Astronomical Data Analysis Software and Systems VI
ASP Conference Series, Vol. 125, 1997
Gareth Hunt and H. E. Payne, eds.

HARP—The Hubble Archive Re-Engineering Project

R. J. Hanisch, F. Abney, M. Donahue, L. Gardner, E. Hopkins, H. Kennedy, M. Kyprianou, J. Pollizzi, M. Postman, J. Richon, D. Swade, J. Travisano, and R. White

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, E-mail: hanisch@stsci.edu

Abstract. The Hubble Data Archive now contains in excess of 2.5 TB of HST data in a system of four optical disk jukeboxes. In addition to providing a WWW-based user interface and removing a custom I/O processor (see Travisano & Richon 1997), STScI has undertaken a high-level effort to improve the operating efficiency, reduce costs, and improve service to archive users. The HARP group is studying data compression, data segregation, large on-line disk caching, on-the-fly calibration, and migration to new storage media. In this paper, we describe the results of our cost-benefit analysis of these and other options for re-engineering the HDA.

1. Introduction

The Hubble Data Archive (HDA) contains over 2.5 TB of near-on-line data. The data are stored on 12-inch WORM optical disks (the OD drives and platters are manufactured by Sony). These disks are mounted in four optical disk jukeboxes (Cygnet). Each jukebox has 131 OD slots. With a storage capacity of 6 GB per platter, the total near-line capacity of the current HDA configuration is 3.1 TB. Data ingest and retrievals are managed by the Data Archive and Distribution System—DADS. DADS operates on a mixed VAX and DEC Alpha architecture cluster, and work is now underway to migrate the entire system to the Alpha OpenVMS environment. Data enter DADS via an FDDI link from the OPUS data processing pipeline. Data are written simultaneously to two optical disks: one that will be used in daily operations, and another than is set aside in a safe area as a backup. Two additional copies of the data are made for the Space Telescope-European Coordinating Facility and the Canadian Astronomy Data Center.

2. Challenges Facing the Hubble Archive

The 12-inch WORM media and optical disk jukeboxes are nearing obsolescence. The jukebox systems are expensive (approximately $150k each), the blank media are also expensive (approximately $300 each), and the optical disk drives are expensive to maintain. We expect this medium to be obsolete once DVD—Digital Versatile Disk—technology is available. We expect DVD media and robotics costs to be much lower. Of course, we have always known that we would have to migrate to new media at some point, and the architecture of DADS should allow for heterogeneous operations as we transition from 12-inch WORM to DVD. The installation of STIS and NICMOS during the 1997 Servicing Mission, and of ACS in 1999, will lead to a factor of 3–6 increase in the volume of data generated by HST. Without changing our approach to archiving, we will be filling optical disk jukeboxes at the alarming rate of one every four months. The alternative is to incur the expense of providing enough operations staff to manage the off-line disk handling and on-demand mounting.

The CADC and ST-ECF have implemented a new "on-the-fly" calibration (OTFC) facility for Hubble data. Using compressed, raw science data as the basic data source, data are calibrated as they are requested by archive users. By calibrating on-the-fly, the most recent or best calibration reference files can be used. Observers are therefore provided with a higher quality result than they might get from analyzing the data archived immediately following the standard pipeline calibration. OTFC is most effective for data from the WFPC and WFPC2, where flat fields and dark-current corrections are often improved in the weeks following an observation. Since the pipeline calibration is normally performed within 24 hours after data are taken, the calibrated data in the archive for WFPC and WFPC2 are almost all sub-optimal.

3. Archive Efficiency Improvements

The HARP team has identified a number of areas in which the efficiency of the Hubble Archive can be increased and costs can be reduced. Areas that are now under consideration for implementation in the coming year are described below.

Data Compression. Currently, data in the HDA are not compressed. Raw data from the current HST instruments can be compressed losslessly by at least a factor of three, and perhaps as much as a factor of ten (depending on instrument and observing mode). Overall, including calibrated science data and engineering data, we expect to achieve compression ratios of about three. Lossy compression techniques can achieve much higher compression ratios, and for WFPC in particular, we are considering at least a simple rounding of low-order bits in the calibrated (floating point) data. Allowing users to retrieve compressed data also eases network loading. Data compression ratios for STIS and NICMOS have not yet been evaluated thoroughly, though our current expectation is that STIS data will compress quite well, and NICMOS data may not compress at all well (owing to strongly non-uniform backgrounds in the IR). Owing to its good overall efficiency, availability of source code, and ubiquitous use in the community, we will use the gzip compression algorithm.

Data Segregation. All types of HST data—raw telemetry, engineering data, and science data—are written to the currently open optical disks. Infrequently accessed data, such as raw telemetry, are intermixed with frequently accessed science data. By writing different types of data to different storage optical disks, or different storage devices, the infrequently accessed data could be moved to off-line storage with little impact on operations. Data segregation has already been implemented via a semi-automated procedure in archive operations.

Secondary Load. Data compression and data segregation can be combined to rewrite the existing archive contents onto a new set of ODs in which all data are compressed, and all infrequently accessed data are not copied. The existing disks would be moved out of the jukeboxes and replaced with disks contained the compressed, segregated data. If off-line data are requested, they can still be provided via the operator. The lifetime of the existing four OD jukeboxes can be extended to well into 1998 if we begin Secondary Load by spring of 1997.

On-the-Fly Calibration. An OTFC facility provides archival researchers with a "best" calibration (the best that has been determined by the time the data are retrieved from the archive) without requiring the researcher to go through a tedious and complex recalibration process themselves. The CADC and ST-ECF have implemented such a facility, and by storing only compressed, raw science data and generating calibrated data on demand, they have reduced the archive of currently public HST data to a volume that fits on ~60 CD-ROMs (Crabtree et al. 1996). The greatest benefit in OTFC, at least as far as efficiency of storage is concerned, comes from WFPC data. Raw WFPC and WFPC2 data compress quite easily by a factor of ten, while the process of calibration increases the total volume of data by another factor of ten. GHRS and FOS data are small in volume, and storing their calibrated data in the archive has virtually no effect on overall storage requirements. FOC data are somewhat larger, but the instrument is used less frequently, and has much more stable calibrations. The computational load of OTFC for the existing HST instruments is quite manageable on a high-end workstation. The case for OTFC is not so clear for STIS and NICMOS. In both cases the calibrated result, i.e., extracted spectra or combined images, is a more compact product. The calibration process requires having large numbers of raw data files available simultaneously, and the CPU requirements are considerably greater than for WFPC and WFPC2.

The support staff at STScI is often required to answer questions from HST guest observers (GOs) concerning the quality of, or artifacts in, their data. Answering such questions requires that the STScI staff has an identical version of the data sent to the GO, or that an identical version can quickly be generated. OTFC can potentially yield different results every time it is used, depending on how frequently the calibration reference files or calibration algorithms themselves are updated.

The HARP team has not yet reached a conclusion on the advisability of OTFC as a means for increasing archive efficiency. It is clearly valuable as a user service, but aside from WFPC and WFPC2, where we can probably achieve sufficient efficiency simply by compression of both raw and calibrated data already in the archive, OTFC does not obviously reduce the overall HDA data volume.

Large Disk Cache. As currently implemented, DADS does not cache data that are frequently accessed or are likely to be accessed. As a result, the optical disk jukeboxes are exercised at a maximum rate, and users must wait periods from minutes to hours before their data are be retrieved. A large disk cache would provide more immediate access to frequently accessed data.

Unfortunately, data access patterns for the HDA are not very simple, and cache implementation is not straightforward (Comeau 1996, Comeau & Park 1997). Current usage patterns indicate that a preloaded cache with a capacity for at least one month's worth of science data (~20–30 GB) would streamline archive performance. Additional disk space would be required in support of data verification and to provide an intermediate time-scale back-up facility.

4. Summary—HARP Study Recommendations

The primary goal of the HARP study group is to reduce HDA operational costs without compromising user service. The key recommendations for achieving this goal include:

- **Segregate science and engineering data.** Engineering data comprise about 1/3 of the archive data storage but are accessed only rarely.
- **Compress all data.** Use gzip on all data sets, both raw and calibrated. Some data sets will not compress very well, but this will apply a standard utility and allows users to retrieve compressed data and uncompress it on their end, thus minimizing network loading.
- **Implement secondary load.** Segregate and compress all data in the existing archive in order to move infrequently used data off-line and conserve jukebox slots.
- **Migrate to DVD technology.** Begin work now on supporting a heterogeneous archive, and eventually migrate all existing data to DVD. CD-ROM will be used as a prototyping medium. We prefer not to transition first to CD-ROM and then to DVD, to avoid operating an archive with three distinct types of media.
- **Implement a Disk Cache.** Handle data retrieval requests that occur within one month of ingest from a preloaded cache rather than requiring full retrieval from the jukeboxes.
- **Implement On-the-Fly Calibration as a User Service.** Provide OTFC to make it easy for users to recalibrate data using the most current calibration reference files and software.

References

Comeau, T. 1996, in ASP Conf. Ser., Vol. 101, Astronomical Data Analysis Software and Systems V, ed. G. H. Jacoby & J. Barnes (San Francisco: ASP), 497

Comeau, T., & Park, V. 1997, this volume, 318

Crabtree, D., Durand, D., Gaudet, S., & Hill, N. 1996, in ASP Conf. Ser., Vol. 101, Astronomical Data Analysis Software and Systems V, ed. G. H. Jacoby & J. Barnes (San Francisco: ASP), 505

Travisano, J., & Richon, J. 1997, this volume, 286