Number 16 August 1991

ST—ECF Newsletter

Space Telescope — European Coordinating Facility

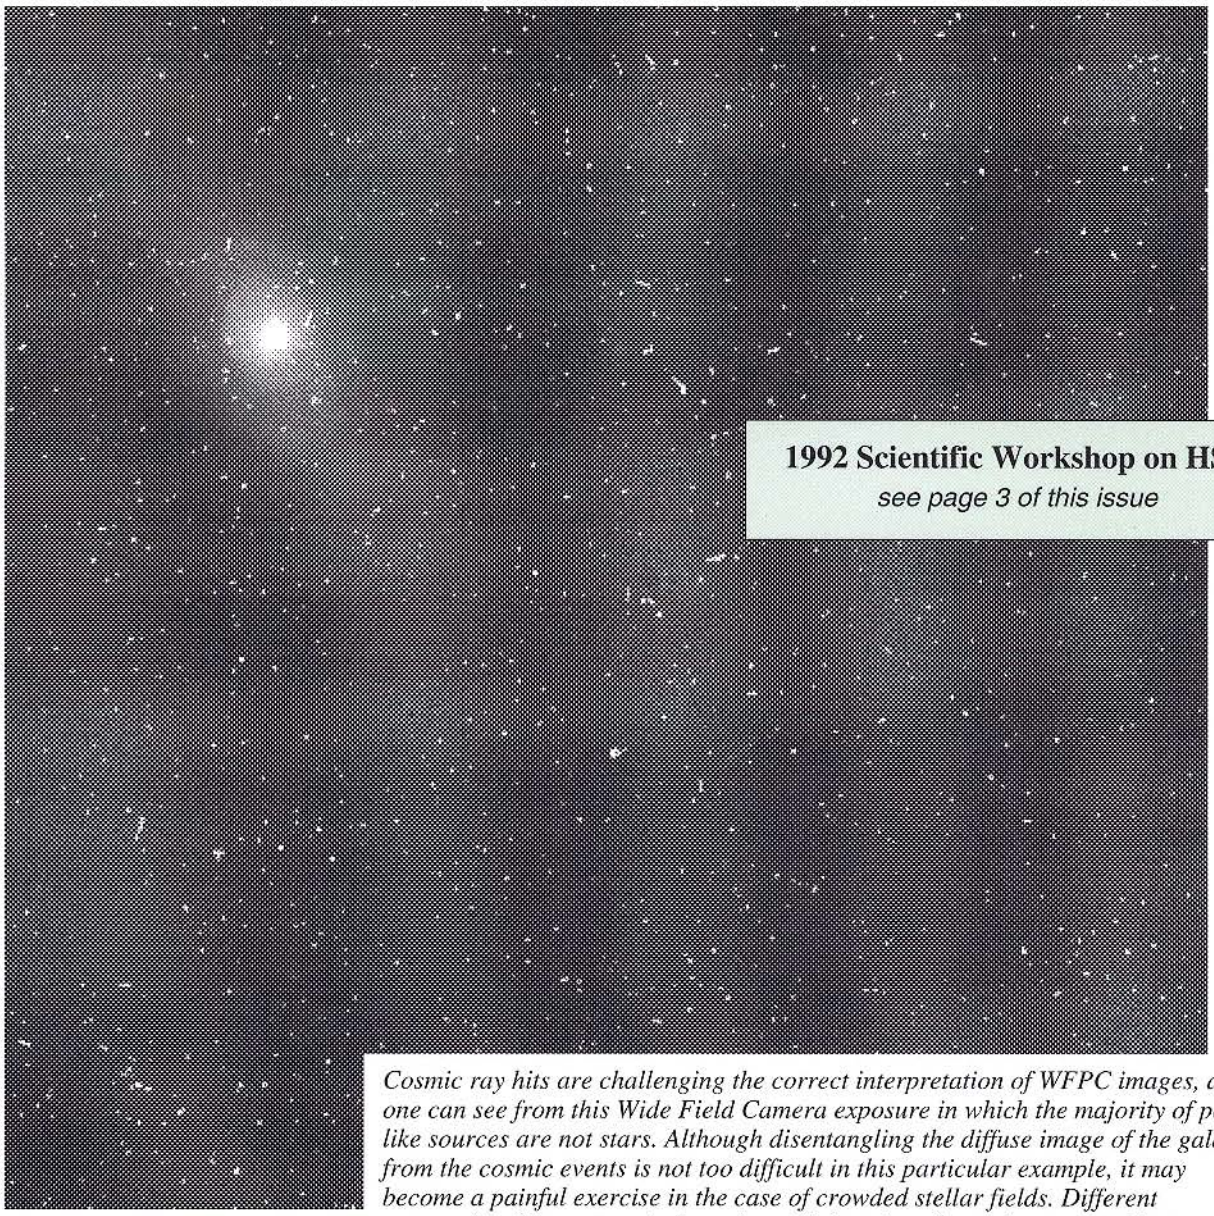

1992 Scientific Workshop on HST

see page 3 of this issue

Cosmic ray hits are challenging the correct interpretation of WFPC images, as one can see from this Wide Field Camera exposure in which the majority of point-like sources are not stars. Although disentangling the diffuse image of the galaxy from the cosmic events is not too difficult in this particular example, it may become a painful exercise in the case of crowded stellar fields. Different approaches for automatic detection and flagging of cosmic rays are discussed in this issue.

HST in orbit — fifteen months

HST has entered its operational phase as an observatory. The archive now contains a considerable number of scientific observations, many of which are in the public domain. There is concern about a number of hardware failures in the spacecraft, some of which may be due to the high levels of solar activity.

Robert Fosbury

After some fifteen months in orbit, the HST is now in what we might call a 'normal' operational phase. Science Verification is winding down and responsibility for the scientific operation has passed from NASA to the STScI. The efficiency of the observatory is already approaching its design value (see page 18 of this issue) and the archive has suddenly become a real resource for research. This became especially so at the beginning of July when the existing SV data—containing many observations of scientific interest—entered the public domain (see page 15 of this issue).

The preparations for the refurbishment mission in late 1993 are proceeding. The WFPC II team are investigating the need for active adjustment of the fold-mirrors to ensure alignment of the OTA pupils with the WFPC optics. This involves a major engineering effort and the current concern is that the additional work may cause a considerable delay. Alternative optical designs for the corrective elements of WFPC, which would not be so sensitive to misalignment as the current one, are under active investigation. COSTAR passed its critical design review in May and is under construction according to schedule. The next critical review will take place in December.

Modifications to the pointing control software have improved the response of the spacecraft pointing to the solar array-induced oscillations induced by thermal effects during terminator crossings. This has reduced the jitter amplitude during this phase by a factor of two and greatly improved the robustness of FGS 'fine-lock'. The focal-plane geometry is now well measured and the success of target acquisitions is high. The recent inclusion of velocity aberration correction (for the orbital motion of the spacecraft) into the pointing software means that the drift rates under gyro-only pointing control should be much lower than before. This introduces the opportunity to make short exposure 'snapshot-mode' observations with a greatly reduced acquisition overhead.

The spacecraft power and thermal systems are working well, indeed there is so much power available from the solar arrays that the instruments can be operated with less thermal cycling than anticipated and there are no unexpected constraints on operations due to thermal effects.

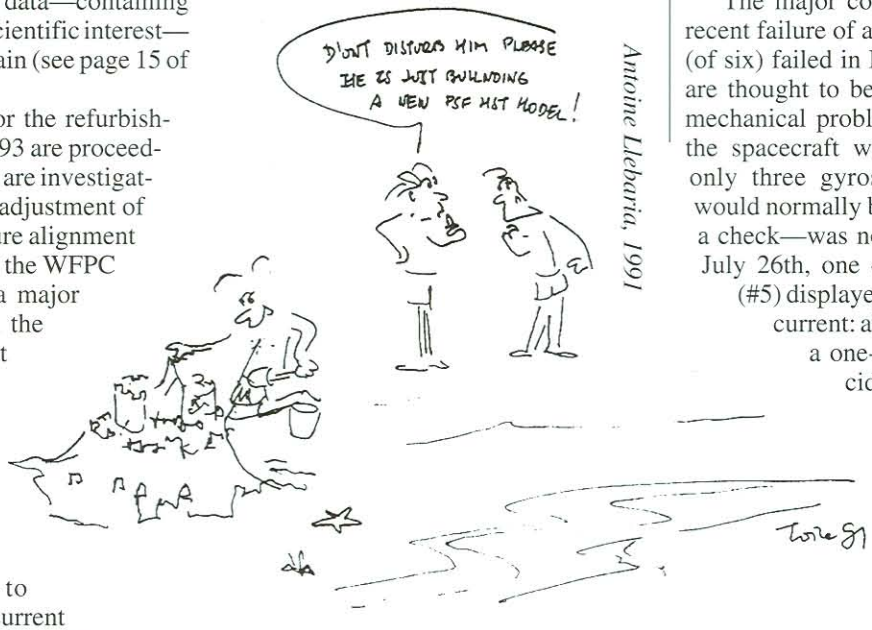

Antoine Llebaria, 1991

During late June and early July, there were some problems with telemetry and telemetry monitors on the spacecraft. These are thought to be due to periods of high solar activity and the susceptibility of some of the on-board electronics to particle hits. Apart from this communications have been excellent with bit-error rates some three orders of magnitude better than specification.

A failed memory unit in the on-board computer caused the HST to enter a deep safemode state on 2 May. This was successfully diagnosed, the computer reconfigured and the observing schedule intercepted by 7 May in what must be considered a very efficient recovery operation.

Both the FOS and GHRS recently experienced safing events: in the FOS case the cause was traced to a software problem and very easily solved, for the GHRS the malfunction seems to be more serious, involving one of the power units. At the time of writing the problem is still under investigation.

The major concern, however, is the recent failure of a second gyro—the first (of six) failed in December 1990. These are thought to be electronic rather than mechanical problems. For a few weeks the spacecraft was left operating with only three gyros. The fourth—which would normally be in the control loop as a check—was not yet switched on. On July 26th, one of the operating gyros (#5) displayed a slightly out-of-limit current: although this may just be a one-time glitch, it was decided to turn the fourth spare gyro on for better safety. The two failed gyros are from the four refurbished IUE units on-board HST (the other two have a different origin). To anticipate the possibility that HST is left with only two operating gyros before the first refurbishment mission, software is being written to operate with two gyros and the sun and star sensors—similar to IUE's current operating mode. NASA is also considering the possibility of an early emergency mission: the decision to proceed or not with such a plan will be taken at the end of August. Hot news on this subject will be posted, as usual, on our electronic Bulletin Board (see page 19 of this issue for the updated procedure to access it).

A Digital All-Sky Survey

The digitized Schmidt plates which served for the construction of the Guide Star Catalog may soon become available to the community. The cost of a copy of the Digital All-Sky Survey will depend on the number of orders.

Riccardo Giacconi † *& Michael Shara* †

As part of the effort to construct the Guide Star Catalog (GSC), STScI has digitized Schmidt plates covering the entire sky. The plate collection consists of materials from the UK Schmidt in Siding Spring, Australia, operated by the Royal Observatory Edinburgh until June 1988, thereafter by the Anglo Australian Observatory, and from the 48-inch Oschin telescope on Palomar Mountain, operated by the California Institute of Technology. These digitized scans (of order 10^{12} bytes of data) are stored on optical disks at STScI. Details of the digitization program, in particular the survey and scan characteristics, are described in The Astronomical Journal (Volume 99, page 2019, 1990).

The purpose of this announcement is to ask if your institution/research group might be interested in purchasing a copy of the digitized sky. The northern sky images that will be offered are from the original POSS E (red) plates. The southern sky will be covered by the SERC J survey. In addition we expect to add other digitized surveys to this collection as appropriate arrangements are completed.

With the cooperation and encouragement of Caltech, the National Geographic Society, the UK Science and Engineering Research Council and NASA, the STScI is currently planning to distribute moderately compressed images of the digital scans. Extensive series of tests have shown that essentially no astrometric or photometric information is lost through the compression and decompression processes. Relative positional accuracies significantly better than 1 arcsecond, and stellar magnitudes accurate to better than 0.5 magnitude are routinely obtained except near plate flaws and edges. The compression algorithms and tests were reported at the Digital Optical Sky Surveys Conference in Edinburgh this past June and will also appear as an ST ScI preprint.

The breadth of the community interest will enable STScI to decide whether to proceed with the survey distribution and to better estimate the costs of production. Any distribution of digitized surveys by STScI will be done on a cost-recovery basis.

The per-copy cost of the all-sky digitized survey will depend on the number of copies ordered. If demand is about 100 orders, then the survey cost should be approximately US $ 6000 per complete set, and if 500 orders are received, the cost per set should be under US $ 1500.

Of the several different mass storage and distribution media in widespread use, CD ROMs (Compact Disk Read Only Memory) appear to be the most stable and cost effective technology available; and we propose to distribute the survey as a set of about 100 CDs. Software to read, to decompress and to display sky images on some workstations, and to obtain coordinates from the images will also be distributed free of charge to survey customers.

If your institution/research group might be interested in purchasing a CD copy of the digitized sky survey please send a message (which in no way commits you to a purchase) to Michael Shara, ST ScI, 3700 San Martin Drive, Baltimore, MD 21218, USA (E-mail: scivax::shara, shara@stsci.edu, shara@stsci.bitnet). Be sure to indicate the type of workstation and operating system you would prefer to use with the digitized scans.

† STScI

STScI NEWSLETTER

The Space Telescope Science Institute publishes a Newsletter at regular intervals (3–4 times per year). The STScI Newsletter contains information of interest to proposers, including updates on the status of the HST and its instruments. Subscriptions are available at no cost to all interested scientists; requests to be added to the mailing list should be sent (by regular or electronic mail) to the User Support Branch at the following address:

User Support Branch
Space Telescope Science Institute
3700 San Martin Drive
Baltimore, MD 21218
USA

E-mail: scivax::usb (SPAN)
usb@stsci.edu (Internet)
usb@stsci (Bitnet/Earn)

Requests should also specify whether the subscriber wishes to receive future Calls for Proposals.

1992 HST Scientific Workshop Preliminary Announcement

The ST-ECF and the ST ScI are planning a workshop to be held from June 29th to July 7th, 1992, in Sardinia, Italy, to discuss the scientific investigations possible with HST in its early years, concentrating on optimal observing and data analysis strategies. We hope to present results from the first year of the General Observer programme, prior to the proposal deadline for Cycle 3. Details will be announced in the next Newsletter.

Piero Benvenuti and Ethan J. Schreier

Cosmic Ray detection on WFPC images

The detection and flagging of cosmic ray hits on WFPC frames is an important part of data processing. Some relatively straightforward approaches are described here which might be considered before incorporating cosmic ray detection into one of the iterative restoration schemes.

Fionn Murtagh &Hans-Martin Adorf

Background

Handling cosmic ray hits on astronomical images may be tackled in two ways. The first is simply to flag the associated pixels as bad data and subsequently to avoid these in model fitting or the application of other algorithms. Alternatively, pixels affected by cosmic ray hits (CRs) can be identified by the poorness of a model fit or as a feature incompatible with the point spread function used in a restoration. The very large number of CRs on HST images has led us to investigate the first approach. Investigations involving simultaneous image restoration and CR detection have been, or are being, carried out by N. Weir (CalTech) and L. Lucy (ST–ECF). The direct search methods for CRs reported here are simple and quick. We have sought to avoid imposing models on the data and have tried to take as much information into account as possible when making the choice between CR and non-CR. A number of practical approaches arise out of these studies.

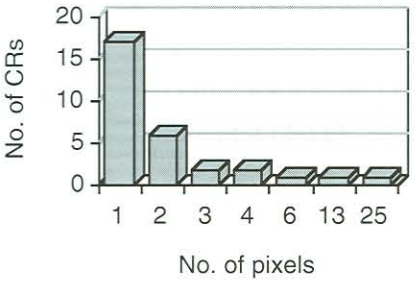

Figure 1: From one set of 87 identified cosmic rays, this shows the number of contiguous pixels per cosmic ray.

The likelihood of Wide Field and Planetary Camera (WFPC) image pixels being contaminated by cosmic ray hits is far greater than would be the case for ground-based detectors. It has been estimated that in a 45 minute WFPC exposure, 10 to 25% of all pixels will be hit by a cosmic ray (Holtzman & Westphal, in Adorf, 1991). The extensive contamination due to such CRs makes it imperative to detect them during data reduction.

For our purposes, cosmic rays can be divided into two classes: high energy 'primary', and low energy 'secondary' particles. High energy cosmic rays pass straight through the detector. On average, they cause electrons to be deposited in two or three pixels. In one set of 87 CRs, the numbers of contiguous pixels per CR are shown in Figure 1.

Eighty to one hundred electrons are deposited per μm travelled through a pixel. A WFPC CCD pixel has approximate dimensions of 15×15×8 μm in length, breadth and thickness. One of the Planetary Camera chips (PC8) is, however, notably less thick. Depending on the traversal angle, between 8 and 20 times (say) 80 electrons are deposited in a pixel. Diffusion along the path traversed by the high energy cosmic ray also takes place but is sufficiently small to be ignored. A conversion factor of 7.5 brings the electron counts to ADU used in the image data.

Secondary particles are produced by cosmic rays hitting the surrounding spacecraft components. In particular, the detector housing and 'shielding' give rise to these low energy secondaries and the number of electrons deposited by them can be very high.

It was originally foreseen and recommended (see Griffiths 1990, Lauer 1989) to find CRs by aligning two frames taken in 'split mode', i.e. repeated observations. However, due to HST's current less-than-perfect pointing characteristics, accurate re-alignment of non-aligned frames would require resampling of one of the frames which, due to the coarse pixel grid, is problematic if not impossible (Adorf, 1990).

Other difficulties with CR-splits include additional overhead time for read-out and pre-flash and the additional read-out noise. Consequently, in addition to the existing approach working on frame pairs, there is a perceived need for a procedure which, given a pixel and some of its neighbourhood, reliably detects CRs on individual WFPC frames and discriminates them from similar looking star images.

The problem of discrimination

CR/non-CR discrimination is a typical pattern recognition problem. Two broad approaches are feasible: feature (or astronomical object) based and image (or pixel) based. For a feature-based classifier one has to define a good set of features describing the image context in the region of interest surrounding the pixel in question. A trainable classifier is then trained on these features together with the correct answers 'CR' and 'non-CR' to be provided externally. An image-based classifier works directly on the pixel data in the region of interest (see Adorf, 1991).

What we have called the feature-based approach is investigated later in this article and in a number of different ways. What has been called the pixel-based approach is considered first.

A 'pixel-based' approach to classifier design

The method employed was to: (1) have cosmic ray hits on images flagged by a human expert ('ground truth'); (2) determine 3×3 subimages which potentially could be CRs or non-CRs; (3) use this information to train and test various classifiers; and (4) apply the classifiers to other images. Calibrated WFC images of the open cluster NGC 1850, with 10 sec exposure and without any preprocessing, were used in the trainable classifier experiments. In addition, other images of the faint Lynx galaxy field were made available by Ivan King (FOC IDT) and a marked version of part of this by Nick Weir (CalTech, Pasadena). Jon Holtzman (WFPC IDT), who initially motivated this work, marked the same 200×200 image on two different occasions, separated by several months. The second marking session used an 800×800 image from which the 200×200 image was extracted. The first frame had 85 cosmic ray pixels marked and 102 in the second, with 74 common to both. If the first marking is taken as ground truth, this corresponds to a redetection rate of 87%, whereas if the second marking is regarded as ground truth then the redetection rate drops to 73%.

For this study we used two variants of a neural network method: the multilayer perceptron (MLP) and the well established k-nearest neighbours (k-NNs) discriminant analysis method. They are both non-parametric techniques and may be viewed as providing piece-wise linear

separation between the CR and non-CR classes in parameter space (here, the 9-dimensional space associated with the 3×3 subimages).

Training, testing and assessment of classifiers

Details of the classifier training and testing—which are quite general and can be employed with classifiers in any domain—can be found in Murtagh & Adorf (1991). The performance of the classifier can be represented as follows: a cosmic ray (however this is defined, e.g. a pixel corresponding to part of a cosmic ray hit) can be (i) detected or (ii) missed by the classifier.

Similarly, a star can be properly rejected as not being a cosmic ray, or can occasion a false alarm. Considering CRs, the rate of detection (RD) is an estimate of the *probability of detection*, $P(CR_c \mid CR_t)$ where CR_c is a CR determined by the classifier, and CR_t is a region of interest classified as a CR by the 'teacher'. We wish to maximize RD. The rate of false alarms (RFA) is an estimate of the *false alarm probability*, $P(CR_c \mid nonCR_t)$. We wish to minimize RFA. However, the quantities RD and RFA are not independent and the *operating characteristic* (*OC*) diagram is the plot of RD, RFA couples. Each point defines an operating point of a classifier.

Figure 2 presents the operating characteristic diagram found in assessing various classifiers. Discussion of computational and other properties of these classifiers may be found in Murtagh & Adorf (1991). Note the curves traced out by the different methods: these reflect the trade-off inherent in maximizing RD while minimizing RFA. The OC diagrams shown here could be used to implement various cosmic ray detection strategies. One possibility, for instance, would be to hold the RFA value fixed, at some upper acceptable limit, and find the best RD value corresponding to this.

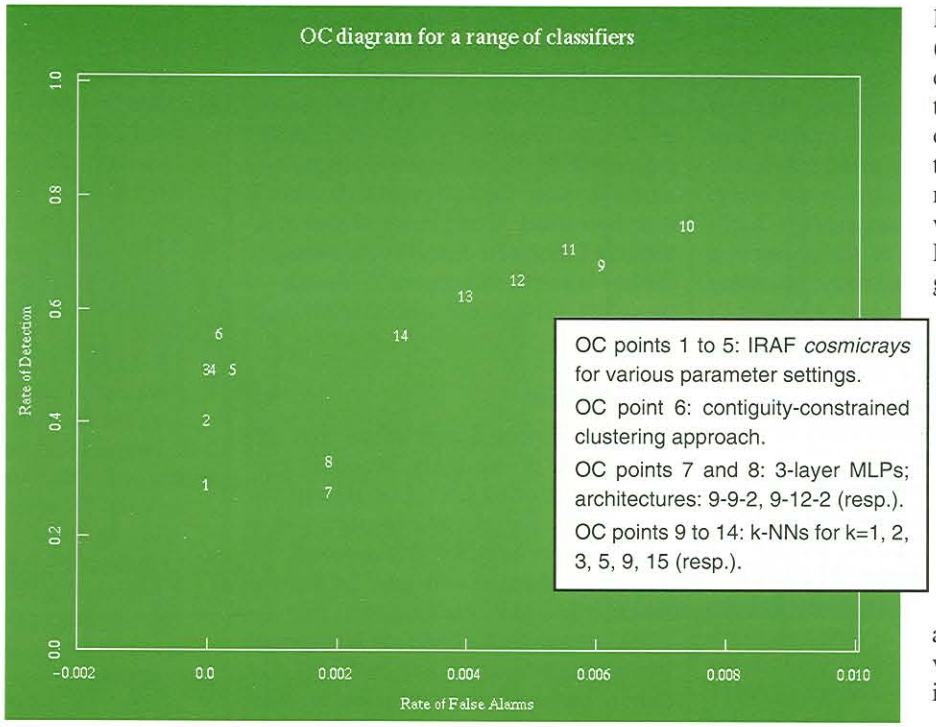

Figure 2: Operating characteristic diagram for application of various trained and non-trained classifiers. Note that most of the results appear on an approximate curve in (RFA, RD)-space.

'Feature-based' approaches to classifier design

Murtagh (1991b, 1991c) used a clustering of stellar image data, discretized in intensity space, where the clusters correspond to star-like or CR-like registrations. Such clusters of discretized points in x/y/intensity space can, with little effort, be contiguity-constrained. In fact, clusters which are not contiguous would make little sense in this instance. In line with the use of 2-dimensional photometry packages described below, the clusters may be characterised on the basis of sharpness and roundedness attributes. Subsequently, straightforward discriminant methods may be applied in attribute space to objectively separate CRs from non-CRs.

Interactive 3-dimensional graphical portrayal of the data is supported by many statistical

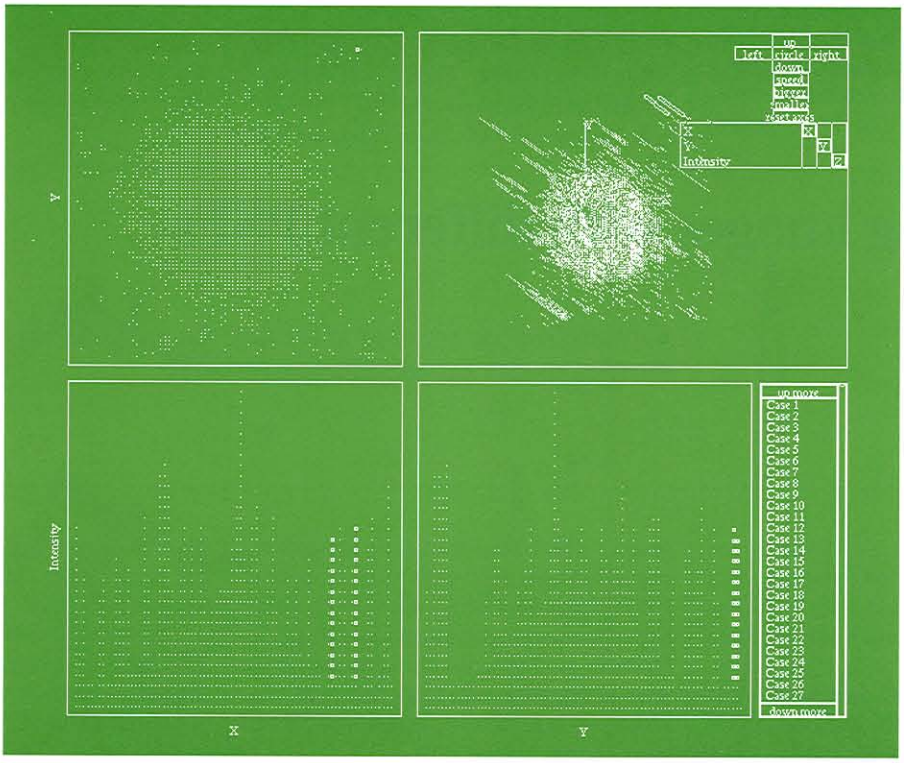

Figure 3: Three views of (discretized) image:
- *upper left: thresholded in intensity,*
- *lower left: intensity profile along X-direction,*
- *lower right: intensity profile along Y-direction.*

Upper right: spinnable 3-D view of image.
Brushed parts of image simultaneously highlighted in all views. Note how an apparently single CR in lower right view is shown to be two CRs in lower left.

packages. The image represented in the discretized x/y/intensity space can be interactively rotated by the user. This allows a multiplicity of different views of the image to be obtained quickly and easily. Highlighting suspected CRs in one view of the image and having the corresponding points instantaneously highlighted in other views of the image is a powerful aid in interactively assessing what a CR is or is not. Figure 3 presents a multiple view window, with a suspected CR highlighted in the y/intensity view. Note from the highlighting in the x/intensity view how this suspected CR is seen to represent two objects.

A different approach to CR detection involves basic image processing operations on the original image. CR hits are sharply peaked, and hence the interest in applying a gradient or edge enhancement filter. Using the Sobel operator was certainly found to emphasize CRs. Some results are described in Murtagh (1991a). It is questionable though how feasible such an approach is in the presence of undersampling.

Two-dimensional photometry packages are based on object finding or fitting. Attribute values for the objects produced by these packages could provide useful discriminants between CRs and non-CRs. A problem with the explicit attempt to detect CRs using such packages is unfortunately that most of the well-known packages do their level best to avoid the user ever worrying about CRs! In MIDAS's INVENTORY package, local averaging is used to make object detection robust with regard to CRs. The DAOPHOT-adapted IDL FIND routine was used (Landsman & Isensee 1990, Stetson 1991) with some success, to produce sharpness and roundness values for a large number of objects. Straightforward discriminating curves in the sharpness-roundedness attribute space then allowed CR/non-CR discrimination. Using more sensitive attributes for objects found could certainly further improve this approach. Results of using IDL FIND are described in Murtagh (1991a).

Conclusion

The work reported here arises from the inadequacy of currently available routines for detecting CRs. Our algorithmic research has the goal of arriving at one or more routines to tackle this problem. Alternative schemes will also be kept in mind. It is clear, for instance, that the more knowledge that one can incorporate into the discrimination task the better. The lines of inquiry followed to date are relatively undemanding in computer time and are quite simple.

Acknowledgements: We are particularly grateful to J. Holtzman. We also thank I. King, N. Weir, R. Hanisch, P. Stetson, J.MacKenty, and others for discussion, suggestions, and 'ground truth' data.

References

1. H.-M. Adorf, "Coping with the undersampling of WFPC images", *ST-ECF Newslett.*, No. 12, 9–10, 1990.
2. H.-M. Adorf, "Cosmic ray detection on individual WFPC frames", ST-ECF TR (DA)-21, 5 pp., April 1991.
3. R. Griffiths, *Hubble Space Telescope Wide Field and Planetary Camera Instrument Handbook*, Version 2.1, Space Telescope Science Institute, Baltimore, 1990. See also: "Hubble Space Telescope Cycle 2 Proposal Submissions: Updated Technical Information", Enclosure 3: Wide Field/Planetary Camera Update, Space Telescope Science Institute, 23 May 1991.
4. W. Landsman and J. Isensee (1990 and 1987), IDL program FIND, UIT Library of IDL software, Goddard Space Flight Center.
5. T.R. Lauer, "The Reduction of Wide Field / Planetary Camera Images", *Publ. Astron. Soc. Pacific*, **101**, 445–469, 1989.
6. F. Murtagh (1991a), "Detecting cosmic ray hits on WFPC frames using gradient operators and 2-d object finding", ST-ECF TR (DA)-25, 16 pp., June 1991.
7. F. Murtagh (1991b), "Contiguity-constrained clustering for morphological analysis of objects in digitized images", ST-ECF TR (DA)-23,15 pp., May 1991.
8. F. Murtagh (1991c), "Determining cosmic ray hits on WFPC images through peak detection", ST-ECF TR (DA)-22, 12 pp., May 1991.
9. F. Murtagh and H.-M. Adorf, "Detecting cosmic ray hits on HST WFPC images using neural networks and other discriminant analysis approaches", in V. Di Gesù et al., Eds., *Data Analysis in Astronomy IV*, Plenum Press, New York, 1991, in press. ST-ECF TR (DA)-24, June 1991.
10. P.B. Stetson (1991), *DAOPHOT II: The Next Generation. User's Manual*, DAO, Victoria, BC, 1991. (Chapter C.V.)

Co-adding images with different PSFs

Synthetic images are used to illustrate the performance of a technique for merging images formed with different PSFs.

L.B. Lucy

In their quest for high S/N images and spectra, astronomers often co-add or stack separate exposures made with the same equipment. With each data set weighted according to exposure time, the final co-added image is regarded as equivalent to a single exposure with the same total time and is still accorded the status of basic data.

An important variant of this problem arises when the instrumental set-up varies from exposure to exposure. For ground-based imaging this effectively occurs already simply due to the variability of the seeing, which results in each image having a different PSF both with respect to FWHM and to mathematical form. More extreme examples of image-to-image differences in PSFs obviously occur if we wish to combine data from different telescopes — in particular, ground-based and space-borne instruments.

A general procedure for treating this variant of the co-adding problem has been developed and successfully tested for 1-D images. The technique gives a final co-added image that combines the resolution of the best image with the S/N of the total data set. In addition, it can be argued that the final image should still be accorded the status of basic data.

An example is shown in Figs. 1 and 2. The field — a single star on a flat background — is observed with two telescopes having gaussian PSFs with $\sigma_1 = 0.1$ and $\sigma_2 = 1.0$; and telescope 2 acquires 5 times the counts of telescope 1. The two raw images are plotted in Fig. 1, with the noise being pure poissonian. The co-added image is shown in Fig. 2, where it is compared with the short exposure, high resolution image obtained with telescope 1 —

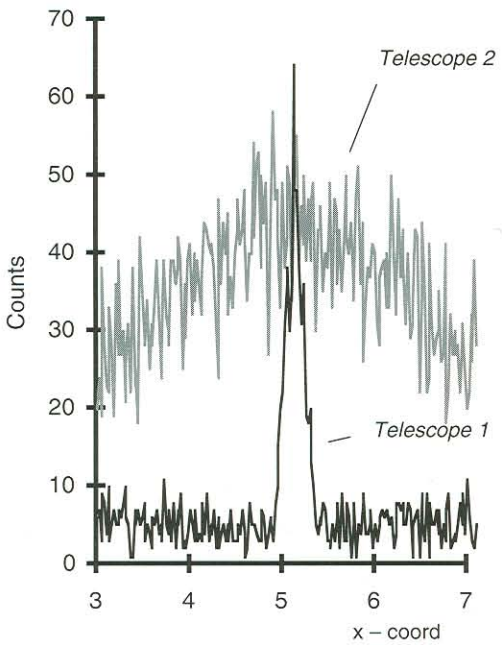

Fig. 1. Images of star on constant background obtained with telescopes 1 and 2.

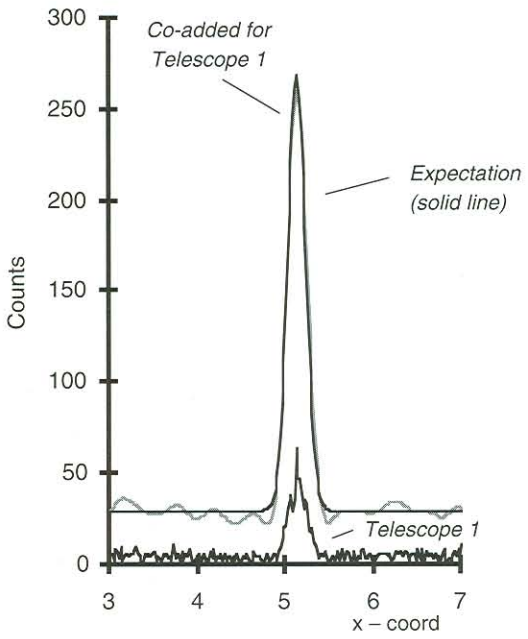

Fig. 2. Co-added version of images in Fig. 1 compared with expectation and with the raw image obtained with telescope 1.

note the changed vertical scale. The high resolution and high S/N of the co-added image are evident. Note also that comparison with the expected image shows that artifacts due to the adopted mathematical procedure are absent or negligible.

The above example illustrates the possibility of combining images from ground-based (low resolution, but many photons) and space-borne (high resolution, but few photons) telescopes. A second example illustrates the possibility of exploiting the variability of seeing at a terrestrial observatory. In Fig. 3, three identical exposures of a double star are shown but, because of changes in the seeing, the gaussian PSFs have σ's of 0.5, 1.0 and 1.5 seconds of arc. Only the best seeing image reveals unambiguous evidence of the object's true nature.

In Fig. 4, the co-added version of these three images is compared with the high resolution image ($\sigma = 0.5$) and with expectation. Again the results are distinctly encouraging.

Tests with synthetic 2-D images as well as with real astronomical data are planned in the near future.

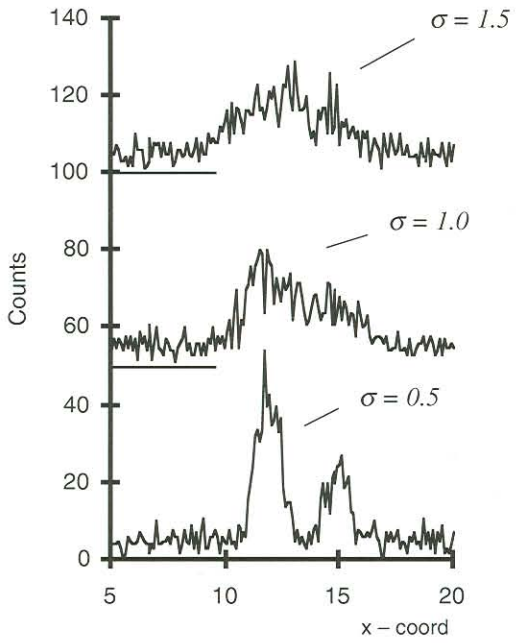

Fig. 3. Images of a double star taken under three different seeing conditions. Plots have been offset for clarity.

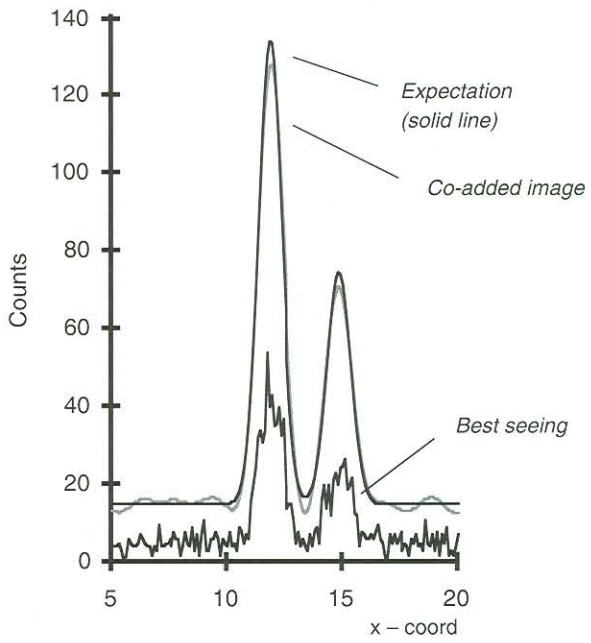

Fig. 4. Co-added version of images in Fig. 3 compared with expectation and with the best seeing image.

Recent developments in Maximum Entropy-based image restoration

Nicholas Weir[†]

Introduction

At the simplest level, image restoration can be thought of as the task of finding the most appropriate estimate of a true image depicted in data. The key to this process is defining exactly what we mean by the term 'appropriate'.

To be appropriate, for example, we should demand that the restored image be somehow compatible with the data when projected into data space. Note the problematic use of a fuzzy word here, 'compatible', which must be defined below. Also note the careful distinction made between the notion of 'data space' and what we will refer to as 'image space'. Our data, which in the case of direct imaging happen to appear in the form of images, need not appear as such in the general framework we wish to establish. To restore to an image, we must understand the process by which quantities in restored image space may be projected into data space where we are able to perform our compatibility tests with the data. In the case of direct imaging, this projector operator is often in the form of a convolution with a point spread function (PSF), which may practically be accomplished with a couple of fast Fourier transforms (FFT). When the PSF is spatially variant, as in the case of WFPC images, a linear projector operator can still be devised but it may no longer be accomplished with a simple two-FFT convolution. General restoration schemes—such as the ones described here—that simply require the definition of this projector operator and perhaps its transpose, are therefore tied in no fundamental way to space invariant PSFs, or even to data sets already in the form of images. These characteristics are simply generally desirable for the speed and ease with which the projection can be computed.

The compatibility test we refer to is generally realized in the form of a likelihood function. Given the known or estimated noise characteristics of a given data set, this function is usually well-defined and provides a quantitative assessment of the degree of misfit between the data and a projected image—hereafter referred to as 'mock' data.

To first order, it may seem reasonable to stop here and simply equate what we have called 'compatibility' with 'appropriateness'. In other words, to produce a restored image, we simply find the image by maximizing the likelihood function and so producing the smallest projected misfit with the data. Unfortunately, of course, life is rarely that simple. In the presence of even the least amount of noise, this inversion problem becomes ill-posed, meaning the maximum likelihood solution is either non-unique or particularly sensitive to the noise. One may devise iterative restoration techniques which essentially converge (typically in a steepest descent fashion) toward a maximum likelihood solution but stop before reaching convergence in order to avoid producing an artifact-ridden image. One such method, which also incorporates a positivity constraint by construction, is the popular Lucy-Richardson iterative approach. Another is the related method employed by Núñez and Llacer and described in the March 1991 *ST–ECF Newsletter* (No. 15, p6).

While reasonably quick and effective, and rationalized on pragmatic grounds, all such iterative methods come with some noticeable drawbacks. For those more concerned with mathematical purity than practical efficacy, the primary complaint is that they are *ad hoc*. They maximize no well-defined objective function, producing results that are justified only in the sense that the heuristic approach by which they were produced seems reasonable. Practically speaking, this may be an important deficiency if one is interested in things like reliable error estimates for the restoration. For those strictly practically minded however, the argument against these methods is more straightforward: we can do better. More information may be buried in the data than can be mined using these techniques. The price one must pay to extract the extra nuggets is computer time—in some cases, lots of it. And in some instances the extra effort may not be worth it.

It is useful to have an intuitive feel for why there might be more information available than is recoverable by these techniques. Ultimately it boils down to a matter of signal to noise. High S/N regions in an image respond to deconvolution to higher levels of resolution than lower S/N features. The key to deriving the maximum degree of reliable information from the image is in appropriately decoupling the high S/N / low S/N and high-resolution/low-resolution portions of the image. Maximum likelihood-based iterative methods fail to decouple these portions of the image effectively. Thus to avoid overfitting the low S/N regions, one is forced to cease the iterative fitting process well before the high S/N regions have been maximally deconvolved, or even appropriately fit. The residuals between the mock data and real data will demonstrate a significant bias in the high S/N regions where the image has not yet been well fitted. Thus we obtain a restoration which not only provides less than optimal resolution where we often want it most (i.e., where the signal is), but which fails to meet fully one of our defining characteristics of appropriateness: compatibility with the data. One may come close to meeting an overall global χ^2 constraint, but given the systematic biases, the residuals are by no means likely in any realistic sense.

The question then becomes one of further decoupling the image, enhancing resolution where appropriate, while suppressing the noise. One possibility is to add a bit of 'glue' or regularization to the iterative approach, making it that much harder for features to pop out in the restoration process. The hope is that one could iterate further towards a better fit to the data without fitting to the noise to the extent that results from maximizing likelihood alone. Mathematical purists tend to frame the arguments for their particular use and choice of glue in terms of Bayesian probability theory. The regularization term, or 'prior', they use is often of an entropic form. There are some well-founded mathematical and physical justifications for using a function of this form but we will avoid them and merely note that it is a reasonable measure of structure in an image. This prior/regulariser is added to the likelihood function to form the *ying* and *yang* of the restoration process. The new objective function now has terms which both reward and penalize the image for better fitting the data with higher resolution features. The obvious next question

[†] CalTech

is how to balance these terms.

A multiplier, typically denoted alpha, is usually attached to the entropy term and is what determines its relative importance. If alpha is very small, the objective function is dominated by the likelihood function; if it is very large, the entropy dominates. For any finite choice of alpha, there is a unique maximum to the objective function. The previous problem of deciding how far to iterate towards maximizing the likelihood now becomes one of choosing a value for this regularization constant. Once it is picked, the solution can be found deterministically.

The proper choice for alpha eluded people for some time. The historic rule of thumb when performing maximum entropy (MaxEnt) based restorations was to use that alpha such that the χ^2 misfit between the projected solution (given that alpha) and the data is equal to n, the number of data points. Many iterations were usually required to reach this alpha as it would involve finding the unique MaxEnt solution for many other alphas along the way. The theoretical justification for this choice of the regularizing constant was always shaky but the results were generally pleasing. With this method, the data could generally be fitted better, thereby obtaining higher resolution in high S/N regions with fewer artifacts in low S/N regions relative to the iterative methods discussed before. In addition, as long as a choice for alpha was made in advance, there was no longer the lingering question of when to halt the procedure.

The problem was far from solved, however. When using the $\chi^2 = n$ choice for alpha, the bias in the residuals, while smaller than before, was still significant. If alpha was pushed to smaller values, resulting in closer fits to the data, the bias would diminish and the resolution in the high S/N regions would go up—in some cases enormously. On the other hand, an increasing number of artifacts would now appear in the image. Alpha such that $\chi^2 = n$ clearly provided too much glue; zero alpha did not provide enough. The balance had yet to be achieved.

Developments from MaxEnt

Around 1989 two leading researchers in the use of MaxEnt-based methods, Steve Gull and John Skilling (1989), hit upon a means of determining alpha using nearly the same Bayesian probabilistic arguments as had been employed to justify the MaxEnt approach in the first place. They showed how one could solve objectively for the most probable value of alpha given a data set—the Bayesian choice—and use it to solve for the restored image. They also showed how to approximate the normalizing constant in the Bayesian-derived objective function, so that one could solve not just for the maximum—the peak of the probability distribution—but obtain an estimate of the size and curvature of the distribution as a whole. Their technique could return the most probable image (what we have the bad habit of calling *the* restored image) based upon an objectively determined alpha as well as an estimate of the reliability of that restored image for any area of interest. These were very significant advances in the general restoration field.

Restorations using this new, what Gull and Skilling termed 'Classic', MaxEnt approach tended to produce less biased residuals. This was an expected consequence of the fact that the Bayesian choice for alpha always turns out to be less than or equal to the historic $\chi^2 = n$ value. For images with high S/N, high-resolution structure such as bright stellar fields, the Classic MaxEnt results were excellent. For more general types of scenes, however, this approach was still inadequate. Given the lower, Bayesian choice of alpha, the restored image would contain in some instances a significant number of artifacts, especially with extended or low S/N structure. In utilising the ability of the method to provide quantitative estimates of the reliability of these spurious features in the restoration, the procedure would almost certainly verify that they were insignificant. But it was frustrating that we should have to invest so much time and effort to check the validity of every structure in our restored image. Given that we have reached such a level of sophistication that we can compute error estimates on our restorations—unless computational or temporal matters dominate—we should not ignore them and rely solely upon the returned image for all subsequent interpretation. Nor should we fully depend on restoration methods which can only produce a single image with no error estimates. Yet it is still not unreasonable to try to refine the latest methods such that the most probable image returned is more in accordance with what we would prefer, what we would really consider most appropriate: as smooth an image as possible which is still compatible with the data. Then any structure which would appear in the image could be guaranteed to be real, or at least necessary to meet the compatibility requirement. The noise would be optimally suppressed but not at the expense of real structure.

The intrinsic correlation function

A first step towards meeting this goal was achieved by Mark Charter in his PhD thesis (1990). He developed the concept of the intrinsic correlation function (ICF). An ICF is some sort of blurring function which must be applied to a maximum entropy solution to produce what we would consider the actual restoration result. In imaging applications, the ICF is typically in the form of a convolution with a Gaussian. Within this new framework, estimates of the Classic MaxEnt solution are produced iteratively within a 'hidden' space and each of these hidden images must be operated upon by the ICF to transform it to 'visible' image space. Entities in the latter space are the only ones we ever actually see or mean to represent the true object distribution. They in turn must be operated upon by the projector operator to enable a comparison with the data. The ICF thereby enforces a minimum degree of spatial correlation in the restored image as the latter is, by construction, the convolution of some image with a Gaussian. The characteristic scale of this minimal correlation is determined by the size of the Gaussian.

By effectively band-limiting the set of all possible restored images, this approach produces results which in most instances are far more reasonable than those achieved without the use of an ICF. The improvement is the result of formally incorporating additional prior knowledge (or preference) into the restoration scheme: that there are pixel to pixel correlations in the image. Of course, if the image is of a star field where pixels in the true image should probably be completely decorrelated, the use of an ICF would not be justified and would be counter-productive. In most real scenes, however, there is at least one and usually many scales on which the pixels are correlated in the true image. By true image, I mean a pixel-gridded view of the true object distribution with absolutely no blurring. Thus the pixel-to-pixel correlations I am referring to are truly intrinsic and not due to blurring by the PSF.

If the scene of interest is dominated by a single preferred scale, the ICF approach works spectacularly. The fit to the data is highly non-biased, resulting in nearly optimum resolution. Artifacts on a scale smaller than the ICF are simply impossible by construction; artifacts on a scale the same size or larger than the ICF are generally prevented by the likelihood function: they disrupt the good fit to the data too badly. But what if there is an unknown scale, or worse yet, many scales present in the true image? Then what is the most appropriate choice for the ICF?

It transpires, just as Bayes' theorem could offer an objective, optimum choice for the regularization constant alpha, it could suggest the most probable form of the ICF. Unfortunately, while a single ICF may be globally optimum, it is virtually never optimum locally. Excessive blurring will exist in the restoration where higher resolution is clearly warranted and artifacts on the same scale as the ICF will appear where the intrinsic image scale length is much larger. Just as before, we

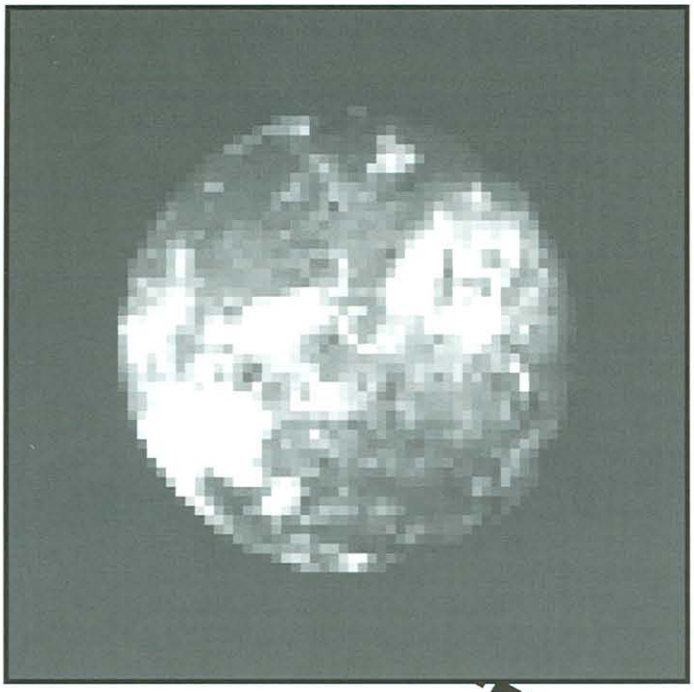

A Voyager image of Io (C1636834, Violet filter) resampled to the 22mas pixels of the FOC f/96 camera

still have not adequately decoupled the various aspects of the image.

The natural next step is to not use only one ICF but to allow for many. We accomplish this by creating multiple hidden images or 'channels', each of which has its own characteristic ICF. Each is operated upon by its respective ICF to create a visible space image. The sum of all these visible space quantities is in turn our estimate of the restored image. We thereby attempt to decouple the image by brute force. Structure in the image of any particular size is meant to be restored into the channel with the most closely matching ICF. Stars and other point sources should be restored to the sharpest channel; the night sky and other extended emission is be restored to the blurriest channel. In order to achieve this effect, however, yet another type of balancing is required: telling the restoring algorithm into which channel to throw the appropriate flux. In addition, we must determine how many channels are adequate and what the best ICF is to use for each.

Initial tests of this approach indicate that it is potentially quite powerful . Intuitive approaches to answering the questions addressed above seem viable. For example, in order to bias the restored flux into the appropriate channels, we scale the entropy defined for each channel differently. We thereby create a penalty in the objective function for allowing flux to be restored to a sharper rather than a blurrier channel which may only be overcome if the subsequent fit to the data increases the likelihood function by more than an offsetting amount. Thus we provide a means for sharp structure to appear in the image, but only if it is sufficiently demanded by the data. From all of our experiments so far, we have found empirically that the best means to scale the entropy in the different channels is by weighting the flux in the channels according to the inverse square of the characteristic linear size of its ICF. The intuitive interpretation of such a scaling is that an effectively independent pixel in any channel is given the same entropic weight. For example, as the pixels in the sharpest channel are the most independent (their ICF imposes the least correlation on them), each of these pixels should weigh the most in the entropy calculation.

As for how many channels are generally necessary, we can only offer empirically determined advice. We choose a sharpest channel whose ICF is at least as small as the sharpest believable features we expect to be able to resolve in the data. For images containing point sources, we might expect to be able to resolve objects down to the pixel level or lower. In such instances, our practice has been to restore to an image pixel grid more finely sampled than the data. We generally oversample by a factor of three in each dimension. The sharpest channel is defined with a small, but finite ICF. Typically a Gaussian of $\sigma = 1$ subpixel works well, producing familiar looking (i.e. non-single pixel) star images in the restored image. This does not compromise anything in terms of resolution: the size of the Gaussians is still smaller than that of the original data pixel. If we can reasonably assume the image

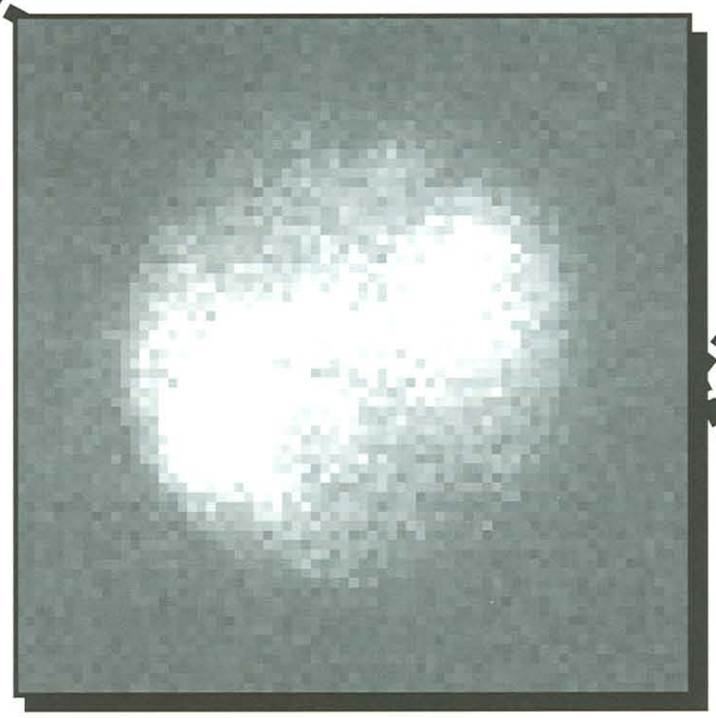

Convolved with the aberrated HST point spread function and with photon noise added.

A restoration experiment using a Voyager image of the Jovian satellite Io

One of the Violet filter images of Io (C1636834) has been resampled to the 22mas pixels of the FOC *f*/96 camera, 'observed' with the aberrated HST PSF to a reasonable exposure level and then restored with both the Lucy and MemSys 5 codes. *Nick Weir, Fionn Murtagh & Robert Fosbury*

contains only point sources, we can restore using this sharp channel alone and achieve rather remarkable results. As an example, see the restorations of the core of the globular cluster M31 (Weir, Piotto & Djorgovski 1990.). When the data image contains some degree of extended structure we have found the following approach to be the most successful: after choosing the ICF for the sharpest channel, create subsequent channels with ICFs progressively larger by factors of two in their characteristic size. The factor of two spacing was chosen only as some approximation to the continuous limit of infinitely many channels with infinitely many sizes of ICFs. In some instances, this spacing may be too coarse; at other times it may be unnecessarily fine. Again, more experimentation is necessary. In general, however, it will probably be found that the largest ICF necessary will be many times smaller than the size of the image itself. The reason is that the larger the channel's ICF, the broader its effect on and sensitivity to the restored image. It is therefore progressively more difficult for blurrier channels to contain spurious artifacts as they well be more easily detected and ruled out by the likelihood function since a greater number of pixels would be affected. With a little experimentation, a feel should quickly be obtained for just 'how much is enough', both in terms of the spacing and total number of channels required.

This multiple-channel/ICF approach represents the state-of-the-art in MaxEnt-based image restoration. More work is certainly required both on the experimental and theoretical sides. The die-hard Bayesian might quite rightfully object to our rather ad hoc approach to deciding how to choose the appropriate number and form of ICFs and especially their relative entropic weights. If alpha or the scale of a single ICF could be appropriately and objectively provided by Bayes' theorem, why should not these other quantities? Well, indeed they should! We are actively investigating the computational issues involved in doing so. Unfortunately, even when the details are worked out, we will still be very dependent upon heuristic guidelines for effective restoration, as the parameter space to be explored will be too large to search blindly. Therefore, even without the multi-channel approach being fully encompassed within a purely Bayesian framework, we perform a useful function by applying it with these various *ad hoc* approaches. We have already stumbled upon many restorations using this method which appear superior to those produced by any other technique yet developed.

In the end, the solution that some purely Bayesian multi-channel approach produces may appear less appropriate than one found using a different choice of channels or scales, or perhaps using a different method entirely. This will not be cause to disregard the Bayesian approach but merely to refine our approach, our image model or our prior probability distribution. Bayes is certainly capable of giving us what we consider the most appropriate image in the end; it is incumbent upon us to define formally and precisely what we mean by appropriate to Bayes.

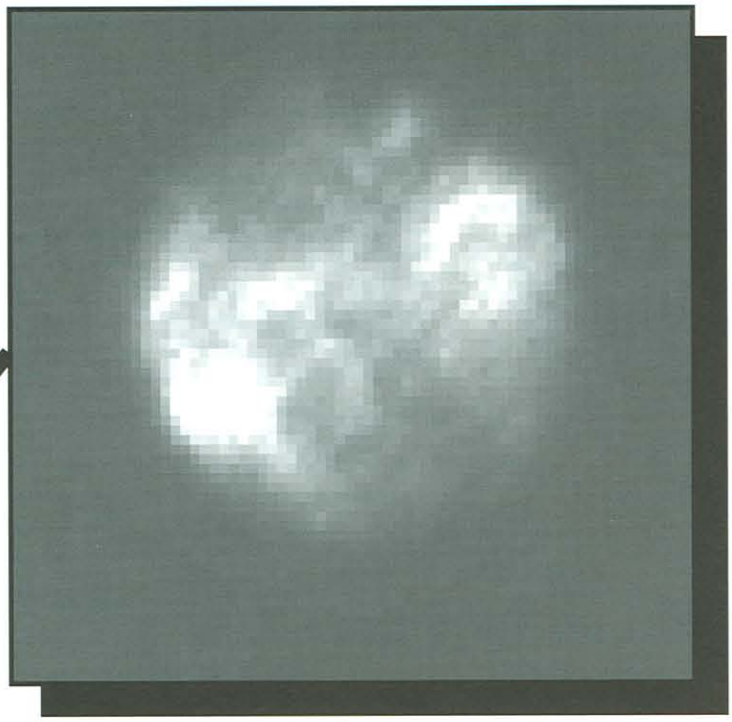

Restored with 40 Lucy-iterations. Many of Io's surface features, visible on the original but not particularly apparent on the 'observation', are revealed.

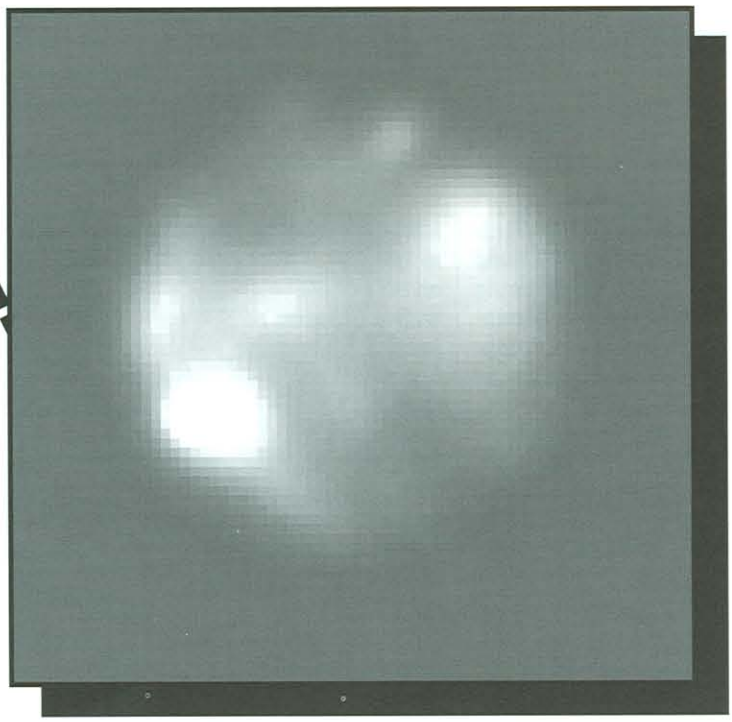

Restored with 58 MemSys 5 iterations using three ICF channels. Again, the more important features are revealed. There is a high level of consistency between the two methods although MemSys 5 is seen to be more conservative at high spatial frequencies.

References

Gull, S.F. & Skilling, J., 1989. *Quantified Maximum Entropy, "MEMSYS 3" Users' Manual*

Weir, N., Piotto, G. & Djorgovski, S., 1990. *ESO Messenger*, No. 62, December 1990, p48

HST Seminars in Europe

Following a suggestion by the ST-ECF Users' Committee, seminars on HST were given in several European Institutes prior to the deadline for Cycle 2 proposal submission.

Piero Benvenuti

During the past months ST-ECF staff gave a number of seminars in several European Institutes. The main purpose of this initiative was to keep the European astronomical community up to date with the recent developments of the HST operations in view of the approaching deadline for the submission of observing proposals for Cycle 2.

The main topics covered by the seminars were the current performance of the scientific instruments, the impact of the spherical aberration on HST science, the possibility of partial recovery by image and spectral restoration, the plans for a long term solution with the installation of WF/PC II and COSTAR and finally, but most important, some early scientific results.

The seminars were well received and attended and we hope they may contribute to a good European response for Cycle 2.

Date	Place	Given by
3-May-91	Lunteren, The Netherlands	G. van Moorsel
3-May-91	Turin, Italy	P. Benvenuti and F. Pasian
17-May-91	Brussels, Belgium	G. van Moorsel and B. Pirenne
11-Jun-91	Marseille, France	A. Caulet and J. Walsh
11-Jun-91	Bern, Switzerland	P. Benvenuti
11-Jun-91	Cambridge, United Kingdom	H.-M. Adorf and R.A.E. Fosbury
12-Jun-91	Heidelberg, Germany	M. Rosa
13-Jun-91	Hamburg, Germany	M. Rosa
18-Jun-91	Madrid, Spain	P. Benvenuti
18-Jun-91	Durham, United Kingdom	H.-M. Adorf and R.A.E. Fosbury
19-Jun-91	Stockholm, Sweden	G. van Moorsel
20-Jun-91	Tenerife, Spain	P. Benvenuti
1-Jul-91	Tübingen, Germany	M. Rosa
2-Jul-91	Garching, Germany	J. Walsh
9-Jul-91	Bonn, Germany	R.A.E. Fosbury
29-Jul-91	Porto, Portugal	L.B. Lucy

A hot day for HST

Despite the "heated" beginning, the HST Joint Discussion at the IAU General Assembly was quite successful in demonstrating the scientific value of the project.

Michael Rosa

By Tuesday night, July 30, the 21st IAU General Assembly (held July 22 to August 1, 1991, in Buenos Aires, Argentina) already had concluded most of its extensive program. Wednesday, the last day of scientific meetings, was to see the 7th Joint Discussion on "First Results from the Hubble Space Telescope". Leaving the hotel in the morning we took a taxi to the venue, and were puzzled by the driver's comments from which we could only grasp "... no es posible ...". Arriving in the crowded street leading to the Cultural Center San Martin we noticed an unusual traffic jam. As we got closer, the crowd resolved into police cars, fire engines and hundreds of astronomers with fire fighters mixed in. Thick clouds of smoke and soot billowed from the building being used for scientific meetings and administrative support of the Assembly. Luckily, the fire had broken out very early in the morning, so only a few people were in danger, and nobody was severely hurt. Was this another link in a chain of troubles associated with the HST project, as some commented? Or was it the hot topics to be presented in a whole day meeting scheduled to start at 9:00 a.m.?

The local organizing committee and the emergency meeting of the IAU Executive did an excellent job of recovery. Lacking a megaphone the word spread by mouth and at 10:30 the Joint Discussion on HST could be opened in an auditorium with a capacity of 500 seats completely packed. Meanwhile, the slides of the next few speakers were salvaged from the Cultural Center by brave fire-men with oxygen masks. Only one and a half hours were lost from the schedule for the day. This was to summarize the significant scientific results from HST, its current performance and capabilities, work on software to analyze data and to improve the current imaging and spectral capabilities through image reconstruction, and finally the future plans for recovery in-

cluding the second generation instruments. Nevertheless it was decided to stick as much as possible to the original schedule in order to present at least the scientific results and the future plans for recovery in orbit to the worldwide astronomical community represented.

Results obtained with the Faint Object Camera and the Wide Field Planetary Camera that had been publicized already either in the astronomical literature or through NASA and ESA press releases were briefly reviewed. Individual new results presented included: the most accurate distance ever determined to an individual object in the LMC obtained by combining the lightcurve of the forbidden line emission of the ring around the SN 1987A with its accurately measured dimensions; jets from AGNs which appear surprisingly identical at radio and optical wavelengths at resolutions down to 0.04 arcseconds; an elongated optical image detected at the position of the binary pulsar PSR 1913+16; photometry of the globular cluster 47 Tucanæ revealing more than 20 blue stragglers in its core region, lending support to the hypothesis that these are binary stars.

The non-imaging instruments, less advanced in the OV/SV/GTO schedule, were given an excellent opportunity to present their first results to the community. The orders of magnitude jump between IUE and GHRS in both sensitivity and resolution was made evident. The Lyman alpha forest in the echelle mode spectrum of 3C 273 resolves into about 15 redshift components, grouped into families of typically 4 each, suggesting that the medium traversed by the line of sight is clumpy. A few Lyman absorptions seem to stem from the Virgo cluster. Circumstellar lines in Beta Pictoris observed in the echelle mode reveal clumpy disk material raining down onto the star and a wind of thin matter blowing away at the same time into the remaining solid angle. The CIV 1548 Å line profile at a few km/s resolution of an LMC early O type star made its impact. NLTE analysis of such stars in local group galaxies will add considerably to the understanding of nucleosynthesis and the physics of stellar winds at differing metallicities. The Deuterium line, clearly resolved from heavily saturated Lyman alpha, was shown so the audience (including project-related colleagues), eager to learn the cosmologically interesting H/D ratio, was given hope that this will be determined in the very near future.

The HSP, the last instrument to become fully operational, could not yet excel with scientific results, but has demonstrated nominal performance as far as the spherical aberration permits. Stability tests show that 1/100 of a magnitude will be achieved routinely, including jitter and aberration, on the largest science aperture.

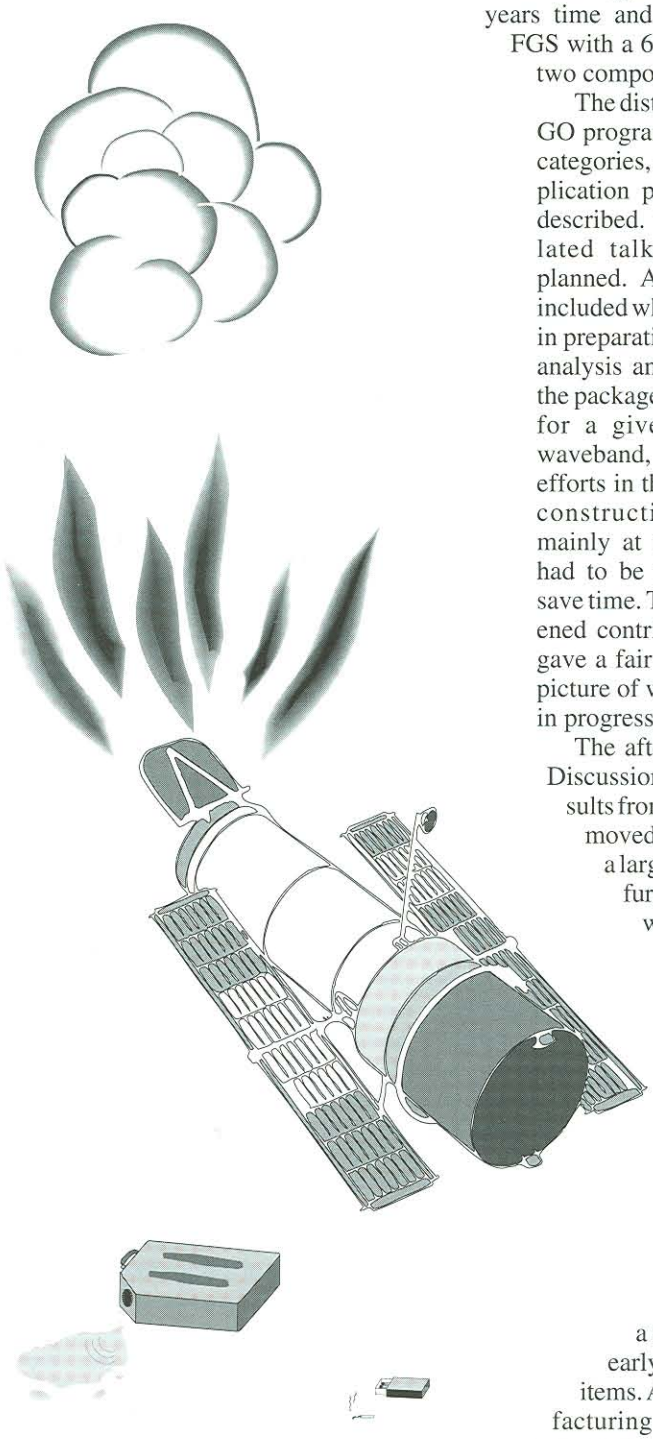

The FGS instruments, in the first place vital to the tracking and guiding of the telescope, have been used to accurately determine the orbit of a highly eccentric close binary. Five measurements taken with the FGS at the current phase close to periastron combined with the ground based data available only around apastron give an orbit with a sigma of 4 mas for the FGS data. Closest approach is expected in a years time and will challenge the FGS with a 6 mas distance of the two components.

The distribution of GTO and GO programs among scientific categories, and details of the application process were briefly described. Three software related talks were originally planned. Areas to be covered included what was available and in preparation at STScI for data analysis and observer support, the package to compute the PSF for a given instrument and waveband, and the continuing efforts in the field of image reconstruction concentrated mainly at the ST-ECF. These had to be merged in order to save time. The drastically shortened contribution nevertheless gave a fair and comprehensive picture of what is available and in progress.

The afternoon part of Joint Discussion VII on "First Results from HST", having been moved once again, now into a large cinema two blocks further away, concluded with a review of the possible, viable and likely options for recovery from the spherical aberration in orbit. It was noted that the current failures of gyros combined with the fear of fatigue in the solar panels might urge the project to consider a servicing mission as early as possible for those items. Also, while the manufacturing of COSTAR still seems to be possible within the deadline for a late 1993, early 1994 mission, WFPC II might not be ready then, if the installation of actuators is felt to be essential.

The speakers and the audience were pleased to cool off later at the closing banquet after this hot day for HST.

The Archive column

The future of the European HST Archive

New developments in the US—as a result of the development of the fully-fledged HST Archive, DADS—will have an effect on the access to archived HST data in Europe

Rudolf Albrecht & Fabio Pasian

The ST–ECF has established an archive facility designed to accommodate the full set of HST data. The archive is known as the HST Science Data Archive (SDA) or, more commonly, as the DMF (Data Management Facility) and is the counterpart of the SDA at the Space Telescope Science Institute. The ST–ECF and the STScI have had a strong and successful collaboration in this area. Why is it a 'facility' rather than an outright archive at the STScI? It is because the full archive is only now being built and it will not be operational before 1993. The archive, called Data Archive and Distribution System (DADS) is being produced by LORAL, a US aerospace company, under contract to NASA.

In order to make the archive acceptable to the astronomer users, the task of developing a user interface for DADS has been given by NASA to the STScI. The situation as it develops is the following: within two years there will be an archive at the STScI which is different from the one at the ST–ECF. There will be a user interface to this archive which will not necessarily be identical to the user interface available at the ST–ECF.

Many of you have been using the ST–ECF-designed user interface to the SDA, a software package called STARCAT (Space Telescope ARchive and Catalogue), in order to browse through the HST catalogue or to access one of the many astronomical catalogues which are available. As more HST data become public and as the network connections to the ST–ECF improve to support the required high data rate, use of STARCAT will be increased. It is at the moment also the user interface to the DMF at the STScI. It works well and is flexible, as demonstrated by the fact that through it users can access many different databases, including some which reside at other institutions. It is transportable—running on VMS and on Unix machines—and it is maintainable: the original author (Alan Richmond, then working at the ST–ECF) has left the HST project, and the software still survives.

There are good reasons for upgrades and new developments. STARCAT was designed for VT-100 type terminals (VTUs), using the TERMWINDOWS software (written by F. Ochsenbein) to makes these devices behave almost like modern window-oriented workstations. Obviously, these efforts can only go so far; actual window oriented workstations offer many more capabilities than can be emulated on a VT-100. Equally obviously, with the technological development going towards distributed processing and high-bandwidth network connections, we want to exploit these capabilities. On the other hand not everybody yet has a workstation and not everybody will have high-bandwidth network connections for a while.

Taking these considerations into account the current plan for the development of the user interface to DADS foresees two different front-ends: one based on STARCAT, with possible changes and/or improvements in order to cover the different functions of DADS for that part of the user community which still has VTU terminals; and another one exploiting the full capabilities of Window-Icon-Mouse-Pointer (WIMP) techniques for those who have hardware to support it. Concepts to achieve this are currently being explored.

While the user interface issue can be resolved between the STScI and the ST–ECF, there are other DADS-related issues which are more difficult to decide. They involve the NASA contractor LORAL who does the actual development work for DADS. For the ST–ECF as part of the ESA Space Telescope program it is extremely cumbersome to insert anything exceeding basic requirements into the development process of a NASA contractor. In addition, the rapidly developing technology in the opto-electronics area will, by itself, necessitate considerable changes.

Trying to extrapolate two or three years into the future we can imagine the following situation: DADS will be taking over from DMF at the STScI, re-ingesting in a massive effort all DMF data into the DADS archive. This will be used as an occasion to streamline the archive, to optimise the data format and to apply corrections and additions to the header entries which will have become necessary. As is already the case, the ST–ECF is receiving copies of all the optical disks which are being produced by DADS. These disks will have a far higher capacity than the ones we use now.

At the ST–ECF we will continue to operate with SDA-DMF, porting the archiving system to the Unix environment. We will continue to support STARCAT as the user interface to our archive and, if at all possible, as the user interface to DADS. We will make efforts to provide for our users a version of the new WIMP-based DADS user interface in order to provide the possibility to access the ST–ECF archive in the same way that users at the STScI access DADS. This will, however, depend on a large number of variables and so it is not yet possible to commit ourselves to this.

In the meantime, we intend to develop STARCAT in the direction of exploiting WIMP capabilities. This enhancement will allow easier integration of graphics facilities into the user interface (e.g., plotting the results of a query within STARCAT and without the production of intermediate files; or the refinement of a query via graphic interaction). Such facilities will also be provided on VTUs with graphics capabilities, but they obviously produce better results in a WIMP environment.

Other developments which are easy to foresee include the graphic browsing of the archive, where data in compressed form (e.g. spectra) are kept on-line directly in the database from where they can quickly be retrieved and displayed.

Another STARCAT improvement will be the availability of the user interface at remote sites. This would provide the remote users with the full screen-managing capabilities of the system, while at the same time minimising network traffic. A project allowing STARCAT to run remotely over DECnet is already being developed in collaboration with ESIS (European Space Information System). Within one year we expect to be ready to distribute STARCAT to remote sites connected to ESO/ST–ECF both on SPAN and through Internet.

Archive user registration

Obtaining authorisation to dearchive HST data

If you wish to have an account set up to dearchive HST data, please complete the form below and return it to us, on paper or via electronic mail. You will be assigned a unique username and password, which will be sent to you at the e-mail address specified on the form. You will be registered for one year. One month before expiration, you will be warned, so that you can renew the registration if you wish to do so. Any changes to the information given in the form should be communicated promptly so that we can take appropriate actions.

ST–ECF Archive Registration Form

❑ new registration ❑ renewal

First name: _
Middle initial(s): _
Last name: _
Title [Dr., Prof., Mr., Mrs.]: _
Institute: _
Address: _

Phone: _
Fax: _
Preferred E-mail address: _

User category:
❑ HST PI
❑ HST Co-I
❑ scientist
❑ post-graduate
❑ student
❑ other (specify): _

Preferred distribution medium:
❑ magnetic tape 1600 bpi
❑ magnetic tape 6250 bpi
❑ Exabyte cartridge
❑ DAT cartridge (not supported yet)
❑ Local Area Network (ESO users only)

FITS blocking factor (1–10): _

Archive news

Benoît Pirenne & Fabio Pasian

This time, we would like to give some information regarding the current status of the HST catalogue and the archive data. An update of the STARCAT connection instructions is also given.

Documentation

A new version of the STARCAT Quick Guide (for STARCAT version 3.4) has been written. An HSTCat Quick Guide explaining the particularities of the HST catalogue is also ready. Distribution to all known STARCAT users is underway. If you would like to receive a copy of the documents, please contact us.

HST Data

All data—with a couple of exceptions—from SV (Science Verification) and OV (Orbital Verification) taken before June 1st became public on July 1st. All new OV and SV data taken after June 1st will be granted a proprietary period of 1 month. This means that from September on, a set of about 15000 datasets will be retrievable through the ST–ECF archive. To become a registered user of the archive, please return the form reprinted on this page to one of the authors.

The connection through STESIS has been improved recently and is now less cumbersome. Data can be produced by STARCAT and retrieved to the user's home computer by using a file copy utility (DECnet or ftp) as explained in the logout message of STARCAT.

STARCAT WAN access

A project is being developed in collaboration with ESIS which will allow STARCAT to be run remotely, at the user's own site, connecting transparently to the database at ESO/ST–ECF via DECnet. This facility will be available for integration and internal testing by the end of September. The extension of the software to include Internet connections is envisaged.

Connection

To connect to STARCAT at ESO/ST–ECF, please use one of these options:

In-house: for staff and visitors (Garching or La Silla)

From any of the Unix workstations: type <starcat> when logged in.

External Users:

connect to the STESIS computer (SPAN: STESIS or 28.99; Internet: stesis.hq.eso.org or 134.171.8.100)

Use the STARCAT account without password.

Scientific use of the HST Archive — the coma of Comet Levy 1990c

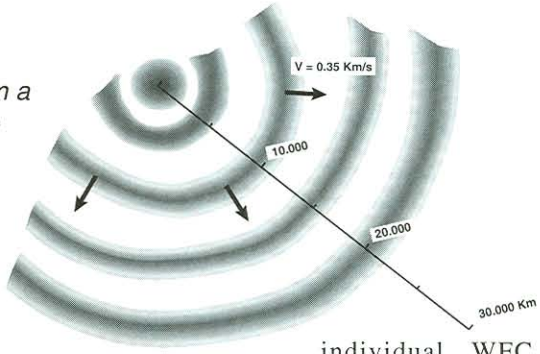

The expanding dust shells from a comet—an example of the use of public data from the HST Archive.

Marco Fulle†, Fabio Pasian & Piero Benvenuti

The observations

Eight frames of the inner coma of C/Levy 1990c were obtained by H. Weaver with

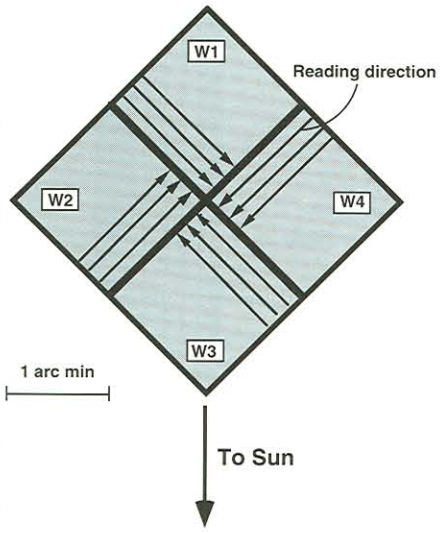

Figure 1: Field of view of the WFC on the sky showing the readout directions for the four 800×800 CCDs.

the Wide Field Camera (WFC) of the Hubble Space Telescope on 27 September 1990 as a part of the Director's Discretionary time programme.

A pair of them were taken with the F439W (blue) filter with an exposure of 2 seconds, three pairs were taken with the F785M (red) filter with exposures of 1, 2 and 4 seconds, respectively. There is a time interval of $6^h\,27^m$ between the two frames composing a pair having the same exposure time. The pixel resolution is 78km at the Earth-Comet distance of 1.055AU. The images have been released into the public domain, and we have retrieved them from the ST–ECF HST Science Data Archive.

The data

WFC images are composed of four 800×800 frames, each corresponding to a CCD chip (see figure 1, showing the field of view of the WFC on the sky). Each image is retrieved from the archive as an 800×800×4 FITS image file and an associated FITS table containing parameters for each frame [1].

Two pairs of images (the red exposures of 2 sec and 4 sec) were selected for analysis, since they should contain the largest amount of information on the dust environment of C/Levy. The FITS files were read into the MIDAS environment, and a MIDAS procedure used to extract the individual WFC frames and to build a mosaic according to the specifications found in [2]. In figure 2, the mosaic of one of the 4 sec exposures in red light is shown. The effects of the junctions, of cosmic rays and of a low quality background calibration are clearly visible. A subset of the whole image has been extracted, containing most of the W2 and W3 chips, and is shown in figure 3.

The dust shells

To detect expanding shells in the coma, as suggested by Weaver [3], difference images were obtained by performing a subtraction between exposures of same length taken ≈ 6.5 hours apart. The image shown in figure 4 is one of such difference images, obtained by subtracting two frames having an exposure time of 4 sec, after a median filtering step aimed at removing cosmic rays. The darker stripe in the cen-

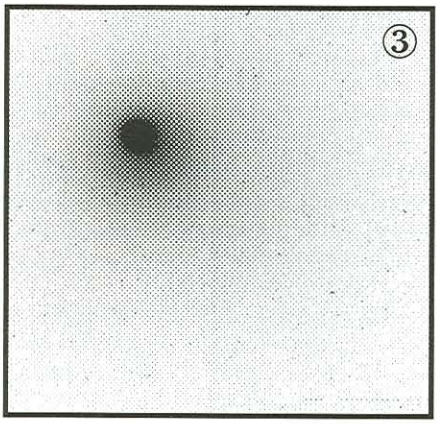

-

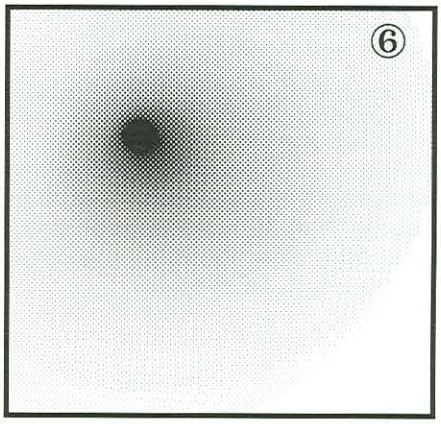

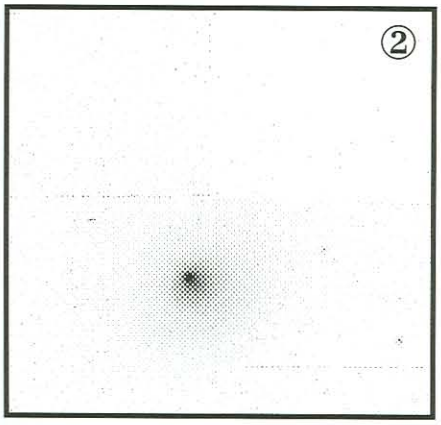

Figure 2: Image W0DC0901T, obtained as a mosaic of four frames (4 sec exposure — size 1600×1600).

Figure 3: Subset of image W0DC0901T (size 512×512).

Figure 6: Model for image shown in figure 3, obtained from fitted ellipses.

Figure 7: Model residuals (figure 3 – figure 6).

=

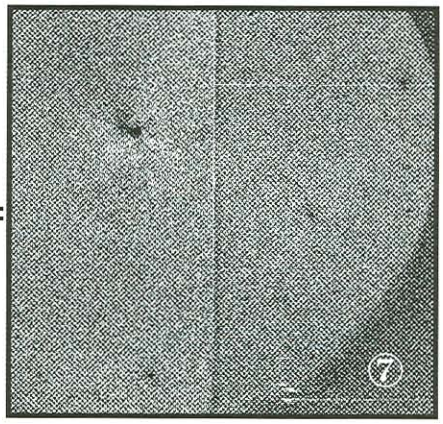

† Osservatorio Astronomico, Trieste, Italy

ter is due to the propagation of the junction defect. Shells of dust propagating through the coma are clearly visible, since they appear as a number of concentric stripes of alternatively higher and lower intensity, centred on the comet nucleus.

An integral radial profile for the difference images has been built by adding the radial profiles centered on the comet nucleus and taken over an interval of 90° with an angular step of 1°. The result for the 4 sec exposure difference image is shown in the graph in figure 5. The ordinate scale has been enhanced by multiplying the integral radial intensity by the distance from the nucleus of the comet, in order to compensate for the hyperbolic decay of the intensity. The still decreasing envelope of the enhanced integrated intensity is probably due to the fountain-motion of the dust in the shells.

The black shells in the difference images were ejected $9^h.0$ and $21^h.5$ before the first set of observations, whereas the white ones were ejected $4^h.0$, $15^h.5$ and $28^h.0$ before the second set of observations (obtained $\approx 6^h.5$ later), so strongly suggesting a rotation period of the comet nucleus close to 24^h or multiples.

Such a period is inconsistent with the periodicity observed with coma photometry (19^h [4, 5]), a case very similar to that which occurred for P/Halley observations. We point out also that in the limiting case of a constant dust velocity—corresponding to the shortest possible period—the subtraction image yields a period of 22^h.

The observations of at least two shells within the rotation period implies a nucleus characterised by at least two active regions placed almost exactly in opposite nuclear hemispheres. Subtracting the two red frames having exposure times of 2^s yields the same results.

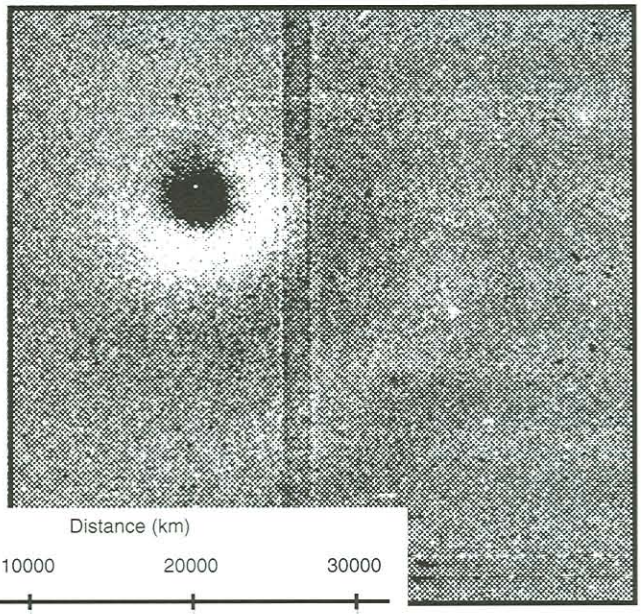

Figure 4: Subset of median filtered difference image W0DC0401T – W0DC0901T (4^s exposures taken $\approx 6^h.5$ apart).

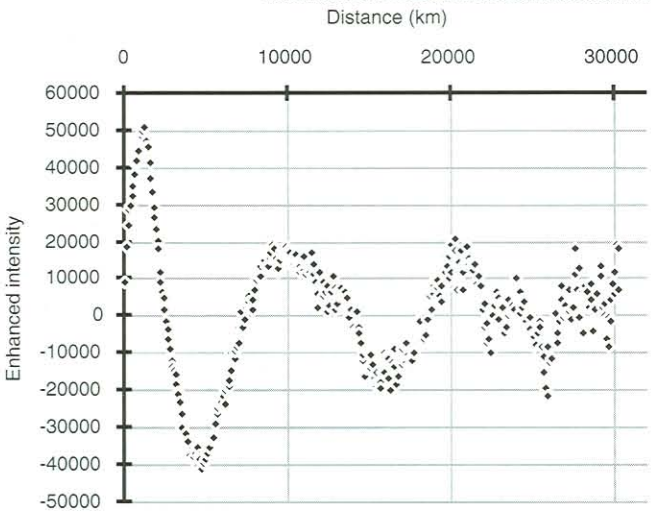

Figure 5: Integral radial profile for difference image shown in figure 4.

Trajectory of the dust shells and physical parameters

The isophotes of the image have been fitted with ellipses using the PLEINPOT package by Ph. Prugniel—normally used for elliptical galaxies. The image in figure 6 shows the model of the comet computed from one of the 4 sec exposures and obtained from the fitted ellipses, while figure 7 shows its residuals. The ellipses are equidistant in intensity. The ellipticity of the fitted curves is fairly low (axes ratio ranging from ≈ 0.85 to 1.0) while the residuals show the good agreement of original and fitted data (maximum error $\approx 5\%$ in the nucleus).

To infer the trajectory of the expanding dust shells, we have fitted (figure 8) the vertices (dashes) of the shells by means of a parabola (grey line) which is characterized by the parameters $a = -6.15$ and $b = 3.65\times10^{-4}$ km^{-1}. The differences between the observed and computed positions of the shells is smaller than 400km for the last two, and smaller than 100km for the others. We have considered as the vertex of the expanding dust shells the intersection of the fitted ellipse with the straight line between the centre of the ellipse itself and the Sun. Due to the low ellipticity, this point is the closest to the Sun.

To solve the non-linear system describing the dust motion in the coma of Comet Levy we need also to constrain the times of dust ejections which are given by the integrated intensity of the difference images. We identify each minimum with the immediately preceding maximum.

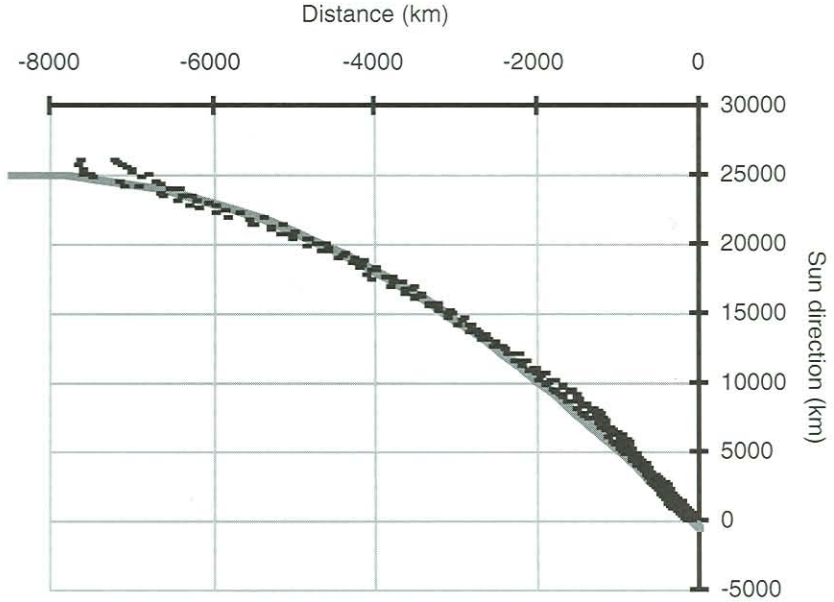

Figure 8: Positions of vertices of expanding shells (-) for four images (W0DC0401T, W0DC0501T, W0DC0901T and W0DC0A01T — 2^s and 4^s exposures) fitted by means of parabola (grey line).

HST observing efficiency

The 'on target' efficiency of the HST is now approaching its nominal value. It is possible to calculate—from the observation catalogue—the actual 'exposure time' efficiency which is more representative of the productivity of the observatory.

Piero Benvenuti & Benoît Pirenne

During recent months, the efficiency of the HST operations and rate of success of target acquisitions have gradually increased thanks to a number of improvements in the ground control software and in the observing procedures. Particularly effective, among the latter, are a better understanding of the FGS as well as a refined alignment of these with the instrument apertures in the focal plane.

Now that the Science Verification phase is being completed and the GTO and GO programmes have begun to be executed routinely, we can attempt to evaluate the overall efficiency of the HST observatory. The efficiency of the ground system, including scheduling, and of spacecraft operations, measured by the time spent by HST 'on target'—that is pointing toward a pre-planned position—is currently about 30% for external targets and 37% if Earth flat-fields, snapshot exposures and internal calibrations are included. The observing time, measured by the actual exposure time during the months of November 1990 to June 1991 is shown in Figure 1 for the different instruments. The actual 'exposure time' efficiency, summed over all the instruments is plotted in Figure 2. All these numbers have to be taken with some caution because a fraction of them still refer to Science Verification observations and they may not reflect a typical sequence of scientific exposures.

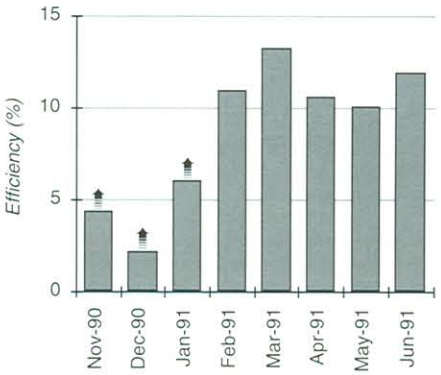

Figure 2: The HST efficiency, summed over all instruments, measured by actual exposure time rather than the more usually quoted 'on target' time—which includes instrument overheads and target acquisition. The first three columns are lower limits due to errors in the observation catalogue which are now being corrected.

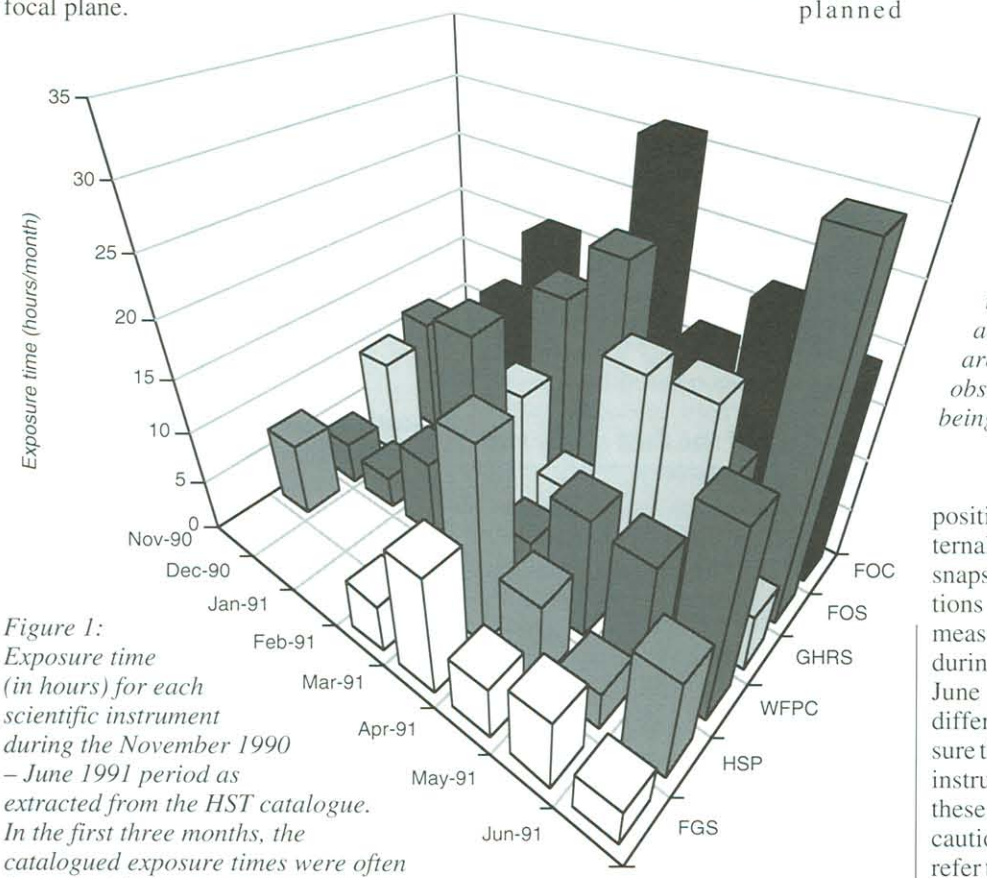

Figure 1: Exposure time (in hours) for each scientific instrument during the November 1990 – June 1991 period as extracted from the HST catalogue. In the first three months, the catalogued exposure times were often incorrect (due to incorrect information in the file headers) which results in some empty locations in the array.

This assumption is justified by the fact that other identifications would imply too high dust velocities ($v > 1$ km sec^{-1}).

The mathematical details can be found in [6]: only the final results are given here. The orientation of the vertices of the expanding dust shells imply a counterclockwise spin. The lag angle ϕ between 10° and 15° implies a maximum dust release within an hour pm.

Dust grains have a size of 3 μm (assuming a dust bulk density of 1 g cm^{-3}) and their initial velocity is 0.35 km sec^{-1}, in good agreement with hydrodynamical model computations. The perspective conditions of the observations have negligible influences on all these results.

References

[1] F. Pasian, B. Pirenne, S. Hill, *ST–ECF Newsletter No. 15*, March 1991

[2] WFPC Instrument Handbook, STScI, May 1990

[3] H. Weaver, *STScI Newsletter No. 8*, March 1991

[4] B.G. Marsden, *IAUC* 5085, 1990

[5] D.G. Schleicher , R.L. Millis, D.J. Osip, *IAUC* 5081, 1990

[6] M. Fulle, F. Pasian, P. Benvenuti, *Annales Geophysicæ*, 1991 (in press)

ST–ECF computer information services

A reminder of how to contact us by computer network. Please note that there will soon be major changes in the ESO/ST–ECF hardware configurations which are bound to affect communications. We are also pleased to announce that we are now connected to the Internet.

Richard Hook

A problem facing all European HST users and prospective users is getting up-to-date and accurate information about the HST. To help overcome this the ST–ECF runs two electronic information services for the astronomical community, a bulletin board service (STINFO) and an electronic mail 'hotline' (STDESK). The purpose of this note is to re-advertise these facilities, describe them in detail and to ask for suggestions for improving them. Some changes to the ways in which they are accessed — due to major changes in the ESO computing hardware and networks — are also mentioned.

Bulletin board access

To access STINFO you need to login to one of the ESO VAX computers, either over the SPAN network or using PSI. In addition we are pleased to be able to announce that we are finally connected to the Internet after many delays. If your computer is on SPAN the command (from a VMS computer) would normally be SET HOST ESO. If you find that the node ESO is unknown you can try SET HOST 28760 — but please also ask your system manager to add the number to your system. An identical service is available on the ESIS μVAX machine STESIS (=28771). If you only have PSI access you can call us on 0262458900924 but please also contact us by telephone to get password information. You may login from the Internet using the command 'telnet stesis.hq.eso.org'. Once you are connected you give the username STINFO, there is no password, you are then prompted for your name and then automatically enter the bulletin board software. You may find out what is available using the DIR command and read bulletins using the READ command.

The great virtue of STINFO is that the information is very up to date — typically reflecting HST operations from the previous day! We post the daily reports of HST operations which are compiled at the Goddard Space Flight Center in the US. These give information about which proposals are currently being executed by HST, how successful they were and what is coming up. Operational information (e.g. guide star acquisition success rates) is also given as well as lists of outstanding problems. In addition to the daily reports more detailed weekly reports are included as well as other items of topical interest — e.g. announcements of new software, proposal deadlines etc. In the early days, sources of reliable news tended to change or dry up from day-to-day but the system is now more stable. If you have trouble accessing this service or have any suggestions for additional contents please contact us. Note also that much other information and software is available from the STEIS service at the Space Telescope Science Institute (but most conveniently for those attached directly to the Internet). See *ST–ECF Newsletter*, May 1990, **13**, p3 for more details of this service.

STDESK may be contacted at any of the following network addresses:

SPAN:	ESO::STDESK (28760::STDESK) *or* STESIS::STDESK (28771::STDESK)
Bitnet/EARN:	STDESK@DGAESO51
Internet:	stdesk@eso.org
PSI:‡	PSI%0262458900924::STDESK
UUCP:	eso!stdesk *or* ...uunet!mcsun!unido!eso!stdesk *or* stdesk@eso.uucp

‡ to be discontinued

Information requests

In addition to STINFO, the ST–ECF has an electronic mail account called STDESK to which any HST related enquiry may be addressed. This account is checked for new mail regularly throughout every working day and you can expect a reply to a query within hours of it getting to us. Please do not hesitate to ask anything — a short message to us may save you a lot of time. If the ST–ECF member manning STDESK when your message is received cannot answer your question they will pass it on to a more appropriate person or advise on who else you should consider contacting. This facility may also be used by telephone or FAX using the numbers given on the back page.

Computer system changes

An important change to the ST–ECF computing resources will occur later this year with the removal of the main ESO VAX/VMS machines (ESOMC1 and ESOMC0). These old VAX 8600 machines will be replaced by UNIX server machines. In addition small VAX server computers will be bought to allow continued VMS support within ESO. The STESIS machine which is connected to both SPAN and the Internet will remain. When the main VAXs go, their SPAN DECnet node names will probably go with them but the name ESO and STESIS will remain and you are strongly advised to use these now. Internet connections will stay the same.

Much of the above information has been included in earlier Newsletters but we hope that by re-stressing it now that HST is in routine operation and the news sources are reliable we will encourage more people to try these services. If you have any ideas for improving them, please let us know.

News

Piero Benvenuti & Richard Hook

4th ESO / ST-ECF Data Analysis Workshop

The 4th Workshop of this series will take place at ESO, Garching, on May 13-15, 1992. This time the main topic will be algorithms and methods for the analysis of spectroscopic data. The preliminary programme and deadline for receiving abstracts for contributions as well as registration details will be announced in the next issue of this Newsletter. Please note that this Workshop will not be followed by the usual ST-ECF Users' Meeting which will more appropriately take place during the HST Scientific Workshop (June 29th - July 7th, 1992, see page 2 of this issue).

Important reminder

European astronomers who have submitted HST proposals for the 2nd Cycle are reminded that they should send one copy of their proposal(s) to the ESA HST Project Scientist, c/o ST-ECF. These copies are used for accounting purposes only, i.e. to evaluate the ESA share of submitted and allocated HST time, and will be treated confidentially.

STINFO trivia

During the period since the launch of HST there have been a total of 6084 logins to the STINFO account. The busiest day was June 29th 1990 (just after the spherical aberration became public) when there were 154 logins. We have had some famous users (if one believes the names people enter when prompted). These include Mickey Mouse, the galactic travellers James T. Kirk and Dr. Who (twice) and even Diego Maradona (during the World Cup). Some other unusual entries cannot be reproduced here on grounds of taste.

Contents

We should like this Newsletter to reach as wide an audience of European astronomers as possible. If you are not on the mailing list but would like to receive future issues, please write to the editor stating your affiliation.

Contacts at ST–ECF

Head: P. Benvenuti
Tel., +49–89–320 06 290
Computer: PBENVENU

Science Data and Software: R. Albrecht
Tel., +49–89–320 06 287
Computer: RALBRECH

Science Instrument Information: R.A.E. Fosbury
Tel., +49–89–320 06 235
Computer: RFOSBURY

Hot-line and information Computer: STDESK (for e-mail)
STINFO (for bulletin board)

Editor: Robert Fosbury

Typeset with Apple® Macintosh™
Printed by Royal Media Werbgesellschaft mbH, Ottobrunn

Published by the
Space Telescope–European Coordinating Facility
European Southern Observatory
Karl-Schwarzschild-Str. 2
D–8046 Garching bei München
Germany
Telephone: +49–89–320 06 291
Telex: 528 282 22 EO D
Telefax: +49–89–320 06 480
Computer: *Bitnet/EARN:* DGAESO51
SPAN: ESO (28.88)
Internet: eso.org

esa SPACE TELESCOPE
ESO European Coordinating Facility