USING DATASPACE TO SUPPORT LONG-TERM STEWARDSHIP OF REMOTE AND DISTRIBUTED DATA

Robert L. Grossman, Dave Hanley,
Xinwei Hong and Parthasarathy Krishnaswamy
Laboratory for Advanced Computing
University of Illinois at Chicago
851 South Morgan Street
Chicago, Illinois 60607
tel: +1-312-413-2176
email: grossman@uic.edu, dave@lac.uic.edu,
xwhong@lac.uic.edu, babu@lac.uic.edu

1 Introduction

In this note, we introduce DataSpace Archives. DataSpace Archives are built on top of DataSpace's DSTP servers [2] and are designed not only to provide a long term archiving of data, but also to enable the archived data to be discovered, explored, integrated and mined.

DataSpace Archives are based upon web services. Web services' UDDI and WSDL mechanisms provide a simple means for any web service client to discover relevant archived data [7]. In addition, data in DataSpace Archives can carry a variety of XML metadata, and the DSTP servers which underly the DataSpace Archives provide direct access to this metadata.

Unfortunately, web services today do not provide the scalabilty required to work with large remote data sets. For this reason, DataSpace Archives employ a scalable web service we have developed called SOAP+.

As the amount of data grows, the ability to explore and browse remote and distributed archived data will become more and more important. For this reason, a requirement of DataSpace Archives is that they support direct browsing of the data they contain, without the necessity of first retrieving the data and then opening a local application. DataSpace Archives also support a type of distributed database keys, which are described below and which enable data sets in different DataSpace Archives to be easily integrated.

Finally, DataSpace Archives use emerging internet storage platforms, such as IBP [1] and OceanStore [6], as a basis for providing long term storage, long past the demise of any individual disk or server.

2 The DataSpace Transfer Protocol (DSTP)

In previous work we have developed a protocol called the the DataSpace Transport Protocol or DSTP [2]. Data in DataSpace Archives can be accessed directly using DSTP or indirectly using web service based DSTP operations. Here is a quick summary of DSTP.

Data model. Data accessible via DSTP servers form a distributed collection of attribute-based data (in contrast to the file-based data that is usually available in data archives) that we refer to as *DataSpace*. At the simplest, data in DataSpace consists of records which may be distributed across nodes either vertically (by attributes) or horizontally (by records). Data records and data attributes are joined using universal keys, which are described next.

Universal Keys. One of the novel aspects of DataSpace is that certain attributes can be identified as universal (correlation) keys or UCKs by associating globally unique IDs or GUIDs to them. The assumption in DataSpace is that two distributed attributes having the same universal keys, as identified by their GUIDs, can be joined. This simple mechanism enables DataSpace to support vertically partitioned data, that is data records whose attributes are geographically distributed across DSTP servers.

In addition, universal keys may be attached to data sets and, in this way, data sets may be horizontally partitioned. In other words, data records may be geographically distributed across DSTP servers.

Metadata. The DataSpace infrastructure provides direct support for metadata. Each data set and each attribute can have metadata associated with it. In general, we assume that the metadata is in XML. Some DataSpace applications, with large amounts of metadata, use alternate formats for metadata for greater efficiency.

Data Access. DSTP-based access to data is via SOAP/XML or what we call SOAP+. SOAP+ is a variant of SOAP we have developed which employs a separate SOAP/XML-based control channel together with a data channel which employs a streaming protocol for moving large amounts of data or metadata efficiently. Depending upon the request, DSTP servers may return one or more attributes, one or more records, or entire data sets. DSTP servers can also return metadata about data sets or a list of universal keys associated with a data set. Each DSTP server has a special file called the catalog file containing XML metadata about the data sets on the server.

AAA Model. The default access and security model in DataSpace s a web-based instead of a grid-based access mechanism. The difference is how authentication, authorization, access (AAA) are handled. For long term stewardship, mechanisms to manage AAA are quite challenging. The default assumption in DataSpace is that data is open to anyone with a browser. Clearly, many data sets require some type of AAA infrastructure. In these cases, a AAA infrastructure can be implemented using of the standard approaches, such as Globus GSI, IETF SSL/TLS, etc.

DSTP clients and servers support the following services:

- Discovery queries. Discovery in DataSpace is via web service's WSDL and UDDI mechanisms. This provides a standards based discovery mechanism for the discovery of data sets, data attributes, metadata, etc.
- Metadata queries. DSTP servers automatically create XML metadata about the data they serve and provide a simple mechanism for user supplied metadata. DSTP Clients can request attribute based metadata, data set based metadata, or metadata summarizing all the datasets managed by the DSTP server. For example, metadata associated with an attribute typically contains the number of data records on the DSTP server associated with that attribute, a description of the attribute, the min and max values, and perhaps the providence of the attribute.
- UCK queries. Several DSTP operations are based upon universal keys or UCks. For example, a DSTP client can request all UCKs from two distributed DSTP servers, set a UCK, and then request all attributes associated with that UCK to join vertically partitioned data. Similarly, a DSTP client can request the UCKs from two distributed DSTP servers associated with data sets on the server, set a data set UCK, and retrieve all data records associated with that UCK to merge two horizontally partitioned data sets.
- Range based queries. DSTP client and servers support range based queries. Ranges may be determined using a single UCK or using several UCKs.
- Server side sampling. It is easy for DSTP Servers to overwhelm DSTP Client applications with data. The DSTP servers support server side sampling so that the appropriate amounts of data can be returned.
- Support for missing values. DSTP servers and clients support missing values as a primitive data type. This is important for exploratory data analysis and many data mining applications.

3 DataSpace Archives

We define a DataSpace Archive to consist of one or more DataSpace DSTP servers with the following the additional structure:

- *Internet-based, replicated storage.* Since we cannot guarantee over long periods of time that DSTP servers will be backed up and that the hardware will be maintained, we have decided instead to rely, in part, on distributed, replicated internet based storage, such as that provided by the OceanStore Project [6] or an Internet Backplane Protocol enabled storage system [1]. We are not suggesting that either of these two projects is ready yet to provide the type of long term storage required for archiving data, but rather that, in principle, this type of approach may play a role in helping to ensure the long term survival of data.

- *XML based metadata.* DSTP servers have a mechanism for attaching auxiliary information, such as XML metadata, to data sets. For example, PMML data may be attached to data sets in DSTP servers. The Predictive Model Markup (PMML) Language is an XML based markup language for working with statistical and data mining data sets and models [3]. It can be used to describe data attributes, data mining attributes, and some common transformations used for preparing data for analysis and mining. Maintaining data over long periods of time is made more complex by the many transformations, normalizations, and aggregations that are generally part of data prepration. The assumption in DataSpace Archives is that these are done using the transformations specified by an open XML-based standard, such as PMML.

- *Physical replicas.* In practice, one of the most durable mechanisms for preserving data is paper, as well as other durable materials such as metal or stone. Although neither practical nor desirable in many cases, it is still useful to be able to associate physical replicas with DataSpace data sets. If an ID (bar code, radio frequency ID, etc.) is attached to the physical replica, then this ID may be associated with the data set in a DataSpace Archive. More generally, GUIDS of various types may be associated with attributes and data sets in DataSpace. These can be used so that reports, laboratory books, disks, tapes, etc. can be associated with attributes, collections of attributes, and data sets in DataSpace.

4 Implementation and Experimental Studies

To understand some of the stewardship issues associated with DataSpace Archives, we have taken the data sets in the University of California at Irvine Knowledge Discovery in Databases Archive (UCI KDD Archive) [5] and created a DataSpace Archive using them. This is available on line at data.dataspaceweb.org.

Prior to the DataSpace Archive which we created, the only way to retrieve data from the UCI KDD Archive was via ftp. Moreover, as far as we are aware, the UCK KDD Archive was available on only a single server. Using a DataSpace Archive, UCI KDD data can now not only be directly viewed, but in addition, it can be explored with simple exploratory data analysis operations. In addition, XML based metadata is readily available.

We have also integrated the DSTP server with the Internet Backplane Protocol or IBP [1], which provides a baseline replication of the physical storage underlying the DSTP server. We are currently exploring additional internet based storage approaches in order to improve the long term survival of the data, such as OceanStore [6]. We have not implemented any physical replicas yet.

Finally, we note that DataSpace's support for UCKs allows UCK KDD Archive data to be integrated and correlated with other DataSpace data for the first time.

Number Records	DSTP: SOAP/XML (sec)	DSTP: SOAP/XML+ (sec)	Speed-up
10,000	0.65	0.21	3.10
50,000	177	0.72	245.83
150,000	673	125	5.38
375,000	3078	301	10.23
1,000,000	21121	823	25.66

Table 1: DSTP servers have two modes for data access. One mode uses SOAP/XML, which works for small data sets, and (small) metadata. The other mode uses a new protocol called SOAP+ with seperate control and data channels which scales much better for large and complex data sets.

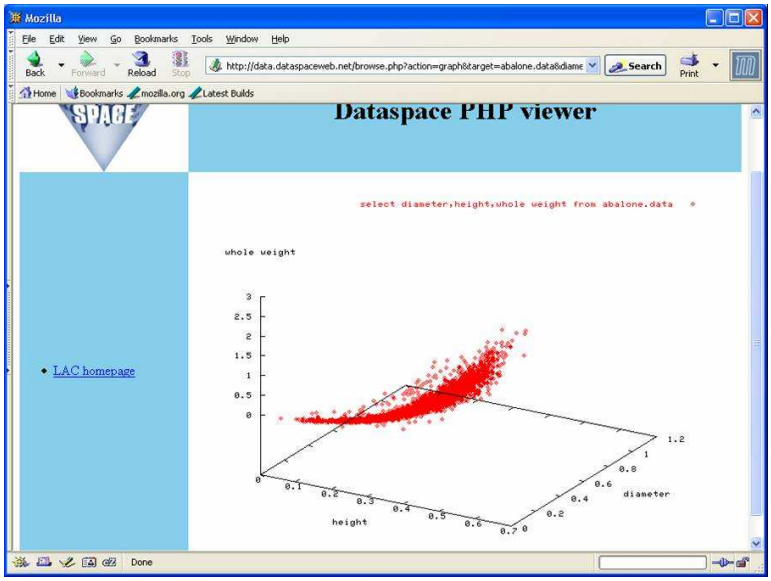

Figure 1: We have put most of the data from the UCI KDD data sets into DataSpace Archives. This figure illustrates the results of browsing one of the UCI KDD data sets and graphing the results.

5 Summary and Conclusion

In this note, we have introduced DataSpace Archives, which are archives for data sets built on top of DataSpace's DSTP servers. DataSpace Archives are designed to support the long term archiving of data in a format in which it is easy to browse and explore the archived data. In addition, DataSpace Archives enable geographically distributed data to be integrated, even if the data comes from separate data sets. In general, most current archives provide ftp access to data, but do not support the browsing, exploration, and integration of large quantities of remote and distributed data.

We have implemented DataSpace Archives on top of Version 3.0 of DataSpace and created an archive containing the University of California at Irvine Knowledge Discovery in Databases Archive.

We are currently creating a DataSpace Archive containing approximately a hundred data sets and about a Terabyte of data in order to test these ideas further.

We are also currently improving the internet based replicated storage for DataSpace Archives, as well exploring additional mechanisms to ensure that long term survival of the data is possible, such as the use of physical replicas for important data sets.

We are also working with the Data Mining Group [4] to standardize web services for data discovery, data integration, and data mining so that DataSpace Archives remain open and non-proprietary.

References

[1] M. Beck, T. Moore, and J. Plank. An end-to-end approach to globally scalable network storage. In *ACM SIGCOMM 2002 Conference*, 2002.

[2] Robert Grossman and Marco Mazzucco. Dataspace - a web infrastructure for the exploratory analysis and mining of data. *IEEE Computing in Science and Engineering*, pages 44–51, July/August, 2002.

[3] Data Mining Group. Predictive model markup language (PMML). Retrieved from http://www.dmg.org, September 30, 2003.

[4] Data Mining Group. Web services for data mining. Retrieved from http://www.dmg.org, September 30, 2003.

[5] University of California at Irvine (UCI). UCI knowledge discovery in databases archive. http://kdd.ics.uci.edu/, retrieved on September 4, 2003.

[6] Sean Rhea, Chris Wells, Patrick Eaton, Dennis Geels, Ben Zhao, Hakim Weatherspoon, and John Kubiatowicz. Maintenance-free global data storage, 2001.

[7] W3C semantic web. Retrieved from http://www.w3.org/2002/ws/, September 4, 2003.