing amount of models and their diversity within the IVOA, keeping models independent of each other is the only viable approach to their evolution and maintenance. By only referencing a simple abstraction of an external model instead of importing every detail of that model, the *Time Series Cube DM* does not need to change every time an *Axis Domain Model* changes or a new one is created - the abstraction of that domain will be still the same.

## 5 Discovery - ObsCore

The ObsCore protocol is the natural option for discovery of our time series data cubes. In ObsCore 1.1, the most crucial parameters are added for the description of axis lengths for the basic types of axes, i.e., `s_xel` for spatial axis length, `t_xel` for temporal, etc., so the client knows what is the size of the cube it is about to download. The usual implementation of ObsCore model discovery is using TAP protocol for access to the tables holding the time series data.

At current point of time, the ObsCore suffices our requirements for discovery and anyone can add new parameters to his ObsCore service - except the *DataLink* (Patrick Dowler (2015b)) metadata discovery. This point is under discussion and, in the end, it may will not demand change to the ObsCore TAP standard.

## 6 Serialization

We decided to discuss the VOTable format as the first one to support as it is the easiest one to form discussions upon and afterwards close them with a deterministic XSD that can validate whether the serialization is strictly following the standard or not. This section focuses on serialization of the metadata and data defined within the *Time Series Cube DM*, not on serialization of the data model structure - that will be handled by VO-DML.

GROUP XML element can be used as an equivalent of the data model classes described in Chap. 4. These are linked to the IVOA data

models by annotations. The *Time Series Cube DM* itself then contains `independent_axes` and `dependent_axes` `GROUP` elements which hold the actual information about the axis.

The metadata, i.e., common information for every point in that axis, is referenced in the group as `PARAMRef` elements and the data (actual measurements that have usually different values for every point of this axis) is referenced by `FIELDRef` elements.

These `FIELDRef` elements point to `FIELD` elements, holding data of the `VOTable`. The order of these is the same as the order of data values stored in the standard `DATA` element of the `VOTable` (here we can find `TR` and `TD` elements of the `VOTable`).

Axis Models are again serialized as `GROUP` elements and are referenced from the independent or dependent axis collections by the model using `GROUPRef`.

The basic example of how such a serialization would look like can be seen on Fig. 5 and more can be found in the Chap. 7.

## 7 Supported science cases

In this chapter we describe how the most typical time series use cases are supported. We show that our approach can support all of the usual use cases, such as simple light curves or groups of light curves, in the same way as completely generic time series data like provenance linkage or frequency-based analysis.

We focus mainly on the support of time series data other than light curves. We will describe the science cases collected and suggest a `VOTable` serialization example for each one of these examples, so we can see that the *Time Series Cube DM* is easily serializable.

### 7.1 Requirements

The formalization of basic time series use cases comes from the *CSPTimeseries* science cases document (Solano (2012)). In addition to these use cases, we collected a few others from LSST or LIGO requirements for time series data models. In this chapter, we take these science cases one by one and explain to implement them with *Time Series Cube DM*.

We took the first three logical groups of light curves from the *CSPTimeseries* document (Solano (2012)) and added two more, the complex light curves and gravitational waves, as examples of what can be accomplished with the model.

## 7.2 Science case definitions

The science case definitions are defined from the requirements and are as follows.

- Simple Light curves - Combine photometry and light curves of a given object/list of objects in the same photometric band. We refer to the Group A in the *CSPTimeseries* document mentioned above as *simple light curves*.
- Groups of light curves - Combine photometry and light curves of a given object/list of objects in different photometric bands. We refer to the Group B in the *CSPTimeseries* document and slightly generalize it as *groups of light curves*, i.e., not only collecting light curves for filters, but also holding other types of axes, e.g., polarization.
- Generic Time series - Time series other than light curves. Basically anything what has at least one time-based axis can be put here. We refer to the Group C in the *CSPTimeseries* document as *generic time series*.
- Complex light curves - Special requirements on light curve data:
  - Provenance tracking: These use cases require to track original image or its cutouts for every point of the light curve, e.g., link to SIAP cutouts of the relevant region in the original image which can be opened in a graphical client. We are able to transfer the data and describe the metadata for this use case.
  - Probability distribution: Another use case computes and transfers probability distribution functions for the time series, e.g., computed simple Gaussian distribution parameters for the coordinate errors from which the light curve was generated). We are able to transfer the data but for metadata description, there would have to be created a new kind of Probabilistic DM as seen on Fig. 4. However, when such model comes, we just add the `model` reference to the *CubeAxis* metadata, without any changes to the *Time Series Cube DM* serialization.
  - *DataLink* for time series: A basic use case for *Time Series Cube DM* combined with *DataLink* would be storing data linkage for the whole time series instance, e.g., *DataLink* to the periodogram of this time series. Another one would be filtering of light curve points on a supported criteria, e.g., filtering only high quality points of a light curve. Here we are struggling with the metadata transfer for the ObsCore TAP discovery described in Chap. 5. The problem is that with *ObsCore DM* we can get also other data

than time series and the *DataLink* does not support metadata restriction based on a `dataProductType`. We are able to construct such a *DataLink* for the light curves, the only problem is with transferring the metadata about such a domain over ObsCore TAP service.

- Power spectra, Gravitational waves: Special use case for generic time series data - we show here how to store not only simple time axis but also its derivatives, e.g., frequency. For Gravitational waves, this would be tracking for example ASD (strain per rtHz) against Frequency (Hz) as seen on the Gravitational wave tutorial (*Signal Processing with GW150914* (n.d.)).

### 7.3 Simple Light curves

This is the basic use case supported up to now with *SimpleTimeSeries* (Matthew Graham (2014)). Now, we can easily support it in the *Time Series Cube DM*. An example of the *DataSet* metadata can be seen on Fig. 5, the definition of the cube on Fig. 6 and the first row on Fig. 7.

The light curve rendered within SPLAT-VO client ((Petr Skoda (2014))) can be seen on Fig. 8 and 10. These images were taken on real data coming within our reference implementation for server and client. More about the reference implementation can be found in Chap. 8.

#### 7.3.1 Examples

The examples of simple light curves can be seen on Fig. 5, Fig. 6 and Fig. 7. The plot rendered within SPLAT-VO can be seen on Fig. 8.

```
<?xml version="1.0" encoding="utf-8"?>
<VOTABLE version="1.2" xmlns="http://www.ivoa.net/xml/VOTable/v1.2" xmlns:spec="http://www.ivoa.net/xml/SpectrumModel/v1.01" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.ivoa.net/xml/VOTable/v1.2 http://vo.ari.uni-heidelberg.de/docs/schemata/VOTable-1.2.xsd">
  <DESCRIPTION>Identified objects on DK-154 surveys</DESCRIPTION>
  <RESOURCE>
    <TABLE name="sInstance">
      <GROUP vodml="ds:sparseCubeDataset">
        <PARAMref id="dataProductType" vodml="ProductType"/> <!-- the value would be timeseries for whole dataset -->
        <PARAMref id="calibLevel" vodml="CalibrationLevelType"/>
      </GROUP>
      <GROUP vodml="ds:DataID">
        <PARAMref id="date" vodml="date"/> <!-- ID may be held in the param itself -->
        <PARAMref id="creator" vodml="creator"/>
        <PARAMref id="observationID" vodml="observationID"/>
      </GROUP>
    </TABLE>
  </RESOURCE>
</VOTABLE>
```

Figure 5: Dataset metadata serialization. Here we can see dataset `GROUP` element. There is a dataset (collection) of sparse cubes. The dataset contains the metadata, e.g., the `DataID`. All of the relevant elements (`GROUP`, `PARAM`, `FIELD`, etc.) are annotated against their respective *Time Series Cube DM* counterparts using the `vodml` attribute.

### 7.4 Groups of light curves

The use case of storing light curves for one object in multiple photometric bands can be supported simply by addition of one independent axis - the



```

<GROUP id="timeseries" vodml="ndcube:TimeSeriesCube">
  <GROUP vodml="ndcube:TimeSeriesCube"> <!-- this is the cube itself -->
    <PARAMref vodml="ProductType"/> <!-- the value would be timeseries for whole dataset -->
    <GROUP id="independent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="dateTimeAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="HJD" id="field"/>
        <GROUPref ref="datestc" id="model" vodml="VOOML Model"/> <!-- ref to group holding STC2 model -->
      </GROUP>
      <GROUP name="spatialAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="raj2000" id="field"/>
        <GROUPref ref="spatialFrame" id="model" vodml="VOOML Model"/> <!-- ref to group holding STC1.3 model -->
      </GROUP>
      <GROUP name="spatialAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="dej2000" id="field"/>
        <GROUPref ref="spatialFrame" id="model" vodml="VOOML Model"/> <!-- ref to group holding STC1.3 model -->
      </GROUP>
    </GROUP>
    <GROUP id="dependent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="fluxAxis" vodml="CubeAxis"> <!-- the name could be also ID here I guess -->
        <FIELDref ref="FLX" id="field"/>
        <FIELDref ref="FLXERR" id="error"/>
        <GROUPref id="model" vodml="VOOML Model"/> <!-- ref to flux model (if existing), otherwise just a GROUP holding together metadata about this axis -->
      </GROUP>
      <GROUP name="magnitudeAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="MAG" id="field"/>
        <FIELDref ref="MAGERR" id="error"/>
        <GROUPref id="model" vodml="VOOML Model"/> <!-- ref to magnitude model (if existing), otherwise just a GROUP holding together metadata about this axis -->
      </GROUP>
    </GROUP>
  </GROUP>
</GROUP>

```

*Figure 6:* Time series data cube metadata serialization. Here we can see the metadata **GROUP** element for our time series data cube. The independent axes like temporal and spatial are referencing the **FIELD** elements holding their data and annotating against their axis models holding their metadata, using the **GROUPref** elements, e.g., **datestc** or **spatialFrame**. The same goes for dependent axes.


```

<FIELD ID="HJD" datatype="double" name="HJD" ucd="VOX:Image_MJDateObs" unit="d" utype="spec:Spectrum.Data.SpectralAxis.Value">
  <DESCRIPTION>Epoch at midpoint of observation in Heliocentric Modified Julian Date</DESCRIPTION>
</FIELD>
<FIELD ID="raj2000" datatype="double" name="raj2000" ucd="pos.eq.ra;meta.main" unit="deg">
  <DESCRIPTION>Observed RA of the object</DESCRIPTION>
</FIELD>
<FIELD ID="dej2000" datatype="double" name="dej2000" ucd="pos.eq.dec;meta.main" unit="deg">
  <DESCRIPTION>Main value of declination</DESCRIPTION>
</FIELD>
<FIELD ID="FLX" datatype="float" name="FLX" ucd="phot.flux">
  <DESCRIPTION>Photon Flux</DESCRIPTION>
</FIELD>
<FIELD ID="FLXERR" datatype="float" name="FLXERR" ucd="stat.error;phot.flux">
<FIELD ID="MAG" datatype="float" name="MAG" ucd="phot.mag" unit="mag">
  <DESCRIPTION>Magnitude of the object</DESCRIPTION>
</FIELD>
<FIELD ID="MAGERR" datatype="float" name="MAGERR" ucd="stat.error;phot.mag" unit="mag">
</FIELD>
<DATA: <!-- the order of data cols are always the same as order of FIELD elements defined above. -->
  <TABLEDATA>
    <TR>
      <TD>56287.11921047</TD>
      <TD>12.9768538030848</TD>
      <TD>-72.9580158132637</TD>
      <TD>7.85796e-17</TD>
      <TD>0.8</TD>
      <TD>19.0829</TD>
      <TD>0.8</TD>
    </TR>
  </TABLEDATA>
</DATA>

```

*Figure 7:* Time series data cube data serialization. Here we can see the **FIELD** elements referenced from the time series cube metadata. This is a basic light curve holding the heliocentric Julian date for temporal axis, equatorial coordinates and two dependent axes with measurements - flux and magnitude.

spectral axis. Having the light curve points in multiple bands, e.g., UBVRI, is nothing more than having one more discrete spectral axis with coordinates that can have 5 different values. If the spectral axis is described in continuous intervals, the error **FIELD** would be meaningful, for just a list of filter names we are storing only data value elements, i.e. **field** attribute of **Cube Axis**, with textual information.



*Figure 8:* Light curve within SPLAT-VO. The light curve is rendered from VOTable coming within time series data cube format. Note that SPLAT-VO is able to distinguish the magnitude and inverts the axis, as well as the units for the temporal axis. These are coming from the axis domain metadata.

It is evident that adding more such "groups" (other independent axes) will produce a lot of duplicated values in the `DATA` element of our VOTable. However, this is just a question of optimization.

If the need arises, for start, compressing algorithms such as EXI (*Efficient XML Interchange (EXI) Format, Version 1.0 (2014)*) can be used for immediate value of eliminating the duplicities. A cleaner approach would be a table normalization during the transfer - the independent axes (coordinates) would be stored in separate `TABLE` elements (the same goes for FITS) that would be referenced by the dependent axes records.

The optimization of serialization is not the main aim of this document but can be expanded in the future if deemed necessary.

#### 7.4.1 Examples

The XML example for groups of light curves can be found on Fig. 9. The plot rendered within SPLAT-VO can be seen on Fig. 10.

### 7.5 Generic time series

The Generic time series are the main motivation for using *Time Series Cube DM*, as any kind of data that can be structured into series can be represented

```

<GROUP id="timeseries" vodml="ndcube:TimeSeriesCube">
  <PARAM vodml="ndcube:TimeSeriesCube" <!-- this is the cube itself -->
  <PARAM-ref vodml="ProductType" <!-- the value would be timeseries for whole dataset -->
  <GROUP id="independent_axes" vodml="ndcube:CubeAxis">
    <GROUP name="dateTimeAxis" vodml="ndcube:CubeAxis">
      <FIELDref ref="HJD" id="field"/>
      <GROUPref ref="dateutc" id="model" vodml="VODML Model" <!-- ref to group holding STC2 model -->
    </GROUP>
    <GROUP name="spectralAxis" vodml="ndcube:CubeAxis">
      <FIELDref id="value"/>
      <GROUPref id="model" vodml="VODML Model" <!-- this is a reference to a GROUP representing the Photometry DM. -->
    </GROUP>
  </GROUP>
  <GROUP id="dependent_axes" vodml="ndcube:CubeAxis">
    <GROUP name="magnitudeAxis" vodml="ndcube:CubeAxis">
      <FIELDref ref="MAG" id="field"/>
      <FIELDref ref="MAGERR" id="error"/>
      <GROUPref id="model" vodml="VODML Model" <!-- ref to magnitude model (if existing), otherwise just a GROUP holding together metadata about this axis -->
    </GROUP>
  </GROUP>
</GROUP>
</GROUP>

```

Figure 9: Groups of light curves time series data cube metadata serialization. Here we have two independent axes, the temporal and spectral. This means that the dependent axis (actual measurements, magnitude in this case) can hold values for any combination of the date and spectral axis values. The *Photometry DM* as an axis domain definition is referenced from the `model` attribute of the `spectralAxis` GROUP element.




Figure 10: Group of light curves within SPLAT-VO. As we are still using a valid VOTable serialization, our SPLAT-VO client can render the time series data cube for a light curve taken in 4 filters without modification. The spectral axis is the 3rd rendered axis and is expressed in colors instead of axis coordinates. The axis labels are taken from axis domain metadata referenced by the *Time Series Cube DM*

by a data cube. In here we also show that we can describe our metadata by using `PARAM` elements which our client can understand by parsing UCDS. Or,

we can describe those by using external data models that the *Time Series Cube DM* only references.

The fact that the *Time Series Cube DM* will not know about the details of every domain metadata is crucial as it enables us to describe the metadata with our own parameters if a standard model is not yet available, without the need of including them in *Time Series Cube DM*. After the model is standardized, e.g., *STC v2.0*, *Provenance DM*, *Probability DM*, etc., we just group the needed metadata under a `GROUP` element and reference it in our `Cube axis` description, without the need to change the structure of the data model or even the serialization within the rest of the *VOTable*.

The axes derived from the basic types can be still tracked as the same type. The frequency for a power spectrum is still a time series in its own sense and we can easily support it within *Time Series Cube DM*. It is true that the frequency is not technically a sparse axis that we chose as a basis for the time series data cubes, but logically a power spectrum or periodogram can be understood as time series. Moreover, the *Sparse Cube DM* places a restriction on *Time Series Cube DM* to support sparse axes, not to restrict itself only to them.

### 7.5.1 Examples

A basic example of a generic time series can be seen on Fig. 11. Focus on the highlighted model attribute of the hardness ratio dependent axis. Even if the axis domain model does not exist yet, we still want to describe the structure of the data (values and errors of this axis) and store the rest in `PARAM` elements, which are not part of the data-perspective view provided by *Time Series Cube DM*. If the axis domain model for describing the hardness ratio existed, we would reference it by the `model` attribute. This reference tells the client which `PARAM` elements are relevant for the domain of this axis.

## 7.6 Complex light curves

This group is already describing science cases gathered apart the ones described in *CSPTimeseries* document. We have collected these requirements from the most organizations most pressed by the needs for a standard, easily extendable solution.

### 7.6.1 Examples

One of the most important things when processing time series data is to be able to link the points of a light curve to the original image it was coming from. An outlier on the light curve could be a static noise or some other kind of error on the image but it could be also a valid burst in the object's magnitude. We can identify this on the first look if we have the link to the

```

<GROUP id="timeseries" vodml="ndcube:TimeSeriesCube">
  <GROUP vodml="ndcube:TimeSeriesCube"> <!-- this is the cube itself -->
    <PARAMref vodml="ProductType"/> <!-- the value would be timeseries for whole dataset -->
    <GROUP id="independent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="dateTimeAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="HJD" id="field"/>
        <GROUPref ref="datestc" id="model" vodml="VODML Model"/> <!-- ref to group holding STC2 model -->
      </GROUP>
    </GROUP>
    <GROUP id="dependent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="hardnessRatioAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="ratio" id="field"/>
        <FIELDref ref="ratio_err" id="field"/>
        <GROUPref id="model" vodml="VODML Model"/> <!-- A reference to the GROUP describing Hardness ratio -->
      </GROUP>
    </GROUP>
  </GROUP>
</GROUP>
<DATA>
  <TABLEDATA>
    <TR>
      <TD>56287.11921047</TD>
      <TD>1.0</TD>
    </TR>
  </TABLEDATA>
</DATA>

```

Figure 11: Hardness ratio time series example. We omit here the actual FIELD elements for simplicity. These would contain most importantly UCDS for the fields and other metadata.

original data for each point of the light curve. The sample serialization of such provenance tracking can be seen on Fig. 12.

The implementation on server side is easy and we plan this to be a very simple operation on the client side too. Just one click on the point of the light curve can send the coordinates via SAMP to Aladin client which displays the original image. The light curve plotted within TopCat client can be seen on Fig. 13. With the metadata *Time Series Cube DM* provides, Topcat could automatically identify what to do when user clicks on a point of the light curve. In our case, we want to take the URL of the original image from which this point of the curve originates and send it to Aladin via SAMP, see Fig. 14. We are also able to send just a cutout instead of the whole image link, as seen on Fig. 15.

## 7.7 Gravitational waves

This example is quite similar to the one described in Chap. 7.5. The only difference is a temporal axis describing a frequency (as time derivative) and then another dependent axis that contains measured data for gravitational wave analysis. This use case is taken from the LIGO tutorial site (*Signal Processing with GW150914* (n.d.)).

In our specific uses case, we describe a time series data cube storing gravitational wave analysis, as can be seen on Fig. 17. We have a temporal axis of **frequency** rendered as **x axis** and the **strain/rHz** dependent axis rendered as **y axis**. The second independent axis holding sampling rate frequency is rendered as 3rd dimension in colors.

```

<GROUP id="timeseries" vodml="ndcube:TimeSeriesCube">
  <GROUP vodml="ndcube:TimeSeriesCube"> <!-- this is the cube itself -->
    <PARAMref vodml="ProductType"/> <!-- the value would be timeseries for whole dataset -->
    <GROUP id="independent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="dateTimeAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="HJD" id="field"/>
        <GROUPref ref="datestc" id="model" vodml="VOOML Model"/> <!-- ref to group holding STC2 model -->
      </GROUP>
      <GROUP id="dependent_axes" vodml="ndcube:CubeAxis">
        <GROUP name="magnitudeAxis" vodml="ndcube:CubeAxis">
          <FIELDref ref="MAG" id="field"/>
          <FIELDref ref="MAGERR" id="error"/>
          <GROUPref id="model" vodml="VOOML Model"/> <!-- ref to magnitude model (if existing), otherwise just a GROUP holding together metadata about this axis -->
        </GROUP>
        <GROUP name="provenanceAxis" vodml="ndcube:CubeAxis">
          <FIELDref id="field"/>
          <GROUPref id="model" vodml="VOOML Model"/> <!-- this is a reference to a GROUP representing the Provenance DM. -->
        </GROUP>
      </GROUP>
    </GROUP>
  </GROUP>
</GROUP>
<DATA>
  <TABLEDATA>
    <TR>
      <TD>56287.11921047</TD>
      <TD>19.0829</TD>
      <TD>0.00000001</TD>
      <TD>http://vos2.asu.cas.cz/getproduct/dk154\_reduced/data/20121225/SHCF3\_000001\_sci.fits.fz?sdec=0.018amp;dec=19.03398amp;ra=-72.95801581326378amp;sra=0.01</TD>
    </TR>
  </TABLEDATA>
</DATA>

```

Figure 12: Linkage towards original image cutout for one point of the light curve. The data can be seen in the last column of the table with cutout parameters taken from the right ascension and declination of this point of the curve.




Figure 13: A light curve rendered in Topcat client. The graph represents magnitude based on the HJD when the star was observed. After clicking on the highlighted point, Fig. 14 or Fig. 15 follows.

### 7.7.1 Examples

The sample serialization and one row can be seen on Fig. 16.




Figure 14: Original image for one point of a light curve. After clicking on the point, the image from which it originates is displayed in Aladin.




Figure 15: Small cutout from original image for one point of a light curve. If a SIA cutout link is sent instead of the SIA link to the whole image, Aladin will render it the same way as the whole original image.

```

<GROUP id="timeseries" vodml="ndcube:TimeSeriesCube">
  <GROUP vodml="ndcube:TimeSeriesCube"> <!-- this is the cube itself -->
    <PARAMref vodml="ProductType"/> <!-- the value would be timeseries for whole dataset -->
    <GROUP id="independent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="frequencyAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="sampleFreq" id="field"/>
        <GROUPref id="model" vodml="VODML Model"/> <!-- ref to group holding the frequency domain metadata -->
      </GROUP>
      <GROUP name="sampleFrequencyAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="sampleFreq" id="field"/>
        <GROUPref id="model" vodml="VODML Model"/> <!-- ref to group holding the frequency domain metadata -->
      </GROUP>
    </GROUP>
    <GROUP id="dependent_axes" vodml="ndcube:CubeAxis">
      <GROUP name="provenanceAxis" vodml="ndcube:CubeAxis">
        <FIELDref ref="ASD" id="field"/>
        <GROUPref id="model" vodml="VODML Model"/> <!-- this is a reference to a GROUP representing model of the Strain of a gravitational wave. -->
      </GROUP>
    </GROUP>
  </GROUP>
</GROUP>
<DATA>
  <TABLEDATA>
    <TR>
      <TD>100</TD>
      <TD>4096</TD>
      <TD>1.05631e-23</TD>
    </TR>
  </TABLEDATA>
</DATA>

```

Figure 16: VOTable example for gravitational Wave data. We omit FIELD elements for simplicity, the important thing is that all Gravitational wave specific axes are referencing their domain metadata in external models, so *Time Series Cube DM* does not need to know about them. The sample row is showing a strain for 100 Hz frequency sampled with 4096 Hz.




Figure 17: LIGO time series plotting example. This graph is showing strain for frequency between  $10^1$  Hz and  $10^4$  Hz for two different samplings of 4096 Hz and 16384 Hz which are expressed in colors.



## 8 Reference implementation

We are working on GAVO DaCHS implementation for server-side service providing time series data cubes in the *Time Series Cube DM* format. The service is ready and providing real data, but we have only a development environment ready right now and is under constant change. If you want to have a look on the final version, you can find it on the following web page (Nadvornik (2017)).

We are also working on the SPLAT-VO modifications so it can understand the *Time Series Cube DM* format. Since we are using standard VOTables for serialization of our data cubes, the modifications are more related to SPLAT-VO understanding the axis domain models like magnitude axis inversion or time axis unit identification. In the current version of the serialization, SPLAT-VO understands the data perfectly.

## 9 Summary

This is still a work in progress and will need a lot of discussion and polishing to make it to an IVOA recommendation. However, we believe that this document creates a solid foundation upon which we can continue building.

## References

Christophe Arviset, S. G. e. a. (2010), ‘IVOA architecture’, IVOA Note.

**URL:** <http://www.ivoa.net/documents/Notes/IVOAArchitecture>

Doug Tody, M. D. e. a. (2012), ‘Simple spectral access protocol version 1.1’, IVOA Recommendation.

**URL:** <http://www.ivoa.net/documents/SSA/20120210/REC-SSA-1.1-20120210.htm>

Douglas Tody, F. B. e. a. (2013), ‘Ivoa image data model, version 1.0’, IVOA Working Draft.

**URL:** <http://wiki.ivoa.net/internal/IVOA/ImageDM/WD-ImageDM-20130812.pdf>

*Efficient XML Interchange (EXI) Format, Version 1.0* (2014), W3C Recommendation.

**URL:** <https://www.w3.org/TR/exi/>

et al., D. T. F. B. (2015), ‘Ivoa n-dimensional cube model, version 1.0’, IVOA Working Draft.

**URL:** <http://www.ivoa.net/documents/NDimCubeDM/20150320/WD-CubeDM-1.0-20150320.pdf>

- Gerard Lemson, O. L. e. a. (2016), ‘Vo-dml - a consistent modeling language for ivoa data models’, IVOA Proposed Recommendation.  
**URL:** <http://www.ivoa.net/documents/VODML/20160923/VO-DML-PR-v1.0.pdf>
- Jesus Salgado, C. R. e. a. (2013), ‘Ivoa photometry data model’, IVOA Recommendation.  
**URL:** <http://www.ivoa.net/documents/PHOTDM/20131005/REC-PhotDM-1.0-20131005.pdf>
- Kimball, R. (1997a), ‘A dimensional modeling manifesto’, Article.  
**URL:** <http://www.kimballgroup.com/data-warehouse-business-intelligence-resources/kimball-techniques/dimensional-modeling-techniques/star-schema-olap-cube/>
- Kimball, R. (1997b), ‘A dimensional modeling manifesto’, Article.  
**URL:** <http://www.kimballgroup.com/1997/08/a-dimensional-modeling-manifesto/>
- Matthew Graham, E. S. e. a. (2014), ‘Employing simpletimeseries for representing time series’, IVOA Note.  
**URL:** <http://www.ivoa.net/documents/Notes/SimpleTimeSeries/20140513/NOTE-SimpleTimeSeries-1.0-20140513.html>
- Mireille Louys, D. T. e. a. (2011), ‘Observation data model core components, its implementation in the Table Access Protocol, version 1.1’, IVOA Proposed Recommendation.  
**URL:** <http://www.ivoa.net/documents/ObsCore/20111028/REC-ObsCore-v1.0-20111028.pdf>
- Nadvornik, J. (2017), ‘Aiascr vo services’, Web page.  
**URL:** <http://vos2.asu.cas.cz/>
- Patrick Dowler, D. T. e. a. (2015a), ‘Ivoa simple image access, version 2.0’, IVOA Recommendation.  
**URL:** <http://www.ivoa.net/documents/SIA/20151223/REC-SIA-2.0-20151223.pdf>
- Patrick Dowler, F. B. e. a. (2015b), ‘IVOA datalink’, IVOA Recommendation 17 June 2015.  
**URL:** <http://www.ivoa.net/documents/DataLink/>
- Petr Skoda, P. W. D. e. a. (2014), ‘Spectroscopic Analysis in the Virtual Observatory Environment with SPLAT-VO’, *ArXiv e-prints*, arXiv:1407.1765.  
**URL:** <https://arxiv.org/abs/1407.1765v1>

Rots, A. (2007), 'Space-time coordinate metadata for the virtual observatory, version 1.3', IVOA Recommendation.

**URL:** <http://www.ivoa.net/documents/latest/STC.html>

*Signal Processing with GW150914* (n.d.), IPython Notebook.

**URL:** [https://losc.ligo.org/s/events/GW150914/GW150914\\_tutorial.html](https://losc.ligo.org/s/events/GW150914/GW150914_tutorial.html)

Solano, E. (2012), 'Science use cases for time series', Science use Cases for Time Series.

**URL:** <http://wiki.ivoa.net/twiki/bin/view/IVOA/CSPTimeSeries>