

Semantically-Enabled Large-Scale Science Data Repositories

Peter Fox¹, Deborah McGuinness^{2,3}, Don Middleton⁴, Luca Cinquini⁴,
J. Anthony Darnell¹, Jose Garcia¹, Patrick West¹,
James Benedict³, and Stan Solomon¹

¹ High Altitude Observatory, ESSL/NCAR
PO Box 3000, Boulder CO 80307-3000
pfox@ucar.edu

² Knowledge Systems, Artificial Intelligence Lab
353 Serra Mall, Stanford University, Stanford, CA 94305

³ McGuinness Associates
20 Peter Coutts Circle, Stanford, CA 94305
d1m@cs.stanford.edu

⁴ Scientific Computing Division, CISL/NCAR
PO Box 3000, Boulder CO 80307-3000

Abstract. Large heterogeneous online repositories of scientific information have the potential to change the way science is done today. In order for this potential to be realized, numerous challenges must be addressed concerning access to and interoperability of the online scientific data. In our work, we are using semantic web technologies to improve access and interoperability by providing a framework for collaboration and a basis for building and distributing advanced data simulation tools. Our initial scientific focus area is the solar terrestrial physics community. In this paper, we will present our work on the Virtual Solar Terrestrial Observatory (VSTO). We will present the emerging trend of the virtual observatory - a virtual integrated evolving scientific data repository - and describe the general use case and our semantically-enabled architecture. We will also present our specific implementation and describe the benefits of the semantic web in this setting. Further, we speculate on the future of the growing adoption of semantic technologies in this important application area of scientific cyberinfrastructure and semantically enabled scientific data repositories.

1 Introduction

Semantic Web technology has the potential to enable new work paradigms and make vast changes to existing and/or emerging paradigms. One emerging area of scientific work practice where semantic technologies are starting to be used is the area of Virtual Observatories (VOs). VOs [3] are distributed resources that may contain vast amounts of scientific observational data, theoretical models, and analysis programs and results from a broad range of disciplines. Semantics in VOs resemble efforts in the Semantic Grid efforts [5] and science more generally[2].

Currently semantics are mostly usable with the aid of an experienced researcher and only in a narrow domain. One goal of Virtual Observatories is to enable not just expert researchers working in narrow domains to make progress but also to make young researchers and interdisciplinary researchers much more efficient and enable new, global problems to be solved as researchers may now access vast amounts of data that they or others have created. The key to the new efficiency is that users (humans and agents) must now be able to understand what the data is describing, how the data (and topic area) relates to other data (and other topic areas), how the data was collected, and what assumptions are being used. These problems are a perfect match for semantic technologies.

We are using semantic technologies to create an interdisciplinary Virtual Solar Terrestrial Observatory (VSTO) [15]. The work to create a scalable environment for searching, integrating, and analyzing databases distributed over the Internet requires a higher level of semantic interoperability than was previously required by most (if not all) distributed data systems or discipline specific virtual observatories. We leveraged existing background domain ontologies (SWEET) [13] and generated our own ontologies in OWL [10] covering the required subject areas. We leverage the precise formal definitions of the terms in supporting semantic search and interoperability.

Our science domain area - solar and solar-terrestrial physics - utilizes a balance of observational data, theoretical models and analysis/interpretation to make effective progress. Since many of the data collections are increasingly growing in volume and complexity, the task of truly making them a research resource that is easy to find, access, compare and utilize is still a very significant challenge to discipline researchers. The datasets can be highly interdisciplinary as well as complex and provide a good initial focus for virtual observatory work since the datasets are of significant scientific value to a set of researchers and capture many, if not all, of the challenges inherent complex, diverse scientific data.

VSTO addresses the next logical and intellectual challenge for scientific cyberinfrastructure work: that of an interdisciplinary virtual observatory which requires advances both in computer science areas such as knowledge representation and ontology development as well as depth in the science areas of concern to provide an appropriate scientific infrastructure that is usable and extensible.

In this article we describe our Virtual Observatory project, including the vision, our design and implementation. We will describe where we are using Semantic Web technologies and discuss our motivation for using them and some benefits we are realizing. We will also briefly describe our deployment setting which started production late summer of 2006.

2 Virtual Observatories

The virtual observatory (VO) vision includes a distributed, virtual, ubiquitous, semantically integrated scientific repository where scientists (and possibly lay people) can access data. The VO data repository appears to be local and integrated. In the VO vision, tools exist that make it easy for users to access the data they want.

Additionally, support exists for helping them understand the data, its embedded assumptions, and any inherent uncertainties in a discipline-specific context.

3 The Virtual Solar Terrestrial Observatory

Our VSTO project inherits all of the goals of a general virtual observatory project [8,14] while focusing on the domains of solar and space physics. Our goal is to provide virtual access to specific data, model, tool and material archives containing items from a variety of space- and ground-based instruments experiments, as well as individual and community modeling and software efforts bridging research and educational use.

We have developed a series of use cases sampling the solar-terrestrial problem domain, drawing on existing experience and either fully or partially developed capabilities in existing domain-relevant scientific services. Prior to beginning this project, the National Center for Atmospheric Research already had the CEDAR-WEB [4] (CEDAR - Coupled Energetics and Dynamics of Atmospheric Regions; an NSF focused program covering aeronomy and terrestrial upper-atmosphere studies), Center for Integrated Space-Weather Modeling at Boston and funded by an NSF Science and Technology Center grant, and the Advanced Coronal Observing System (ACOS) at the Mauna Loa Solar Observatory (MLSO) operated by the National Center for Atmospheric Research's High Altitude Observatory. Our team includes members who are key contributors to those existing efforts, some of whom started and/or maintain those systems. This team makeup helps with our efforts involved both in gathering requirements as well as in providing a transition plan for deployment and acceptance in the communities.

3.1 Solar-Terrestrial Use Cases

We have developed a number of use cases for VSTO and in this section we will present the first of these which addresses the state of the neutral terrestrial upper atmosphere. We will describe how this use case (and how some aspects of the other use cases) have contributed to our ontology and semantic web architecture requirements.

Our general use case is of the form "Plot values of a particular parameter as recorded by a particular instrument subject to certain constraints in a particular time period, in a manner that makes sense for the data." An instantiation of this pattern that may be asked of our implemented system is: "Plot the observed/measured Neutral Temperature as recorded by the Millstone Hill Fabry-Perot interferometer while looking in the vertical direction during January 2000 in a way that makes sense for the data."

This use case serves as a prototypical example for our target scientific community, that if answered will help the scientists do their research more efficiently and in a more collaborative manner. Our goal from a semantic web perspective is to demonstrate the development of the semantic framework for a virtual observatory while leveraging existing data sources and (catalog and plotting) services.

The anticipated result is a successful return of a graphical representation of the specified data.

The second use case is in the field of solar physics, specifically assisting a user in finding images of the solar atmosphere from the advanced coronal observing system during a particular time period. This process involves a series of selections similar to the first use case but includes an additional stage of finding preview images prior to the actual data selection. Our goal is the successful identification and selection of solar image data with the successful outcome that the user finds preview images and downloads image data.

In the analysis of the second use case, a substantial similarity with the first use case - observatory, instrument, parameter, dataset, etc. was identified and we added only a few properties and sub-classes to the instrument ontology. As a result of this conceptual similarity in the workflow we were able to generalize the workflow interface. Figure 1 displays the general workflow within the context of the semantic web framework, including connections to information sources.




Fig. 1. Generalized workflow for VSTO production release - integrating two use cases from two different disciplines

1. User accesses the portal application (or otherwise accesses application with or without authenticating)
2. User may select from three generic workflows (combination of instrument selection, date-time selection and parameter selection) which are first class

objects in the ontology. At each step, the user selection determines the range of available options in the subsequent steps. The remainder of the information to proceed to a catalog query and thus a data selection request is inferred (using Pellet) from the ontology. At the final stage of the request, we inferred the return data types and possible ways of plotting the data, which includes whether they are time-series, height and time, images, etc.

3. The framework validates the user request: first it verifies that the user is authorized to access the specific kind of data, then it verifies the logical correctness of the request, i.e. that Millstone Hill is an observatory that operates a type of instrument that measures neutral temperature (i.e. check that Millstone Hill is an observatory and check that the range of the measures property on the Millstone Hill Fabry Perot Interferometer subsumes neutral temperature).
4. The application processes the user request to locate the physical storage of the data, returning for example a URL-like expression: find Millstone Hill Fabry-Perot Interferometer data of the correct type (operating mode; defined by a specific operating model since the instrument has two operating modes) in the given time range.
5. The application plots the data in the specified plot type (inferred to be a time series). This step involves extracting the data from records of one or more files, creating an aggregate array of data with independent variable time (of day or day+time depending on time range selected) and passing this to a procedure to create the resulting image.

3.2 Architecture

One of the overriding principles to virtual observatories is to be able to find and retrieve a wide variety of data sources. As a result, the ability to rapidly develop the semantic framework, deploy and test it is essential. Fortunately, the availability of the OWL language, and software tools and plug-ins such as Protégé supported rapid ontology building and additional tools, such as Pellet [11] also supported reasoning and queries for testing.

In Figure 2 the current VSTO architecture is represented graphically. It utilizes the Jena [7] and Eclipse [6] plug-ins for Protégé to generate the Java stub code for the ontology classes and allows the incorporation of existing calls to the CEDAR catalog service for the date and time coverage for the data from the instruments (the remainder of the previous calls to the catalog, implemented in MySQL, are encoded as individuals in the ontology).

The user interface is built on the Spring [12] framework, which encodes the workflow and navigation features. The examples of the prototype implementation are displayed in later figures. The initial implementation includes the need for reasoning which is implemented via the Pellet reasoning engine which will operate on over 10,000 triples and typically returns results in a few seconds on our deployment platform.

As a part of the implementation, we utilize an existing set of services for returning selections over a large number (over 60 million records) of date/time




Fig. 2. Overall VSTO software architecture

information in the CEDAR database. We also utilize a set of existing services for plotting the returned data which are currently operating in the production CEDARWEB. These services utilize the Interactive Data Language as well as the Open Source Project for Network Data Access Protocol [9] to access the relevant data elements from the data archive. The ability to rapidly re-use these services is an essential and effective tool in our effort to deploy a production data-driven virtual observatory environment.

3.3 Ontology Focus Areas

We began our ontology development process after carefully analyzing our use cases to look for important classes, instances, and relationships between terms. We also looked at critical controlled vocabulary starting points that were either already included in our base implementations of the existing CEDAR and Mauna Loa Solar Observatory services. One such starting point was the controlled vocabulary associated with the CEDAR database which has a long history in the upper atmospheric and aeronomy communities. For a history of the CEDAR program and the CEDAR database, visit the current website - <http://cedarweb.hao.ucar.edu>. Data in the CEDAR database was arranged around date of observation and a combined observatory/instrument classification. Within each dataset, a series of tables is encoded in a so-called CEDAR binary format which holds the parameters. Each observatory/instrument and parameter has a long name, a mnemonic name and a numeric code.

In developing the ontology, we drew upon the vocabulary of the use case, the existing vocabulary of CEDAR and wherever possible the terms and concepts in the semantic web for earth and environmental terminology (SWEET) ontology. In the case of SWEET, to date there has been limited application to the earth's upper atmosphere (i.e. realms in SWEET terminology) so we adopted parts of SWEET that applied to our needs and for the time being, developed our ontology separately from SWEET but keeping in mind that our aim is to merge much of what we develop back into SWEET for broad use. Our goal was to keep our ontology development separate until we believed it was stable and vetted at two different workshops which brought together domain scientists to discuss foundational earth and space science ontologies and related issues.

One of the first classes to be discussed in the use case was the concept of an instrument, in this case a Fabry-Perot Interferometer (see description below). One of our contributions both to our domain specific work on VSTO and to general work on virtual observatories is our work on the instrument ontology. We constructed a class hierarchy of Instrument (see Figure 3), OpticalInstrument, Interferometer and then Fabry-Perot Interferometer (as known as FPI, for which the Millstone Hill FPI is an individual of the last class). With each class for the initial prototype we added the minimal set of properties at each level in the class hierarchy. The production release features a more complete but still evolving set of properties across all classes. In the next few paragraphs, we elaborate on a few of the ontology classes in order to give enough background for the impact discussion later.




Fig. 3. Schematic of VSTO Ontology 1.0 indicating a variety of classes: for data, service, service implementation and value restrictions. We also indicate a few properties/associations, inheritance and inference.

Instrument: description - A device that measures a physical phenomenon or parameter. At a minimum, it possesses a detector which produces a signal from which the desired quantity is calculated or inferred.

OpticalInstrument: description - Instrument that utilizes optical elements, i.e. passing photons (light) through the system elements which may be reflective and transmissive and may include filters.

Interferometer: description - An instrument that uses the principle of interference of electromagnetic waves for purposes of measurement. Note: Interferometers may be used to measure a variety of physical variables, such as displacement (distance), temperature, pressure, and strain.

Fabry-PerotInterferometer: description - A multiple-beam interferometer. Due to their optical and mechanical configuration, Fabry-Perot interferometers can also be used as spectrometers with high resolution. This description highlights one important attribute of this instrument sub-class which we highlight in a later section: that a “Fabry-Perot interferometer scan be used as a spectrometer with high resolution”.

We also have built an initial Instrument class hierarchy as a result of all the instruments utilized in generating the CEDAR and MLSO data holdings. This hierarchy is encoded in OWL and is part of the VSTO framework. Below is an excerpt from the list of the OpticalInstrument class with some subclasses abbreviated in parentheses.

- OpticalInstrument
 - Heliograph {SpectroHeliograph, ... }
 - Interferometer
 - * Fabry-PerotInterferometer
 - * MichelsonInterferometer {InfraredMichelsonInterferometer, DopplerMichelsonInterferometer, ...}
 - Imager {AirGlowImager, AllSkyImager ...}
 - Lidar {AerosolLidar, CalciumLidar, DifferentialAbsorptionLidar, DopplerLidar, IonLidar, OzoneLidar, RamanLidar, RayleighLidar, SodiumLidar, StrontiumLidar, ...}
 - Photometer {SingleChannelPhotometer, MultiChannelPhotometer, SpectroPhotometer, ...}
 - Polarimeter {SpectroPolarimeter, ... }
 - Spectrometer {SpectroPhotometer, SpectroHeliograph, SpectroPolarimeter, ...}

In all cases, the class properties are associated with value restrictions, but these are not discussed here.

The next important class is the InstrumentOperatingMode with depends on the Instrument and leads to a particular type of physical quantity (parameter) being measured and an indication of its domain of applicability and how it should be interpreted. Its description is: A configuration which allows the instrument to produce the required signal.

In practice for terrestrial atmosphere use case the instrument operating mode indicates which direction the FPI is pointing, i.e. vertical, or “horizontal” - actually 30° or 45°. Knowing these modes is critical for understanding and using the data as different quantities are measured in each mode and geometric projection, i.e. north component of neutral wind has to be calculated correctly depending on the mode.

Also shown in Figure 3 is the Observatory class, whose description for the sub-class GroundBasedObservatory is: A facility which houses and operates one

or more instruments either synoptically or periodically. It has a physical location and operating hours. It can be either manned or remote.

An important part of the use case is the actual quantity that a user (scientist) is seeking. This entity is captured in the class `Parameter` (also known as `PhysicalQuantity` in SWEET). Its description is: A measured physical property in signal units. It has units and possibly a timestamp. The signal units are not the physical ones, like Gauss or Kelvin, but are something inherent in the detector, like volts. The physical units are either calculated or inferred from the signal units.

In developing a production implementation of the VSTO it has been essential to make the connection between the high-level concepts of the ontology classes all the way to the data itself. This entails the data files, the data constraints, and the underlying catalogs, and data and plotting services - all of which have been in existence for some time and are made available from distributed network sites and accessed via common internet protocols (ftp, http, web services, etc.). Thus we fill out the ontology with data-related classes (see Figure 3).

Dataset: description - A collection of observations of the physical quantity of interest. They usually have a location, observatory, instrument, and parameter (or set of parameters) associated with them. They also have a format along with an epoch over which they were taken.

Perhaps the most important property of the `Dataset` class is: `hasContained-Parameter`, which is the asserted association with the `Parameter` class which in turn connects to the instrument, etc.

For a user of the VSTO, the creation of a data product based on the series of user choices and available data constraints. We represent this at the `DataProduct` (and associated `Request`, `Service` and `Metadata`) classes which we will not give details on in this paper.

Additional, data-related classes are as follows:

DataRequest: description - Generic class representing a request for data. The class contains both the information necessary to define the data to be extracted (input), and the form of the resulting data product (output).

Dataservice: description - Generic class representing the outcome of a data request to a service. It acts as a wrapper for a collection of `DataProduct` objects.

MetadataService: description - Generic class that defines the functionality for querying metadata information from a data archive. The results of the query may be constrained by an associated `DataConstraint` object. Instances of `MetadataService` and `DataConstraint` are created on demand to support a transient query session.

As a final integrating theme, the parameters and instruments of interest have a physical domain of influence which needs to be represented. In use case for the terrestrial atmosphere, the SWEET-equivalent class `AtmosphereLayer` contains layers known as: Thermosphere which ranges from 80-85 km (i.e. the upper boundary of the Mesosphere) to greater than ≈ 640 km.

A final note for the ontology development is that for the classes included in the current implementation, we encode all the individuals within the ontology

except for the date and time (class: `DateTime`). The latter is a practical choice due to the large number of specific instances of date and time records associated with the diverse set of instrument datasets associated with both use cases. Thus, we implement a set of service classes to execute queries and retrieve results from the underlying (SQL) catalogs for each set of data holdings.

3.4 Discussion

Ontologies are used throughout the workflow to guide the user through the consecutive selection steps leading to the final service request. By representing physical instruments and their output streams as concrete instances of classes the application is able to follow the relationships between classes so to always present to the user a range of sensible valid options that greatly reduces the amount of specific knowledge the user needs to already possess about the data.

Our application includes an ontology covering important domain concepts (observatories, instruments, operating modes and parameters). We have found that the ontology can be easily reused in related efforts and additionally found it to be more flexible and extensible than a traditional database-like system.

Our ontologies include annotations at the class, instance, and property level and these annotations contain formal as well as informal descriptions. These descriptions may be used by domain experts (scientists and researchers) as well as by other application users (other scientists from the same or different domain as well as teachers, students, etc.).

We are not simply using ontologies for straightforward class subsumption and instance recognition. With the help of a reasoning engine we are also using them to infer the possible plot type based on the selected parameter. Plot type possibility deduction reduces the level of knowledge required from users and is possible because of the combination of reasoning and declarative background knowledge encoding.

In addition to this simple inference, there are many related and valuable inferential requirements for our application areas. One example inference is the selection of instruments that measure the ‘same’ parameter. Previously, users needed to know a significant amount of domain-specific information to be able to guess or choose which other classes of instruments or specific instruments were relevant. The semantic framework can not only infer this information but also explain how the inference was made.

A second example of inference is highlighted by our first use case (see the description of FPI above) where the FPI is able to operate as a spectrometer, i.e. an interferometer operates as something else in the `OpticalInstrument` class hierarchy. As a result, we can infer this in such a way that the framework uses inheritance but does not give up or need to override any properties on the spectrometer. Thus, a user seeking a particular type of spectral intensity (parameter) measurement, e.g. over a certain wavelength with high spectral resolution would be able to find not only data from spectrometers but also from Fabry-Perot Interferometers, the latter being an unknown source of data.

Our work in ontology-supported virtual observatories in two fairly distinct discipline-specific use cases has come together to allow us to produce an integrated ontology for semantic integration of solar terrestrial scientific integration. The resulting overall ontology which is used to generate the semantic framework is thus based on the core set of ontologies and then includes discipline-specific classes and instances for each of the solar and terrestrial upper atmospheres.

Prior to the production release we made the new portal available for internal testing to a group of science and data literate users who were very familiar with the existing services and had specific functional requirements. We also solicited input and evaluations from domain experts on our ontology developments both at small workshop and large national and international conferences where we presented, talks, posters and demonstrations. Now that the portal is released we will perform an evaluation study in about six months.

3.5 Status

The generalized workflow for our use cases (see Figure 1) implement data service capabilities for two significant scientific overlapping but distinct communities. The CEDAR community has over 1200 participants, ≈ 600 of which are registered and active users of the CEDAR data holdings which comprise ~ 1370 datasets, and over 320 distinct instruments/data sources. The ACOS instrument suite has an active user community base (ranging from individuals to agencies) of ≈ 120 and features $\sim 300,000$ datasets with a total archive size of ≈ 10 TBytes.

The VSTO portal supercedes the existing operation of both the CEDAR instruments and ACOS web-based data search and retrieval. As such, the semantic framework described above and implemented within the initial VSTO immediately has a large user base and delivers a wider range of functionality over the existing datasets.

After initial work on designing use cases and our architecture, and identifying existing data services that could be leveraged (such as the CEDARWEB data retrieval services), we implemented a prototype interface. After testing and demoing that prototype, we implemented a second use case and continued to populate the ontology and add the required additional services to support the production implementation.

The production interface features a full population of the ontology classes in Figure3, including all individuals (except date and time). As a result of encoding these parts of the ontology the performance of the interface to progression through the workflow is much faster than the existing CEDARWEB interface which queries the SQL catalog as required for all steps in the workflow.

The production VSTO portal also accommodates the security and audit mechanisms in place for the existing CEDARWEB site. At present, we utilize existing data and plotting services for the production VSTO portal and document these services in OWL-S.




Fig. 4. VSTO 1.0 production portal for instrument selection, with possible domain and semantic filter operations

3.6 Conclusion

We have prototyped and deployed a production quality semantically-enabled data framework to support two large, heterogeneous, online repositories in the area of solar, and solar-terrestrial physics. We have utilized numerous semantic web technologies in the process of searching and accessing the data and created an interoperable and dynamically configurable framework. We see this as a major step toward a virtual integrated evolving scientific data repository.

We have found significant benefit in encoding the formal descriptions of terms in OWL and using inference to provide completion. The primary benefits are (i) reducing input specification requirements thus decreasing input burden and more importantly (ii) allowing users to be able to create correct (and non-over constrained) queries without needing to have expert-level knowledge. Previously, we noticed non-experts having difficulty generating error data requests.

In this implementation we also made practical choices as to what level of detail of the science and processing concepts we encoded in OWL and what aspects of the search, access and services we defer to, and thus re-use, existing services.

We are presently implementing the next series of use cases which are enabling us to further populate the ontologies and validate them. For example, on instantiation of the next use case is: “Find data which represents the state of the neutral atmosphere anywhere above 100km and toward the arctic circle (above 45N) at any time of high geomagnetic activity.” The vocabulary of this use case has much less direct mapping to the classes in the first use case and thus additional terms and additional reasoning based on properties of the existing classes is required.

Finally, in the medium term we are exploring options for using our semantic framework and rapid prototyping environment to develop a configurator within specific disciplines to enable the assembly of a virtual observatory within that discipline, or specific to a project/task based using a subset of our ontology.

Acknowledgements

VSTO is an NSF Shared Cyberinfrastructure project under award 0431153 and SESDI is a semantic science data integration project sponsored by NASA Advancing Collaborative Connections for Earth-Sun System Science (ACCESS) and NASA Earth-Sun System Technology Office (ESTO) under award AIST-QRS-06-0016.

References

1. Advanced Coronal Observing System (ACOS), <http://mlso.hao.ucar.edu>
2. Tim Berners-Lee, Wendy Hall, James Hendler, Nigel Shadbolt, and Daniel J. Weitzner 2006, Enhanced: Creating a Science of the Web, *Science*, **313** #5788, pp. 769-771, DOI: 10.1126/science.1126902
3. The US NVO White Paper: Toward a National Virtual Observatory: Science Goals, Technical Challenges, and Implementation Plan, 2001, Virtual Observatories of the Future, *ASP Conference Proceedings*, **225**, Ed.; Robert J. Brunner, S. George Djorgovski, and Alex S. Szalay, San Francisco: Astronomical Society of the Pacific, p.353
4. <http://cedarweb.hao.ucar.edu>
5. De Roure, D. Jennings, N.R. Shadbolt, N.R. 2005, The semantic grid: past, present, and future, *Proceedings of the IEEE*, **93**, Issue: 3, pp. 669-681, DOI: 10.1109/JPROC.2004.842781
6. <http://www.eclipse.org/>
7. <http://jena.sourceforge.net/>
8. <http://www.us-vo.org>
9. Open source Project for Network Data Access Protocol (OPeNDAP), <http://www.opendap.org>
10. Deborah L. McGuinness and Frank van Harmelen. OWL Web Ontology Language Overview. World Wide Web Consortium (W3C) Recommendation. February 10, 2004. Available from <http://www.w3.org/TR/owl-features/>
11. <http://www.mindswap.org/2003/pellet/>
12. <http://www.springframework.org/>
13. <http://sweet.jpl.nasa.gov>
14. <http://virtualsolar.org>
15. <http://www.vsto.org>, <http://vsto.hao.ucar.edu>