Query driven visualization of astronomical catalogs

Hugo Buddelmeijer · Edwin A. Valentijn

Received: 8 August 2011 / Accepted: 6 October 2011 / Published online: 3 December 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract Interactive visualization of astronomical catalogs requires novel techniques due to the huge volumes and complex structure of the data produced by existing and upcoming astronomical surveys. The creation as well as the disclosure of the catalogs can be handled by data pulling mechanisms (Buddelmeijer et al. 2011). These prevent unnecessary processing and facilitate data sharing by having users request the desired end products. In this work we present query driven visualization as a logical continuation of data pulling. Scientists can request catalogs in a declarative way and set process parameters directly from within the visualization. This results in profound interoperability between software with a high level of abstraction. New messages for the Simple Application Messaging Protocol are proposed to achieve this abstraction. Support for these messages are implemented in the Astro-WISE information system and in a set of demonstrational applications.

Keywords Data mining · Visualization · Virtual Observatory · SAMP · Astro-WISE

1 Introduction

Large astronomical surveys require novel ways for handling the data they produce. For example, the ongoing KiDS and VIKING surveys will cover

H. Buddelmeijer (✉) · E. A. Valentijn
Kapteyn Astronomical Institute, Postbus 800,
9700 AV, Groningen, The Netherlands
e-mail: buddel@astro.rug.nl

E. A. Valentijn
e-mail: valentyn@astro.rug.nl

1,500 square degree in optical and infra-red wavelengths [1] and the upcoming Euclid mission will cover 20,000 square degree [5]. These surveys will detect billions of galaxies for which hundreds of parameters will be quantified, leading to terabytes of data to explore.

Data pulling mechanisms can be used to achieve the scalability to create catalogs ([3], hereafter Paper I). The essence of data pulling is that processing steps necessary to create a catalog are determined by specifying the required target catalog. The information system will determine how existing catalogs can be used to fulfill the request and will initiate the creation of new catalogs only when no suitable ones exist. This maximizes reusability of the catalogs and minimizes unnecessary calculations. This requires full *data lineage*, which means that catalogs are stored with all the information required to process them.

Query driven visualization is a methodology to explore large data sets by limiting the processing required for visualization to the subsets of the data deemed “interesting” as defined by the user [8]. Related work focuses on limiting the processing of the visualization itself [8], the fast identification and retrieval of data [2, 7] or on the data representation [6].

In this paper, we see query driven visualization as the logical continuation of data pulling in an information system with full data lineage. The main contributions of our work follow from applying this novel viewpoint to source catalogs: (1) We limit the processing required to create the requested catalog itself, instead of the processing required for the visualization. (2) We permit requests in a more declarative form than direct database queries would allow. (3) We allow the user to inspect and influence the processing from within the visualization by exporting the data lineage. (4) We achieve a high level of abstraction that allows close interoperability between software.

We demonstrate our techniques with our Astro-WISE implementation and by designing new messages for the Simple Application Messaging Protocol.

1.1 Data pulling and declarative querying

Data pulling is an excellent opportunity for query driven visualization. The autonomous discovery and creation of catalogs permits requests that are very declarative. A scientist can request parameters of sources without having knowledge of whether these parameters have already been calculated or not. This functionality can be implemented in external software and an example program to pull catalogs is given with the ‘Simple Puller’ of Section 3.3.

Compare this for example with an SQL-based system [4]: to formulate an SQL query it is required to know which tables contain the required parameters, how to identify the relevant rows and columns, and often how to join tables. This becomes a non-trivial problem when catalogs are shared between multiple users and the number of catalogs and their sizes grow. At a certain stage it becomes too time consuming and error-prone to find required data by hand, especially when it is unknown whether it exists at all.

1.2 Full data lineage and exploration

An information system with data pulling often has *persistent objects* with *full data lineage*: Each data set is represented as an object—in computer science terminology—that persists between sessions and users. In Astro-WISE these objects are called *process targets*. A process target contains all the information required to create the data it represents from other process targets, its *dependencies*. This is called *backward chaining* and links every data product back to the raw data.

The data lineage can be utilized in query driven visualization by having the visualization software request it. This allows the visualization software to show this information, either directly or processed in the visualization. An example of the former is given with the ‘Tree Viewer’ in Section 3.3.

Furthermore, exporting the data lineage makes it possible for scientists to influence the processing by permitting the visualization software to change processing parameters. An example of this is given with the ‘Object Viewer’ in Section 3.3.

1.3 Abstraction and interoperation through SAMP

Data pulling mechanisms are well suited for abstraction on different levels: firstly, pulling data does not require detailed knowledge of every processing step; secondly, these processing details themselves can be abstracted, because of the standardized data lineage.

Such an abstraction allows query driven visualization to be performed between any visualization package and information system. The thoroughness of the interoperation will depend on the level of abstraction supported by both applications. We extended the Simple Application Messaging Protocol¹ to facilitate such interoperation by designing new message types (Section 2).

1.4 Astro-WISE

Query driven visualization requires an information system responsible for creation, storage and delivery of the data. We choose to use Astro-WISE for this, although any information system with data pulling and persistent objects would be suitable, because of the abstraction through SAMP. In Section 6 we describe the details of our Astro-WISE SAMP implementation.

2 Interoperability through SAMP

The *Simple Application Messaging Protocol* (SAMP) is an International Virtual Observatory Alliance (IVOA) standard for interoperation between

¹<http://www.ivoa.net/Documents/latest/SAMP.html>

astronomical software. The idea behind it is akin to the UNIX-philosophy that tools should do one thing, should do that thing well and communicate with other programs for things they cannot do.

2.1 Simple application messaging protocol

We give a short description of SAMP before discussing our additions. For details we refer to Section 5 and to the official documentation.² The protocol uses a client-server model based on application defined messages. Clients register with the SAMP HUB and subscribe to certain types of messages. Clients can then send messages to individual clients or to any client that has registered for that kind of message. The receiving application should subsequently perform the action it has associated with the message. Lastly, the HUB will relay a response back to the sender if necessary.

The expected action that corresponds to a message is determined by the type of the message. Both default administrative messages and widely accepted application defined messages can be found on the SAMP wiki.³ In the rest of this section we first describe relevant existing messages and subsequently introduce our proposed messages. We list the type of the messages and a description of the intended action of the receiver. Details of the messages and their parameters are given Section 5.1.

2.2 Existing catalog related messages

Several existing catalog related messages can be used in conjunction with our new messages:

- `table.load.votable`: Load a table in VOTable format.
- `table.load.fits`: Load a table in FITS format.
- `table.highlight.row`: Highlight a single row of an identified table.
- `table.select.rowList`: Select a list of rows of an identified table.

Exactly what ‘highlighting’ or ‘selecting’ means, is left to the receiving application. Tables have three identifiers in SAMP: a table-id that is unique within the SAMP session, a URI where the catalog can be found and a human readable name. These identifiers are set with one of the `load` messages and used as a reference in the other messages. Rows are identified by their position in the table using zero-based indexing. Note that these messages can refer to any tabulated data set. In this paper we limit ourselves to source catalogs only.

²<http://www.ivoa.net/Documents/latest/SAMP.html>

³<http://www.ivoa.net/cgi-bin/twiki/bin/view/IVOA/SampMTypes>

2.3 Data pulling messages

We designed new SAMP messages to create a system independent way to perform pulling of catalog data. The messages should be sent from visualization software to the information system handling the data. The new message types start with `target.`; this is the name that the Astro-WISE information system uses to describe data objects that can be pulled:

- `target.catalog.pull`: Pull a catalog and send it over SAMP using one of the `table.load.*` messages. The result could be an existing catalog or a new catalog created by the pulling mechanisms. Any new data that is necessary to produce the required catalog is created automatically. This message requires the following parameters, detailed below: an identifier of a catalog to select the sources from, a selection criterion and a list of requested attributes of the sources.
- `target.catalog.derive`: Derive a catalog in the same fashion as with `target.catalog.pull`, but do not create any new data or send the catalog data over SAMP.

Support for the `.pull` message is the minimum required to request catalog data from the information system. The `.derive` message is useful when it is necessary to inspect or modify the derivation of the catalog—using the messages in Section 2.4—before visualization, for example to determine whether all required data is processed already or whether new data has to be created. These two messages require three parameters which we should elaborate on (see also Section 5.2):

- `catalog-id`: An identifier of the base catalog to select the sources from. It is left to the information system to inform scientists how to refer to a specific catalog. The `catalog-id` can be a unique identifier of an existing catalog, but could also be a reference to a catalog that does not yet exist, e.g. a photometric catalog for an observation that has not yet been reduced. It is also possible to designate identifiers for special catalogs, e.g. to denote the latest version of a catalog of an ongoing survey.
- `query`: A selection criterion to specify which sources of the original catalog are requested. This should be a logical expression referencing the `attributes` below. The exact specification of this expression is left to the information system. A logical choice would be the syntax of an ADQL⁴ `WHERE` clause (without the ‘`WHERE`’ itself).
- `attributes`: A list of requested attributes (parameters) of the sources. It is not required that the catalog corresponding to the `catalog-id` contains these attributes. The data pulling mechanisms of the information system should try to find the requested attributes in related catalogs and

⁴<http://www.ivoa.net/Documents/latest/ADQL.html>

should create new data sets if necessary. How an attribute should be specified, is left to the information system.

2.4 Object messages

Several SAMP message types are defined for interaction with an information system with persistent objects. These messages allow the visualization software to gain information about the objects and inspect or even influence its processing. The persistent object related messages are:

- `target.object.highlight`: Highlight an object.
- `target.object.info`: Return information about an object, see below.
- `target.object.change`: Change the value of a property of an object such as a process parameter or a dependency.
- `target.object.action`: Perform an action related to an object or property. Possible actions are retrieved using the `target.object.info` message.

The `target.object.highlight` message can be sent to any application, the others are supposed to be sent to the information system only.

A specific SAMP map is defined as a return value for the `target.object.info` message, containing information about the object and its properties (see Section 5 for details). For the object itself it includes information about what properties it has, its processing status and whether the object can be modified.

The properties of an object include process parameters and references to the dependencies of the object. The returned information about a property include its name, current value and optionally other values it can be set to. Furthermore the information system can define actions that can be performed on the object or its properties.

3 SAMP HUB and clients

The new SAMP messages are implemented in the Astro-WISE information system and demonstrated by a set of proof-of-concept applications. We first describe relevant existing SAMP applications, subsequently the Astro-WISE SAMP connectivity and end with the applications to demonstrate the new messages. Figure 1 shows a diagram of the interoperability between Astro-WISE and several SAMP applications.

3.1 Existing SAMP applications

We list existing SAMP applications that are relevant to catalog data.

- **SAMP HUB**: The *HUB* is the center of SAMP to which the other applications connect. The HUB can be a standalone program or can be integrated in one of the clients, e.g. Aladin and Topcat include one.

Fig. 1 The connectivity between Astro-WISE and SAMP. The SAMP HUB in the *center*, the Astro-WISE system on the *left* and other SAMP enabled applications on the *right*

- **Topcat:** *Topcat*⁵ is a table viewer/manipulator written in Java. The visualization power of Topcat lies in its interactivity. Selections performed in one window propagate to other windows and by the use of SAMP messages to other applications.
- **Aladin:** *Aladin*⁶ is an interactive software sky atlas allowing the user to visualize digitized astronomical images up to 50K by 50K pixels, superimpose entries from astronomical catalogues or databases, and interactively access related data and information online archives for all known sources in the field.

3.2 Astro-WISE and SAMP

Astro-WISE has SAMP connectivity in the interactive Python prompt and on the webservices.

- **awe-prompt:** The Astro-WISE awe-prompt is an interactive Python prompt that forms the primary user interface to Astro-WISE. We developed a module for SAMP connectivity in the awe-prompt and other Python applications. All messages from Section 2 are supported.
- **DBViewer:** With the Astro-WISE *DBViewer* one can view all content of the database and can send query results over SAMP. The DBViewer is beyond the scope of this paper.

3.3 Query driven visualization prototype

A set of proof-of-concept applications has been developed to demonstrate different ways in which SAMP clients can use the query driven visualization messages. They interact with the Astro-WISE awe-prompt through SAMP only and have little knowledge about Astro-WISE, if at all.

⁵<http://www.starlink.ac.uk/topcat/>

⁶<http://aladin.u-strasbg.fr/>

Fig. 2 The Simple Puller application for pulling catalogs over SAMP. It can pull data from any information system that accepts the `target.catalog.pull` message

Catalog Pulling

Starting Catalog:

Selection Criterion:

Attributes:

- The *Simple Puller* (Fig. 2) represents the most basic way an application can pull catalog data. Its sole capability is to send a `target.catalog.pull` message, it cannot receive messages. It requires a minimum amount of input (Section 5.2):
 - An identifier of the base catalog from which the sources are selected.
 - A list of required attributes.
 - A query to select the sources.

The only knowledge the user needs to have about the information system is how these parameters should be specified. This service could be built into existing visualization tools quickly. The demo application uses a web-based interface with the server running locally and relies on other SAMP applications for the actual visualization.

- The *Tree Viewer* (Fig. 3) shows how a SAMP application can use the `target.object.info` message to give the user more information about the data lineage and derivation of a particular dataset. This demo application recognizes several of the classes used in Astro-WISE and is able to interpret some of their properties. The application allows exploration of the dependency graph of a pulled catalog by presenting it as a dot⁷ graph. Clicking on a node sends the `target.object.highlight` message, allowing interaction with the *awe-prompt*.
- The *Object Viewer* (Fig. 4) demonstrates how an application can use the object related messages (`target.object.info`, `target.object.change` and `target.object.action`) to influence the properties of process targets and other objects. It has knowledge about the Astro-WISE Source Collection classes—used to represent astronomical catalogs—and allows many of the actions that can be performed in the *awe-prompt* to be done through the web-based GUI.

These applications rely on other SAMP applications for the actual visualization. For example, Topcat is used in Fig. 5 to visualize the data requested in Fig. 2.

⁷<http://www.graphviz.org/>

Fig. 3 SAMP application for exploring dependency graphs of catalog objects in Astro-WISE. Every node shows the catalog identifier on the *top left*, the class of the catalog in the *top center* and an identifier for the set of sources on the *top right*. The attributes of the sources are shown in the rest of the box

Fig. 4 SAMP application to view and modify details of individual catalogs or other objects. The highlighted catalog from Fig. 3 is shown

Process Target Editor

End of Tree > 271681> -110> -109

-109 FilterSources b

- status: unknown [change](#)
- name: [change](#)
- parent_collection: -107
- query: "R" < 300 [change](#)

Actions

- [Send over SAMP](#)
- [Commit](#)

Attributes

Name	Origin	Actions	Options
SLID	100511	Search	
SID	100511	Search	
SDSS_petroMagErr_g	100511	Search	
SDSS_modelMag_i	100511	Search	
SDSS_petroMagErr_i	100511	Search	

4 Example usage

The figures depicting the prototype applications show a simple use case of the new messages. First the Simple Puller (Fig. 2) is used to request absolute magnitudes and the inverse concentration index for the sources in a specific catalog for which a specific logical expression ($R < 300$) holds. Catalogs in Astro-WISE that can be used for data pulling are called Source Collections and are identified by an integer, in this case 100511. Other information systems might use different identifiers. Attributes are referred to by their name only in Astro-WISE.

Subsequently the Tree Viewer (Fig. 3) is used to inspect the dependency graph that is proposed to provide the requested catalog. The Source Collection that is responsible for the selection of the sample is highlighted. The highlighted object is shown in the Object Viewer (Fig. 4), where the selection criterion is checked and changed if required.

The dependency graph is stored persistently once the scientist has verified that the proposed is suitable for his or her scientific goals, which can be done from the Object Viewer as well. The dependency graph is then optimized automatically before being processed, as described in Paper I. The catalog data of the last node in the dependency graph is send to Topcat for visualization (Fig. 5) once it has been processed.

This example shows how a relatively simple request can result in a complex dependency graph. Nonetheless, this graph can be navigated and changed quickly due to the new SAMP messages. Furthermore, the catalogs created to fulfill the request are created such that they are most suitable for reuse for later requests and at the same time processed in such a way that minimizes the

Fig. 5 A Topcat scatter plot showing a color-concentration diagram of the catalog pulled Fig. 2. A slight bimodality between *red*, concentrated, galaxies and *blue*, extended galaxies can be seen

required calculations. Newly created catalogs are shared implicitly between collaborating scientists.

Therefore, large datasets can be explored quickly with a high level of flexibility, without requiring the scientist to know details of how the information system handles these large datasets.

5 SAMP protocol and messages

We give the details about SAMP that are necessary to describe our proposed messages and present our extensions and their Astro-WISE implementation.

5.1 SAMP protocol and data types

SAMP is in principle language-agnostic and is based on abstract interfaces. That is, it specifies which functions the HUB and the clients must have in order to send and receive messages, but not the exact protocol that the applications use to call those functions. The rules which describe how SAMP functions are mapped to the internally used protocol is described in a *SAMP Profile*. One standard profile based on XML-RPC⁸ is described in the official documentation, and this is what is used in Astro-WISE and in the prototype applications. XML-RPC is a remote procedure call protocol which uses XML to encode its calls and HTTP as a transport mechanism and is platform independent.

Only three data types are supported in SAMP, because it is language- and even communication-protocol-agnostic:

- `string`: A scalar value consisting of a sequence of ASCII-characters.
- `list`: An ordered array of data items.
- `map`: An unordered associative array with a string as key.

Other scalar types have to be mapped to strings, and there is a specification to represent integers, floats and booleans as strings. These data types can be nested to any level: e.g., it is possible to have a map with lists as values.

SAMP applications communicate through messages of specific types. Message types that start with `samp.` are administrative messages defined by the protocol, the others are defined by application authors. Clients are supposed to give a general reply with success or failure of a requested operation, even if no response is required.

5.2 Query driven visualization messages

We designed new SAMP messages and data structures to enable query driven visualization through data pulling mechanisms. The `target.object.*` messages assume that the information system uses an object oriented model for science products such as catalogs (Section 1). The proposed messages are:

- `target.catalog.derive`: Create a catalog through data pulling.
Arguments:
 - `catalog-id` (string): Identifier of the catalog to select the sources from.
 - `query` (string): Selection criterion for the sources.
 - `attributes` (list of strings): Names of the attributes.

⁸<http://www.xmlrpc.com/>

- `target.catalog.pull`: Perform the same action as `target.catalog.derive` and send the data over SAMP. Arguments:
 - `catalog-id` (string): Identifier of the catalog to select the sources from.
 - `query` (string): Selection criterion for the sources.
 - `attributes` (list of strings): Names of the attributes.
- `target.object.highlight`: Highlight an object. Arguments:
 - `class` (string): Class of the object.
 - `object-id` (string): Identifier of the object.
- `target.object.info`: Returns a SAMP map with information about an object as described below. Arguments:
 - `class` (string): Class of the object.
 - `object-id` (string): Identifier of the object.
- `target.object.change`: Change a property of an object. Arguments:
 - `class` (string): Class of the object.
 - `object-id` (string): Identifier of the object.
 - `property-id` (string): Identifier of a property of the object.
 - `value` (string): New value of the property.
- `target.object.action`: Perform an action related to a an object. Arguments:
 - `class` (string): Class of the object.
 - `object-id` (string): Identifier of the object.
 - `property-id` (string, optional): Identifier of a property of the object.
 - `action-id` (string): Identifier of the action.

5.3 Query driven visualization data format

SAMP data structures are defined to send information about objects between applications. The structures are designed generic enough that they could be used for any information system. Information about an object itself, such as the response to an `target.object.info` message, is communicated through a map with the following keys:

- `class` (string): The class of the object. A client that has knowledge about the used classes could handle known classes in a special way.
- `id` (string): Identifier this object, unique in combination with the class.
- `status` (string): Indication the processing status of this object (see below).
- `properties` (list of maps): Properties of this object (see below).
- `actions` (list of maps): Actions that can be performed on this object (see below).
- `readonly` (boolean): Flag to indicate that the object cannot be modified.

Properties of an object, for example process parameters, are described with a map with the following keys:

- `name` (string): Name of the property, as used by the object.
- `class` (string): The class that the value of the property should have, or a primitive such as ‘int’.
- `description` (string): A human readable description of the property.
- `value` (string): The used value for the property. This is the `id` of the object if the property refers to another object.
- `options` (list of maps): Possible values for the property, if applicable.
- `actions` (list of maps): Actions that can be performed on the property.
- `readonly` (boolean): Flag to indicate that the property cannot be modified.

An action that can be performed on an object or property is defined by a map with the following keys:

- `id` (string): A unique identifier for this action.
- `name` (string): A human presentable name for this action.

5.4 Query driven visualization object status

The status value of an object refers to the processing status of the object. It can have the following values:

- `ok`: The object has been processed, or can be processed while retrieving the result.
- `automatic`: The object has to be processed before it can be retrieved. This can be done without user interaction.
- `new`: This is a non persistent object, which can be processed without user interaction.
- `depends`: This is a new object, which can be processed only after human intervention. For example, to set a process parameter that has no proper default.
- `not`: As it is, this object cannot be processed, e.g. because a dependency cannot be fulfilled. The scientist might be able to solve the problem, but whether this is the case is not clear to the information system.
- `unknown`: The status is unknown.

6 SAMP in the Astro-WISE `awe-prompt`

The Astro-WISE `awe-prompt` is an interactive Python prompt that forms the primary user interface to Astro-WISE. We developed a Python module for Astro-WISE to use SAMP from the `awe-prompt`. This allows an astronomer to combine the large scale data handling from Astro-WISE with the visualizations from other SAMP applications.

This section is most interesting for readers already familiar with Astro-WISE. All relevant terms are introduced briefly for readers new to Astro-WISE. The SAMP-related functionality that is not query driven visualization specific, is included as well for completeness.

6.1 SAMP classes and metadata

The SAMP module is split up in two classes, a stand-alone Python SAMP client and a derived class with Astro-WISE specific functionality:

- **SampProxy**: An instance of the *SampProxy* class is a basic SAMP client. This class contains all SAMP code that is not Astro-WISE specific, and can therefore be used by other Python applications as well.
- **Samp**: The *Samp* class is derived from *SampProxy* and contains all Astro-WISE specific code. The metadata that the class declares to the HUB—as stored in its `metadata` property—is:

<code>author.email</code>	buddel@astro.rug.nl
<code>author.name</code>	Hugo Buddelmeijer
<code>home.page</code>	http://www.astro-wise.org
<code>author.affiliation</code>	Kapteyn Astronomical Institute, Groningen
<code>samp.name</code>	Astro-WISE
<code>samp.description.html</code>	<p>Astro-WISE</p>
<code>samp.documentation.url</code>	http://www.astro-wise.org
<code>samp.icon.url</code>	http://www.astro-wise.org/ pics/logo-samp-astrowise.png
<code>samp.description.text</code>	Astro-WISE.

6.2 Sending data

All Astro-WISE objects that represent catalog or image data can be send over SAMP, using the `table.load.votable` and `image.load.fits` messages respectively.

Source catalogs that can be used for data pulling are called *Source Collections* (Paper I). There are different Source Collection classes, depending on the operation used to create the catalog. For example, an *Attribute Calculator* Source Collection is used to calculate new attributes (parameters) of sources. Other catalog related Astro-WISE classes that can be send over SAMP are: the *SourceList* which is primarily used to derive parameters directly from images, the non-persistent *TableConverter* to manipulate tabular data in Python and the *PhotSrcCatalog* used for photometric calibration.

Image data in Astro-WISE is handled by various *Frame* classes. These are beyond the scope of this paper, because its focus is on catalog data.

6.3 Catalog interaction

The SAMP client supports sending and receiving of both the `table.highlight.row` and `table.select.rowList` messages. Sources can be highlighted either by their SAMP row id, or through their Astro-WISE identifiers.

A SourceList has a SLID as identifier, and sources in a SourceList are labeled with a SID. The SLID-SID combination uniquely identifies a source. A Source Collection has a SCID as identifier and can contain sources from multiple SourceLists.

6.4 Query driven visualization data structures

Only Source Collection (Paper I) instances can currently be exported over SAMP through the `target.object.info` message type. The following properties are sent as a reply to such a message:

- All persistent properties that do not relate to data caching. References to other Astro-WISE objects are exported as the unique identifier of the object.
- Process parameters of Attribute Calculators (Paper I) are exported as if they are regular properties.
- The names of the attributes are exported as the `attribute|%i` properties, where the *%i* are consecutive integers. The SCIDs of the Source Collections that the attributes originate from are exported as the `origin|%i` properties.

The actions that can be performed on Source Collections are:

- `commit`: Commits a transient Source Collection.
- `copy`: Creates a copy of a Source Collection.
- `make`: Process the Source Collection. The exact composition of sources and values of the attributes are determined.
- `send`: Broadcasts the catalog data corresponding to a Source Collection over SAMP.

Only the `attribute|%i` properties have an action:

- `search`: Search for Source Collections that could be used as a dependency to provide the attribute. These will be listed in the `options` of the property.

6.5 Receiving query driven visualization messages

The query driven visualization messages from Section 5.1 are supported, but only with respect to Source Collections. The actual data pulling is performed

with the non-persistent *Source Collection Tree* class. The parameters of the messages are interpreted as follows:

- `catalog-id`: The SCID of a Source Collection.
- `query`: An Oracle SQL `WHERE` clause, with attributes in double quotes.
- `attributes`: A list of attribute names.
- `class`: The name of an Astro-WISE class. Only Source Collection classes are supported at the moment.
- `object-id`: The SCID of a Source Collection.
- `property-id`: The name of a property. These are either persistent properties as stored in the database, or transient properties that are derived on the fly.
- `action-id`: The identifier of an action that can be performed on an object or property, as defined by the instance itself.

The query driven visualization messages are handled as follows:

- `target.catalog.derive`: The `derive()` of a Source Collection Tree instance is called to derive a new Source Collection from the specified one.
- `target.catalog.pull`: Performs the same action as `target.catalog.derive` after which the resulting Source Collection is processed and broadcasted.
- `target.object.highlight`: Stores a reference to the highlighted Source Collection as a member of the SAMP instance.
- `target.object.info`: Returns information about a Source Collection.
- `target.object.change`: Change a property of a Source Collection, either directly by the SAMP instance, or by the object itself.
- `target.object.action`: Perform an action related to a Source Collection, either directly by the SAMP instance or by the object itself.

7 Conclusions

In this paper we see query driven visualization as an extension of data pulling, with a focus on catalog data. This allows scientists to discover existing datasets and create new datasets by requesting data directly from within the visualization. New datasets are automatically created in such a way that they are most suitable for reuse in future requests, preventing duplications of data. The subsequent processing of the datasets is limited to those parts that are necessary to create the data for the requested visualization, achieving implicit scalability.

Requesting existing data and creating new data is done through the same process, because data is found and processed automatically. The same mechanisms ensure that scientists have control over the methods and parameters that are used to process their data, achieving flexibility. This allows a high level of abstraction in the interoperation between software, because requests for data can be done in a conceptual way.

The Simple Application Messaging Protocol is an excellent mechanism to provide such a layer abstraction and we proposed new message types to perform query driven visualization. Support for these messages is implemented within Astro-WISE and several prototype applications.

Query driven visualization allows scientists to interact with their data in a conceptual way and allows them to focus on *what* they want to do with the data, because *how* the processing is performed and *where* the data is stored is implicitly taken care of. Current wide field surveys such as KIDS will produce such large datasets that this automation of administration and implicit scalability is essential. Therefore, query driven data visualization is not only a bright possible future, but perhaps even an inevitable one.

Acknowledgement This research is part of the project “Astrovis”, research program STARE (STAR E-Science), funded by the Dutch National Science Foundation (NWO), project no. 643.200.501.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Arnaboldi, M., Neeser, M.J., Parker, L.C., Rosati, P., Lombardi, M., Dietrich, J.P., Hummel, W.: ESO Public Surveys with the VST and VISTA. *The Messenger* **127**, 28 (2007)
2. Bethel, E., Campbell, S., Dart, E., Stockinger, K., Wu, K.: Accelerating network traffic analytics using query-driven visualization. In: *IEEE Symposium On Visual Analytics Science And Technology*, pp. 115–122 (2006)
3. Buddelmeijer, H., Boxhoorn, D., Valentijn, E.A.: Automatic optimized discovery, creation and processing of astronomical catalogs. Paper I. *Exp. Astron.* (2011). doi:[10.1007/s10686-011-9272-z](https://doi.org/10.1007/s10686-011-9272-z)
4. Codd, E.F.: A relational model of data for large shared data banks. *Commun. ACM* **13**, 377–387 (1970)
5. Laureijs, R.: Euclid Assessment Study Report for the ESA Cosmic Visions. [arXiv:0912.0914](https://arxiv.org/abs/0912.0914) (2009)
6. Smith, G., Czerwinski, M., Meyers, B., Robertson, G., Tan, D.: Facetmap: A scalable search and browse visualization. *IEEE Trans. Vis. Comput. Graph.* **12**(5), 797–804 (2006)
7. Stockinger, K., Shalf, J., Bethel, W., Wu, K.: Query-driven visualization of large data sets. In: *Proceedings of IEEE Visualization 2005*, pp. 167–174 (2005)
8. Stockinger, K., Bethel, E.W., Campbell, S., Dart, E., Wu, K.: Detecting distributed scans using high-performance query-driven visualization. In: *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing*, ACM, New York, NY, USA, SC 06 (2006)