European Southern Observatory

MIDAS Courier

Newsletter of the MIDAS Users' Community

Volume 5, Number 1 July 1995

Editorial

During the last half year the MIDAS group experienced an increasing impact of the VLT construction and its planned instrumentation on its activities. The result is that we have entered the - anticipated - phase where most of the manpower in the MIDAS Group is spent on VLT related issues. Examples are the in-house development of the data reduction software for UVES, ISAAC, as well as further development of the on-line functionalities of MIDAS, in particular in the area of interprocess communication. Because these on-line capabilities will be part of the core of the forthcoming 95NOV release sufficient time has been reserved for thorough testing. Therefore, the release cycle for the 95NOV release has already started internally with a freeze of the α-release last April. In the first half of August the β-test version will be frozen and will be sent to our external test sites, allowing for two extra months of test time. More about the 95NOV release can be found under General News on page 2.

Although the VLT instrumentation now requires our attention also, improvements and developments of data reduction software for La Silla have to continue. To set priorities here, in the beginning of the year the software tools ESO-MIDAS offers for off-line calibration and reduction of data coming from the La Silla telescopes have been reviewed. Without going into details, the overall situation in the optical part of the spectrum, both imaging and spectroscopy, seems satisfactory. However, a clear lack of calibration and reduction software was noted in the area of infrared imaging. To improve in this area we are now working on new on- and off-line software for the IRAC2 camera. We anticipate to have a basic version available in the 95NOV release. Whether or not we will be able to support other infrared instruments on La Silla (*e.g.* TIMMI or Come-On Plus) will depend on available manpower.

As a result of limited manpower, and after consultation of the MIDAS sites involved, it has been decided to reduce the support of VAX/VMS sites. However, since MIDAS is built using standards like X11, POSIX we believe that this policy will only have a limited impact on current VMS users of MIDAS. More detailed information about this subject can be found on the next page.

To improve in the area of user support, a number of additional services have been implemented in the first half of 1995. Among those are the `midas-users` and `midas-announce` mailing lists, meant for general discussions and announcements. More about these lists can be found below. Another service is the availability of a MIDAS FAQ (frequently asked questions) file on our anonymous ftp account. The FAQ contains answers to all sorts of questions from installation related subjects to reduction and calibration packages. The (compressed) FAQ file `midas-faq.doc.Z` can be obtained via anonymous ftp `ftphost.hq.es.org` (134.171.40.2). It is located in the directory `/pub/midaspub/94NOV/doc`.

In this ESO-MIDAS Courier:

ANNOUNCEMENT

Testing of new ESO-MIDAS releases

With the increased pressure from the VLT project with respect to the on-line version of MIDAS, we find it difficult to maintain the high level of verification of releases on VAX/VMS platforms, similar to the one we maintain for UNIX systems. In addition, our access to and experience with VMS as well as OpenVMS systems have been significantly reduced especially since the ROSAT/MPI project has moved to Unix OSF/1 on their DEC/Alpha systems.

In order to save internal manpower for maintenance and release, thereby strengthening the development of a high quality on-line version of MIDAS for the commissioning phase of the VLT, we will base future development, maintenance and distribution of MIDAS on the following software requirements:

- Operating system conforming to POSIX.1 and POSIX.2;
- X11 and OSF/Motif display systems;
- ANSI FORTRAN-77 and ANSI-C compilers;
- PostScript for hardcopy;
- RPC and TCP/IP sockets for interprocess communication.

New developments and releases will be verified only on a limited set of UNIX platforms available to us locally and conforming to the specifications above.

This policy will only have a limited impact on current VMS users of MIDAS as VMS/OpenVMS is unlikely to change significantly, the current version of MIDAS runs on these platforms, and OpenVMS offers POSIX compatibility. The main concern would be that new applications are not tested automatically.

The reduced level of testing for VMS/OpenVMS has been discussed with all current MIDAS VMS sites and in the ESO Users' Committee before the decision was made. We will make a beta-release of MIDAS available to a user site willing to perform the validation of the VMS installation of MIDAS.

General, System and Application News

The MIDAS Group, *European Southern Observatory*

General News

New Release 95NOV

The new release of MIDAS, release 95NOV, is now in its alpha state. This release will be of importance for the commissioning of the VLT instruments since it will include the communication routines needed for the on-line use of MIDAS at the VLT. Because of the importance we have anticipated an extended test period and plan to release the β-release in August. Test reports for the β-release are, as usual, expected to come until the beginning of November. After final verification, the official 95NOV will then be announced and distributed.

Patches for the 94NOV release

Till now, three patch levels for the 94NOV version have been released, 94NOVpl1.0, 94NOVpl1.1 and 94NOVpl1.2. They can be obtained from the anonymous ftp account from the directories 94NOV/unix and 94NOV/vms. More details about these patches can be found in the back of this *Courier.* New patch levels will be announced via the midas-users mailing list (see below).

MIDAS mailing lists

The `midas-users` and the `midas-announce` mailing lists are installed to facilitate discussions about ESO-MIDAS and to inform its users.

The `midas-users` mailing list is un-moderated. Articles will be sent to the e-mail addresses of the subscribers as they are posted. All messages to the list are archived. The archives are available via Majordomo using the `get midas-users <filename>` command (send `index midas-users` in the body of a message to `majordomo@eso.org` for the index of files you can `get`), or via anonymous FTP from host ftphost.hq.eso.org in directory `/midaspub/midas-users`. The archives are broken down by year and month, and are stored in files named `midas-users.YYMM`. The copy of the archive available by anonymous FTP is updated every night at 2am local time.

The `midas-announce` list is moderated by ESO and meant for announcements of releases and patches of ESO-MIDAS. Messages are archived in the same way as for the `midas-users` mailing list and can also be retrieved by ftp.

To subscribe to these mailing lists, send a mail to `majordomo@eso.org` with the following command in the body of your email message: `subscribe midas-users [email]` or `subscribe midas-announce [email]`, where `[email]` should be your e-mail address.

System News

Several important features have been added in the MIDAS system software and will be available in the 95NOV release.

The command `SET/MIDAS` has been upgraded to support data files, MIDAS procedures and executables in different directories. `SET/MIDAS PATH=directory` specifies a search path for MIDAS procedures executed with the `@@` command - similar to the PATH variable in Unix. `SET/MIDAS DPATH=dir` and `SET/MIDAS EPATH=dir` do the same for data files and MIDAS executables (program modules).

All MIDAS applications accept images in FITS format as input, *i.e* you don't need to convert from FITS to the internal MIDAS format before accessing an image. For example, to display the FITS file *myimage.fits* you can use `LOAD/IMAGE myimage.fits` directly. New images, however, are always created in MIDAS binary format. Tables in FITS format are not supported at the moment due to their (possible) very complex structure, *e.g.* one FITS table can encapsulate several MIDAS tables.

Subframes of MIDAS images (specified in the usual MIDAS syntax via `name[x1,y1,z1:x2,y2,z2]') can now be used in all commands instead of the full image. *E.g.* `Midas>fft/freq itest[<,<:@128,@128]` will do an FFT only on the lower-left 128x128 window of frame *itest.bdf*.

Pipes (symbol '|' as in Unix) are implemented to connect MIDAS to Host commands as well as Host to MIDAS commands. *E.g.* `Midas> read/descr lola | $grep CUNIT` will just display the contents of the descriptor CUNIT of image *lola.bdf*. And `Midas> $pwd | write/keyw inputc` fills character keyword *INPUTC* with the current directory string.

As an important step towards On-line MIDAS, which is MIDAS plus additional functionality required for using MIDAS on-line at the VLT, we implemented in the 94NOV release a set of *Communication Interfaces* which provide access to MIDAS in a client/server mode. With these interfaces user processes (clients) can use services of a separate MIDAS session (server) running on the same Unix system or on a different system in a network. The commands `CONNECT/BACKGROUND`, `SET/BACKGROUND` and `DISCONNECT/BACKGROUND` supporting connection and interaction of several MIDAS sessions in parallel are also based on the communication interfaces.

A complete description of these interfaces can be found in the last revision 1.3 of the MIDAS Environment document (June 1st, 1995). An example of a program using these interfaces is also included in the 94NOV release (`/midas/94NOV/monit/communi.c`).

Application News

Upgrades of CCD context

The CCD package has been upgraded. The pipe line processing is now exclusively based on the Data Organizer association table. A separate reduction table is not needed any longer. Also, the flat sky and illumination frames can now be used in the pipe line calibration procedure. As a result of the CCD monitoring programme of the La Silla CCDs the MIDAS test procedures to analyse the test data have been upgraded. For more information about the CCD test package see the previous issue of the *Courier*.

FITLYMAN package

A new package for the analysis of absorption spectra has been implemented in the Lyman context. The context has been developed at the Osservatorio Astronomico di Roma and allows for fitting of Voigts profiles to a large number of absorption lines in a normalized spectrum. A detailed description can be found on page 6.

Echelle package and XEchelle

The Echelle context has been refurbished to support world coordinates, allowing to use different sub-windows of the CCD in the same calibration run. A Graphical User Interface has also been developed for this version of the package. An article describing the new features in more detail will appear in the next issue of the *Courier*. This new version of the package and the GUI will be distributed with the β-release of MIDAS and in 95NOV.

Developments for the IRAC2 camera

A new IRAC2 context is being worked on in order to improve the calibration and reduction tools for infrared imaging. The new package will be geared towards the reduction of IRAC2 data and will contain on- and off-line routines, like sky determination and subtraction, correction for dithered frames, and mosaicing. Like the CCD package the IRAC2 package will use the MIDAS Data Organizer for classifying and association of the calibration and science frames. A basic implementation will be available in the 95NOV release. Comments and suggestions for this new context can be sent to `rwarmels@eso.org`.

Once again: MIDAS Problem Reporting

Rein Warmels, *European Southern Observatory*

Introduction

The previous issue of the ESO-MIDAS Courier (Vol. 4. No. 2, October 1994) contains an overview of the GNATS problem report system that is used by the MIDAS Group to keep track of reported problems. After having worked with the system for about one and a half year and having dealt with over 700 reports we have made a brief evaluation. Below some statistics will be presented about the reports received in the first year of usage. In addition, the new graphical user interface (GUI) for forwarding MIDAS Problem Reports is presented.

Some Statistics

During the period May 1994 till February 1995 the MIDAS account received a total of 364 Report Reports (PRs). In the start the usage of the GUI in the XHELP (button FEEDBACK) was rather disappointing: most of the incoming PRs (more than 90%) were composed using normal editing tools. That situation has slightly improved: currently around 30% of the PRs conform to the GNATS format standard that facilitates automatic handling by the GNATS system and faster processing by the MIDAS Group. With the improved GUI (see below) we hope that this number will further increase.

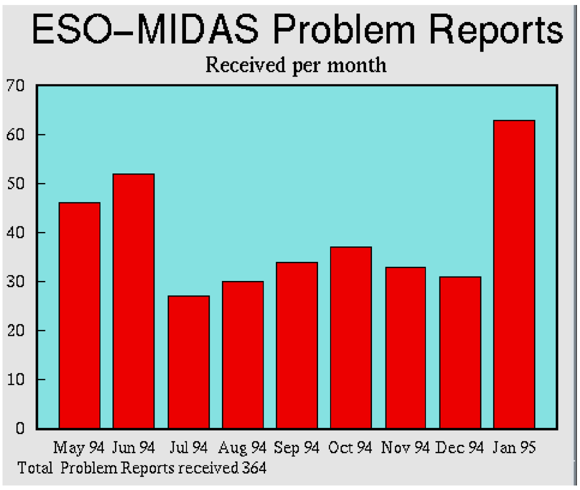

Figure 1.: Histogram distribution showing the number of MIDAS Problem Reports received in the period May 1994 till January 1995.

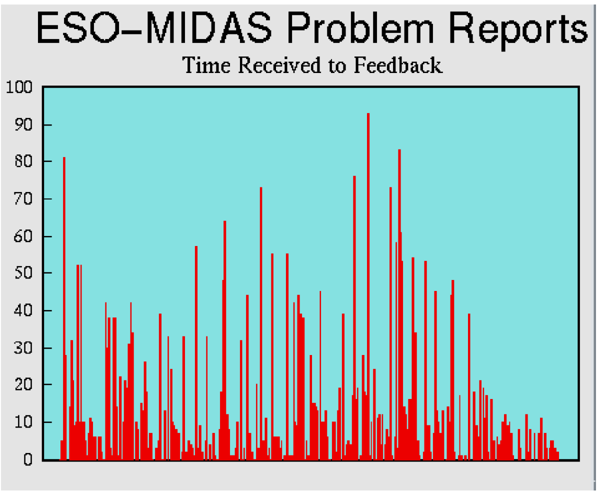

Figure 2.: Distribution of the feedback time, i.e. time between open and feedback state. The feedback time (ordinate, in days) is plotted for all 364 PRs (abcissa) received in the period May 1994 (left) till January 1995 (right).

In spite of the fact that the majority of incoming PRs were actually requests for support, two thirds (66.2%) of the PRs were sent using the class 'sw-bug', which almost certainly resulted from the default setting of the class field. The other scores in the class field were: support: 25.4%, doc-bug: 3.6%, change-request: 3.4%.

Around 50 percent of all PRs come from ESO member states, and 35.7 percent have originators within ESO. The remaining 15 percent of the PRs come from other countries with the US as the main source. The relatively high number of PRs from ESO can be explained by the usage of MIDAS at the telescopes at La Silla and software development by others than MIDAS Group members at ESO.

In Figure 1 the number of PR's in the period May 94 - Jan. 95 is displayed. One can clearly recognise the two peaks in the distribution in June 1994 and January 1995. They reflect the PR's related to the installation of the 94MAY and the 94NOV releases. Apart from these two peaks the average number of PRs is about 30 a month, *i.e.* one or two per day.

Obviously, for MIDAS users all what counts is the actual time that elapses between sending the PR and receiving a solution or at least a work-around. The GNATS PR system, as we use it, knows four states of the PR: open, analysed, feedback and closed. To measure the response time we have used the time between when the PR has been received and has entered the open state, and the time when it has turned into the feedback state. When entering this state the problem is supposed to be solved and a note is sent to the originator for checking.

Figure 2 shows the feedback time for each PR. The high peaks in the figure with feedback times of over 20 days have been identified as PRs that either were duplications in the PR data base Because of the default setting or otherwise incorrect or incorrectly treated PRs. As can be seen in the last part of the figure these administrative problems have been solved after November 94. The average feedback time is, provided the outliers in the figure are

Figure 3. T*he new layout MIDAS Problem Report Graphical User Interface generated by the FEEDBACK button on the XHelp GUI. With respect to the previous version the layout has been simplified. The buttons on the right now show less ambiguous definition of the various PR categories. The context button sets the specific MIDAS context to which the PR applies.*

excluded, around one to two days. Problems that take more time to resolve are generally not related to applications but mostly fall in the category 'system'.

A New Graphical User Interface

Clearly, the introduction of the GNATS system has helped in managing the incoming problems more efficiently than in the past. In addition, the GNATS PR database and tools offer possibilities to query through the PR's and do statistics that improve our understanding of the MIDAS usage and problems. However, the other goal, *i.e.* to have a largely automatic administration of the PR is not reached yet. As mentioned in the introduction, a prerequisite for that is a correct format of the PR. This can be obtained by using the PR GUI in the XHELP GUI using the FEEDBACK button.

We would like, once again, to draw your attention to the XHELP GUI and its FEEDBACK button, in order to facilitate easy forwarding of PR's and automatic processing. We have improved the layout of the interface by eliminating less relevant buttons, and by redefining some category buttons. In addition, the default PR class is now 'support'. PRs related to a context can be indicated by pushing the context button which pops up a listing of the available contexts. Figure 3 contains a reproduction of the new PR FORM. We hope it will enjoy rapidly growing usage.

FITLYMAN: A New Tool for Absorption Spectra Analysis

Adriano Fontana, *Osservatorio Astronomico di Roma*
Pascal Ballester, *European Southern Observatory*

Introduction

The introduction of Echelle spectrographs on 4-m telescopes has allowed the observation of faint objects ($m < 18$) with high resolution ($R = \lambda/\Delta\lambda > 20000$). In the study of interstellar and intergalactic medium, where the absorption lines from tenuous and relatively cold ($T \leq 5\ 10^4$ K) clouds are sought in the spectrum of bright background sources, this resolution allows to observe these narrow lines with an instrumental FWHM lower than their intrinsic width. This makes it possible to derive independently the column density N and the Doppler width from the fitting of the line profile, rather than estimating them from the curve of growth analysis. A typical application of these techniques is the study of the Galactic clouds, that trace the gas content and the metallicity inside the Galaxy, and of the absorption systems found in the spectra of high redshift quasars, that trace the evolution of the primordial structures.

While the reduction of echelle spectra to 1-dimensional calibrated spectra may be accomplished with several packages (like the Echelle context), no similar tool existed for the analysis of their absorption features. This is in part due to the difficulty of designing an application which fulfils the requirements of providing a flexible, easy to operate user interface and allowing to fit large systems composed of hundreds of lines. The fit of heavy element systems, indeed, may require the simultaneous fit of different lines spread far away in the spectrum. For these lines, it is often necessary to test different configurations (*i.e.* number of components and constraints on them) before a satisfying fit is achieved. In the analysis of Lyman-α forest, on the contrary, hundreds of independent lines are fitted, that must be identified by the user Introduction

The lack of a software tool properly tailored for all the different aspects of this problem has moved us to the realization of FITLYMAN, a new package available in the Lyman context of the MIDAS 94NOV release, that we present in this paper.

This code is the evolution of a former version, developed at the Osservatorio Astronomico di Roma (Giallongo et al.). It allows the fit of Voigt profiles to a large number of absorption lines in a normalized spectrum through the numerical minimization of χ^2. Its main features are the flexibility in specifying the line configuration, a fast user interface, the full documentation on all the fit performed (ensuring also recursivity). Special care has been taken to maximize the efficiency of the user operations: FITLYMAN provides recursiveness and allows to recover previous configurations.

Main features

We describe here the most important features of FITLYMAN. For a more complete description of the commands and algorithms please refer to the ESO-MIDAS User's Manual.

Each absorption line is modelled by a Voigt profile, whose parameters are the central wavelength λ_c, the column density N and the Doppler broadening b. In general, the Doppler broadening of a line may be due to the intrinsic thermal broadening $b_K = (2kT/m)^{1/2}$ and to a turbulent motion of the gas $b_T = \sqrt{2}\sigma_T$ where σ_T is the inner velocity dispersion. When transition from different ions are observed, it is possible to discriminate between the two cases: for the thermal component of the broadening, indeed, the b-values of different atoms scale with the masses as $b_{ion1}/b_{ion2} = (m_{ion2}/m_{ion1})^{1/2}$, while in the turbulent case they remain essentially equal. To allow for a complete analysis of the physical conditions of the absorbing clouds, both types of b-values can be accounted for in computing the line profile: each line is thus parameterized by *four* different parameters, (λ_C, N, b_K, b_T) and the Voigt profile is computed using a b-value $b = (b_K{}^2+b_t{}^2)^{1/2}$. This unique feature, combined with the possibility of constraining the different parameters in a flexible way (see below), allows for the simultaneous fitting of different ionic transitions to yield meaningful values for both the turbulent and thermal broadening.

An extensive set of atomic parameters is supplied: it allows for the fitting of several ionic transitions found in absorption spectra, and can be easily expanded by the user himself.

It is possible to fit simultaneously different ions, or from different transitions of the same ion, specifying several links (relations) among them that reflect the physical condition in the cloud responsible for the absorption. The lines can be forced to share the same velocity relative to the earth (redshift), the same column density N, and/or the same b value. Of course, both the thermal (b_K) and/or the turbulent (b_T) component of b will be properly scaled with the masses of the ions. Any combination of lines and links is allowed, the only limit being the maximum number of independent parameters that the program can handle (currently 100). In general, the number of free parameters depends on the kind and number of link imposed: for instance, when the lines are forced to have the same absorption redshift, there is only one actual free parameter (*i.e.* the redshift), not one for any individual central wavelength.

The absolute freedom in specifying links with these rules is a powerful tool when the ionization and metallicity of a system is investigated, and also to correct for possible systematic errors: for instance, two lines far away in the spectrum (*e.g.* $SiII_{1260}$ and $SiII_{1526}$) might be forced to have the same b and N but different redshifts, to allow for some wavelength calibration uncertainties.

The input data have to be stored into a MIDAS table, containing wavelength, normalized flux, standard deviation and instrumental FWHM for each pixel. A tool for the conversion of the usual image data to this format is supplied. Both pixel size and resolution can vary along the spectrum. Gaps may exist in the spectrum.

The numerical minimization is performed using the MINUIT package, which has been developed at CERN. It allows for the simultaneous minimization of up to 100 free parameters (equivalent to at least 33 independent absorption lines). The user can choose among different minimization algorithms or strategies. For instance, some parameters may be fixed at the first iteration, to find a reasonable "first guess" solution, and then released to achieve the final result: this procedure is particularly useful to find meaningful solutions in complex situations. The original MINUIT package can also be used in interactive mode: since this option is not compatible with the MIDAS command language, it has been disabled in FITLYMAN. Nevertheless, all its relevant functions can be invoked from the User Interface: this solution eliminates also the need of learning the MINUIT command language. MINUIT may perform an accurate error analysis, by computing either the covariance matrix and/or more detailed statistics, as non parabolic errors or scans and contour levels in the χ^2 space.

In principle, the results of each minimization depend on several parameters, that specify: *a*) the line configuration, *b*) the regions of the spectrum where the χ^2 is computed, and *c*) the MINUIT commands used. To keep track of these parameters, each minimization is associated to a unique identifier, whose value is stored in the output table. At the same time, the full initial configuration, together with the same identifier, is stored in three “log” tables. At a later time these tables can be either consulted for the interpretation of the results, or used to re-start the minimization. Furthermore, a batch procedure (BATCH/LYMAN) is supplied, which allows for the automatic, non-interactive fit of a whole spectrum using the configurations stored in the “log” tables.

The User Interface of FITLYMAN has been designed to satisfy two different -and almost opposite- requirements: it must allow the user to handle and modify all the parameters needed for the minimization, reducing as much as possible the time required to do it. To fulfil the MIDAS requirements about compatibility with any kind of terminal and environment, we have designed a menu-driven User Interface that uses both the standard I/O devices (keyboard + screen) and the standard MIDAS graphic window.

An example of a working session with FITLYMAN is shown below. It introduces the user to all the functions available in the program, such as modifying the input parameters, loading/saving previous configurations, modifying the set-up options of the program, or accessing the graphic device. The parameters that specify a line configuration may be entered or modified either by answering at a sequence of questions that guide the user to the correct sequence of operations, (to be preferred when new lines are defined) or by

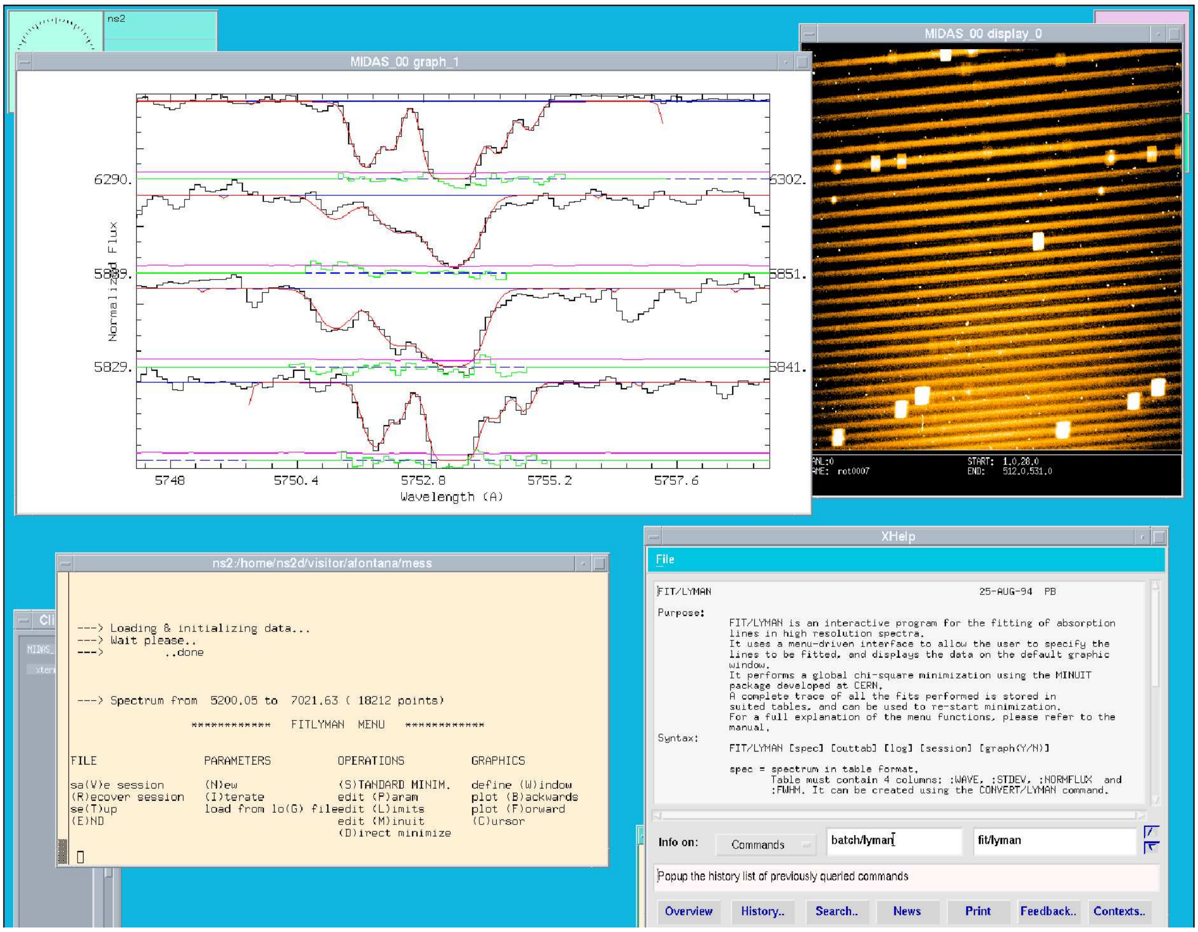

Example of a FITLYMAN session. The Main Menu is shown in the bottom left window. The graphic window over it displays a fit performed on a complex metal system at z=2.768. Lines fitted are, from the top, Al_{1670}, CIV_{1550}, CIV_{1548}, $SiII_{1526}$. CIV doublet has been fitted with three components, lower ionization elements with five. Data from Giallongo et al. 1993 and J. Wampler.

directly accessing to their values (to modify few parameters).

The input spectrum and the fitted profile are plotted on the active graphic window. Optionally, the standard deviation and/or the residuals may be plotted too. You can also use the cursor in the graphic window to enter the wavelengths needed to specify some parameter (*e.g.* "first guess" central wavelength or the intervals over which the χ^2 is computed).

MIDAS on a Metacomputer

M. Stiavelli, I. Lisi, *Scuola Normale Superiore, Pisa*

Astronomers visiting ESO in the eighties would often do so to make use of the computing facilities for data reduction (the so-called MIDAS Room). At the time, only few institutes were able to afford the dedicated hardware needed to properly reduce the large amounts of CCD data produced by a typical observing run on a big telescope. With the advent of UNIX workstations with X11 interfaces, data reduction facilities became much more affordable and the need for such centralized data reduction systems diminished. In the near future a new generation of instrument will become available and start pouring out 8192 x 8192 pixel images. These kinds of images will be essentially intractable on present day workstations and will require either the use of new algorithms and data reduction procedures or the use of larger, not so-common, CPU-servers. Take for instance a Fast Fourier Transform, which is the basis of many deconvolution and filtering algorithms. Its complexity scales as $N \log N$ where N is the *total* number of pixels, *i.e.*, FFTs on 8192 x 8192 pixel frames take 260 times the number of operations needed by a 1024 x 1024 pixel frame. Although one cannot exclude that top line workstations available in a few years will be able to cope with this amount of operations, it is likely that they will be affordable only by a few lucky ones. In fact, it was relatively easy for all institutes to install the first generation of workstations bought with the money made available by the dismissal of central computers. Once everybody had a workstation on his desk, further improvements slowed down and it is easy to check that workstations old by even two or three generations are not uncommon even in the best institutes in Europe and in the US. What will happen to the unlucky ones who will find themselves with 100 8192 x 8192 pixel images from their latest VLT run and will need to reduce the data on a three generations old UNIX box? The answer to the question lies in metacomputing.

A metacomputer is simply a number of inhomogeneous computers interconnected by a hopefully fast network able to solve a single, suitably well-posed, problem. The Italian National Research Council (CNR) and the national phone company (Telecom) have experimentally installed in Pisa and Florence a fiber optics network running at 110 Mbit/s (with a 34 Mbit/s interconnection between the two cities). Such a network has been the ideal ground for experimenting with a metacomputer version of MIDAS. Our goal was to be able to run MIDAS on a metacomputer in a fully transparent way, *i.e.* the user does not know, and does not need to know, which computer is actually computing for him. The system must take care of all this. The user sits quietly on his good old workstation and issues commands as usual, the most computationally expensive commands, however, are not executed locally. The command and the relevant data are sent through the network to some suitably powerful computer which returns the result automatically when finished. Our initial goal is to test the principle rather than produce a complete Metacomputer MIDAS.

Our implementation is based upon the Parallel Virtual Machine (PVM) message-passing library, which is available in the public domain. PVM is written in C language with also a Fortran interface. Such a library provides the necessary tools to write multiplatform applications. PVM has been developed jointly by University of Tennessee, the Oak Ridge National Laboratory, and Emory University. Although other similar products are available, we have focused our attention on PVM because it is available for a large number of platforms.

A metacomputing environment is based upon four basic tools: a Metacomputer Definition tool, a Scheduler, a Master process and one or more Slave processes. We have attempted to define a parametrization for the whole system allowing the system manager to easily redefine the Metacomputer. For each Extended MIDAS *command* one enters one or more descriptive entries in the Metacomputer Definition Table. Each of these entries includes *the command name*, the target *hostlist*, the target *executable pathname*, the maximum allowed *memory occupation,* and a weight parameter defining, in some so far arbitrary scale, the host *performance* in executing the given command. The Metacomputer Definition Tool should in principle be used to maintain a dynamic definition of the meta system to account for new CPUs, machines temporarily unavailable restrictions on the Metacomputer processes on given CPUs, etc.

The Scheduler receives a request from the Master process concerning a given command to be executed. It checks the Metacomputer definition and chooses the most suitable CPU to execute the command based on the given restrictions and, ideally, on a scheduling algorithm aimed at achieving CPU load balance. Once the most suitable entry in the Metacomputer Definition Table has been selected, the Scheduler returns this information to the Master which activates the processes on the (usually) remote host(s) and transfers the relevant data.

The Master process is from the final user point of view simply the Extended MIDAS command which is issued. However, in our scheme, the Master manages inter-process synchronization and data transfer. All dialogue between the Master and the other processes and tools remains fully hidden and the user is totally unaware of which CPU(s) is(are) actually performing the calculation he needs.

The Slave processes are the ones that actually carry out the calculation. In the simplest case, there is only one slave process active. In some cases the calculation itself can be parallelized and therefore several processes can be activated on different hosts.

Let us focus on some concrete examples. A user in Arcetri with a limited CPU power workstation needs to carry out 50 iterations of a Richardson-Lucy deconvolution on a large image. He specifies simply the input frame, the output frame, the Point Spread Function, and the number of iterations. The command first checks the loads of the various CPU servers predefined for the command. The most suitable one could be another workstation in Arcetri or, *e.g.* the Connection Machine 200 at the Scuola Normale Superiore in Pisa. The weight factor specifying the performance of the CM200 should in principle include also an estimate of the larger cost of transferring the data to Pisa and back, so that a slightly slower CPU in Arcetri would be preferred if available. The user will be unaware of the computer actually performing his calculations.

In another case a user might want to carry out a median filter operation, which is available in a fully parallel multi-CPU version. In this case, the Scheduler would find in the Metacomputer definition table both entries for single CPUs and multiple CPUs entries. The most suitable entry therefore could be in principle either a single machine where only one Slave process carries out all the calculations or a number of inhomogeneous CPUs each elaborating a portion of data. The actual splitting and pasting back of the data of the various Slave processes must be carried out by the Master.

At the current stage of our implementation the Metacomputer definition and the scheduling algorithm are static and only very few commands have-been implemented.

Deadline for the next ESO-MIDAS Courier

The next issue of the ESO-MIDAS Courier, Volume 5, Number 2, will be published in the month of December 1995. This issue will include information concerning on-going projects of the MIDAS Group.

Please send your contributions, preferably in computer-readable format, to the MIDAS account or to the editor before October 31 1995 (see back page for mail address).

Expeditions in the MIDAS Jungle: Counting Pixels

Klaus Banse, *European Southern Observatory*

This time we want to have a closer look at the command which most likely is the very first command you execute on a newly acquired frame, namely the `STATISTICS/IMAGE` command. It is quite a complex command with numerous options which we will explore in some detail.[1].

The first one is the *area* parameter which specifies on which part of the image we want to compute the statistics. By default we work on the complete frame; specifying a subframe (in the usual MIDAS syntax) restricts the statistics to only that part of the image. The other options do the statistics on multiple subframes. If the image is currently displayed we can use the option *CURSOR* to choose via a cursor rectangle the subframes/regions of the image where we calculate the statistics in. Another way to work on different subframes is to use the *table* option by giving a table name. This table should contain (standard) columns containing the limits of the regions. Tables with such columns are, for example, generated by the `GET/CURSOR` command with table output. For 2-dim images the *ROW* and *COLUMN* options indicate that the statistics should be done on all rows or columns individually. For 3-dim images the *PLANE* option does the statistics on all xy-planes of the cube.

Since these options have to be distinguished from table names, they have to be entered exactly, *e.g.* *COLUMNS* will be interpreted as using the *table* option with table *COLUMNS.tbl* as input; however, upper- or lowercase does not matter.

The next two parameters *bins* and *lo,hi_exc* determine the no. of bins used in the histogram calculated from the input image. You either enter the no. of bins via *#n* for ‘*n*’ bins or the size of each bin, then the no. of bins is calculated from that value inside the `STATISTICS` command. The default is 256 bins. If you want to avoid outlying pixel values you can specify the beginning and end of the data interval you want to work in via the *lo,hi_exc* parameter which specifies low and high excess values, *i.e.* the end of the lower excess bin and the start of the higher excess bin in the histogram. Then the histogram (and all other statistics) will be calculated only in that interval. The number of bins will be incremented by two, *e.g.* in the default case we have 258 bins in the resulting histogram.

The following *options* parameter consists of 2 characters, the first character, the *statistics level* serves as a flag for different statistics options and the second character, the *display level* specifies the amount of output generated by the `STATISTICS/IMAGE` command.

The default for the *statistics level* is `F` which stands for full statistics producing the following statistical quantities (11 values): minimum, maximum, mean value, standard deviation, 3rd moment, 4th moment, intensity, median, smallest mode, no.of bins, binsize of histogram - calculated over the part of the input image specified via the *area* option. These values are displayed on the terminal and also stored in the real descriptor `STATISTIC` of the image. Level `R` (reduced statistics) will compute only the first 7 quantities of the list above and no histogram is produced for the input image. Setting the level to `S` (short statistics) leads to the calculation of the first 4 quantities only, if the level is `M` (minmax) only the first 2 quantities are computed and again no histogram is generated. Since no histogram is produced, the parameters *bins* and *lo,hi_exc* are ignored for levels `R`, `S` and `M`.

For level `H` (histogram) a histogram is generated and the minimum, maximum as well as the last 4 quantities of the list above are computed.

Except the median value, all the statistical quantities are computed in a straightforward manner. With level `F` and `H`, the median value is approximated via the histogram of the image to speed up the execution time. This works pretty well if there are no ‘runaway’ data, like *e.g.* saturated columns in a CCD frame which are far outside the rest of the data intensity range. In that case you get a meaningless approximation of the median value because of the large size of the bins of the histogram[2].

Therefore, two more *statistics levels*, `G` and `X` are

1. For table freaks the `STATISTICS/TABLE` command plays a similar role.

supported which calculate the precise value of the median, implying that the complete image (or chosen area of it) has to be sorted. Level G is like F (full statistics generating all statistical quantities, the histogram and the exact median; level X is like R (reduced statistics) but additionally produces the exact median.

The difference in execution time for the F or G statistics level is negligible for small frames, *i.e.* images with up to 512*512 pixels. For larger images, however, the exact calculation of the median takes significantly longer: approx. 4 times for a 1024*1024 image, and about 7 times for a 2048*2048 image.

The default *statistics level* is still F because of backwards compatibility. You can change that easily via the CREATE/DEFAULTS command, *i.e.* put the line CREATE/DEFAULTS STATISTICS/IMAGE? ? ? ? GF into your private *login.prg* and STATISTICS/IMAGE will always calculate the exact median[3].

The *display level* is defaulted to F for full display, and can be set to S for short and X for very short display. Setting it to N produces no display.

Besides the descriptor STATISTIC, some other descriptors are written/updated in the STATISTICS/IMAGE command. The integer descriptors WINDOW_FROM and WINDOW_TO are filled with the start and end pixels of the area specified in the *area* parameter. The real descriptor LHCUTS which contains the minimum and maximum of an image in LHCUTS(3,4) is only updated when the statistics of the complete frame are computed. If a histogram is calculated also the real descriptor HIST_BINS and the integer descriptor HISTOGRAM are written to the image. The number of bins and the binsize, the low excess and high excess boundaries and the number of excess bins (= 0 or 2) are stored (in that order) in HIST_BINS. Descriptor HISTOGRAM holds the histogram, *i.e.* HISTOGRAM(1) = number of pixels in first bin, which is the lower excess bin, if *lo,hi_exc* parameter is given; HISTOGRAM(2) = number of pixels in second bin, ..., and HISTOGRAM(n) = number of pixels in last bin, which may be the high excess bin, n = HIST_BINS(1).

Whereas STATISTICS/IMAGE produces one set of statistical quantities when used on the complete frame or single subframe of it, all other *area* options generate this set for each of the specified subframes, *e.g.* for all rows with the *ROW* option. Setting the parameter *outtab* to a table name indicates that all the output should be stored in that table. In particular, the minimum will be stored in a column labelled :MIN, the maximum in a column labelled :MAX and so on. The *outtab* parameter is ignored if the full frame or just one subframe is processed.

The last parameter is the *plotflag* which controls the plotting of the computed histogram (so it is only applicable to *statistics levels* F, G or H). It can be set to P (Plot) or N (NoPlot) to switch the plotting on or off. If no graphics window exists it is created with default sizes. Specifying the plotflag as P with the area options *ROW* or *COLUMN*, is normally not very useful[4].

The values which are stored in descriptor STATISTIC are also saved in the real keyword OUTPUTR so they can be used by other MIDAS applications. Also, the total no. of pixels and the pixel numbers of the minimum and the maximum are saved in the integer keyword OUTPUTI.

Like STATISTICS/IMAGE the command FIND/MINMAX[5] finds the minimum and maximum value of an image and the corresponding pixel coordinates, since it does not bother with other statistical quantities it is usually faster.

The application procedure *histogram* (re)calculates the histogram of an image and copies it to a 1-dim frame. For example, @a histogram tiburon balena will calculate a histogram (with 256 bins) of frame tiburon.bdf and store it in the frame balena.bdf.

The context STATIST provides more specialized statistical applications and tests. SET/CONTEXT STATIST will enable the commands ASSOCIATE/RANK, COMPARE/2SAM, FTEST/VAR, KSTEST/1SAM and STEST/MEAN which work on columns of tables. See the help of these commands for more detailed information.

2. By carefully choosing the excess bins, also in that case one could get a good approximation of the median.

3. If you convince your MIDAS site manager to add that line to the general *login.prg* (in $MIDASHOME/$MIDVERS/prim/proc you have that default for all users at your site.

4. Unless you want to check the drawing speed of your X-terminal

5. For finding other (specific) pixel values, check-out the FIND/PIXEL command.

MIDAS 94NOV Patch File Summary (from the ftp account)

The following text contains a listing (June 1995) of patches for the ESO-MIDAS for the 94NOV and release. The patch files can be obtained from `ftphost.hq.eso.org` *or 134.171.40.2).*
(Message courier:168)

MIDAS 94NOV Release Patch Information

FTP location for 94NOV and patches:
- Ftp host: ftphost.hq.eso.org (134.171.40.2)
- Ftp account: “midas”. Password: “????????”
- Directory: 94NOV
- Files: README.vms
  README.unix
  94NOVpl0/patches_94NOVpl0.1.tar.Z (593 KB)
  94NOVpl1.tar.Z (11 MB)
  94NOVpl1/patches_94NOVpl1.1.tar.Z (150 KB)
  94NOVpl1/patches_94NOVpl1.2.tar.Z (305 KB)

Notes:
- 94NOVpl0.0 is not available any more by ftp.
- patches_94NOVpl0.1 upgrades your old 94NOVpl0.0 to 94NOVpl1.0
- 94NOVpl1.tar.Z is the complete MIDAS distribution for 94NOVpl1.0
- patches_94NOVpl1.1 is an upgrade for 94NOVpl1.0
- patches_94NOVpl1.2 is an upgrade for 94NOVpl1.1

The next patch release 94NOVpl2.0 will be available in July 1995.

New in 94NOVpl0.0:

Compared to 93NOV:
- ANSI-C and POSIX-1 compliance for MIDAS Standard Interfaces.
- 64bits CPU architectures supported (DEC/OSF1).
- PLOT package in ANSI-C, including 3-dim tables.
- Redirection of Input/Output for MIDAS commands supported.
- Refurbishment of the DO (Data Organizer) package.
- The GUI for the DO package.
- Access to tapes: implementation of the Generic tape driver.

Compared to 94MAY:
- The GNU readline library as the default line-editor.
- Data Decompression on the fly for compressed FITS files.
- Access to Remote Tapes with Machine-independent Data Structures.
- Upgrade of the CCD context.
- Upgrade of ECHELLE and LONG contexts.
- The GUI for the IRSPEC package.
- A New Astrometry context.
- GNATS: The New Problem Report System.
- Upgrade of GUI XHELP including a problem report form.

- Upgrade of Display Subsystem.

New or fixed in 94NOVpl1.0:

- Compilation problem for “spec” package solved for VMS.
- Disabling interrupts for background processes (Linux)
- Help on line for Linux and Solaris.
- Script “cleanm.sh” works for OSF/1.
- Bug fixed in EDIT/TABLE on VMS.
- LD_LIBRARY_PATH not necessary any more for Solaris.
- Simplification of CCD keyword initialization.
- Correction for CCD read out direction (COLUMN/ROW).
- Fix LOAD/LUT with non-interpolated LUTs.
- Error on Solaris solved for EXTRACT/OPTIMAL in the ECHELLE context.
- Segmentation violation fixed for IBM/6000 with CREA/GUI HELP.
- Corrected long-slit command RESPONSE/LONG.
- New GUIs: XFilter. Requires calibration data.
- Speed up LOAD/IMAGE.
- More precise help text and error message for DECONVOLVE/IMAGE.
- Bug fixed for APPEND option in OUTTAPE/FITS on OSF/1.
- Initialize IDIDEV correctly.
- STATIST/IMAGE with cursor option had strange character trailing in output (Solaris 2.3).
- Make parameter type check in procedures more robust.
- Add support for zoom windows to the EXTRACT/TRACE command.
- Modify output format for values to fix problems on IBM/AIX
- Skip warnings on missing files when using ‘*’ in INTAPE command.
- Fix setting of graphics cursor if display window open.
- CENTER/GAUSS can handle data with low values (e.g. 10e-15)
- Correction row/col numbers on top of plot.
- Monitor does not print “ERROR IN CHILD” for executables located on midwork.
- New GUI: CREATE/GUI FILTER (calibration data in calib directory)

New or fixed in 94NOVpl1.1:

- Tutorial for the DO package reads data from ./demo/data instead from ./demo/raw
- Tutorials for WAVELET and OPTOPUS accept the AUTO argument.
- Fix LOAD/IMAGE for images of data type “unsigned short int” with scaling factor > 1.
- MIDAS can be compiled with GNU C compiler on Silicon Graphics.
- Fix ‘active’ image channel settings.
- Removing obsolete commands in 94NOV.
- Test for Fortran to C interface more robust.
- Cosmetic changes in the inmidas script.
- AUTO option for the TUTORIAL/ECHELLE command.
- Fix problem with param P3 only ‘nomem’ without mem-sizes in CREATE/DISPLAY.
- Increase max. size of PSF to 101x101 pixels.
- Enable input of ASCII data files with very large record size.
- Correction for upper/lower case in plot device names.
- Bug fixed of the command MODIFY/GCURSOR for VAX/VMS.
- Corrections for NULL device use in plot packages.
- Fix MAGNITUDE/... commands when bright objects are in the background.

- Fix EXTRACT/CURSOR with xpix,ypix specified.
- Fix for incorrect descriptor TBLCONTR for table plotting.
- Fix for lower case display and graphics window names.
- Bug fixed for assignment of incorrect plot window.

New or fixed in 94NOVpl1.2:

- Fix no. of pixels for loading images with scale < -1.
- Avoid MINMAX calculation for LOAD/IMAGE cuts=a,b
- Fix problem with catalog option in SUBTRACT/ICAT
- Fix problems with positive declination in COMPUTE/AIRMASS
- Help file in the ECHELL95 package was missing.
- Include /opt/SUNWspro/lib in LD_LIBRARY_PATH to install MIDAS
- Correction in the command FILTER/ECHELLE
- Bug fix for MODIFY/GCURSOR
- Bug fix for GET/GCURSOR output format for DECterm windows
- Bug fixed for XHELP on PC/Linux.
- Bug fix for overplotting coordinate box and contours on display window
- Implementation of 2dim gauss fitting with calculation of angle of major axis (CENTER/IQE).
- Solves bug in COPY/DISPLAY for SG/IRIX
- Modifications in the inmidas script to use Tektronik emulator.
- Correct the command COPY/TABLE for tables in RECORD format.
- Bug fix for the INTEGRATE/APERTURE command.
- Another bug fix for the command MODIFY/GCURSOR
- make_options for HP does not include -lPW library.
- CREATE/GUI IRSPEC added to the list of GUIS.
- Fix for plotting of error bar when using logarithmic axes
- Generate MIDAS shared libraries on IBM/RS6000 AIX
- Generate MIDAS shared libraries on SG/IRIX 5.3
- Correction in the XHelp resource file for X11 R4.
- Bug fixed in Long context command ERASE/LONG.
- Keep cursor rectangle size for different applications
- Fix LOAD/IMAGE for 1-dim frames
- Command EXECUTE/CATALOG: Do not ask for confirmation with delete option
- Make CENTER/GAUSS more robust.
- Bug fixed in OUTTAPE/FITS with the APPEND option on HP and Linux.
- Fix COMPUTE/IMA res="blabla+12.bdf" - 20
- Correct keyword naming (max. 8 chars).
- Bug fixed in MODIFY/GCURSOR for 1-dim frame.
- Fix column-area mode with exact median for STATIST/IMAGE.
- Make the IDIserver more robust.

The ESO-MIDAS Courier is published twice per year by the ESO-MIDAS Group of the European Southern Observatory.

Contributions as well as suggestions and comments are invited and can be sent to the editor. In particular, authors of ESO-MIDAS application software that would be of general interest for the ESO-MIDAS community, are invited to make this software available.

If you are not on the mailing list and want to receive future issues of the Courier, contact the editor. More updated information about ESO-MIDAS can be obtained from the ESO-MIDAS Bulletin Board, accessible via anonymous ftp, or via WWW

The ESO-MIDAS™ Courier:

Editor: Rein H. Warmels
Typeset by FrameMaker ™
Published by the European Southern Observatory
Karl-Schwarzschild-Straße 2,
D 85748 Garching bei München,
Federal Republic of Germany

ISSN 1018-3051

The following MIDAS Support services are available for sending suggestions, comments or to obtain help when problems arise.

Telephone:	+49-89-32006456
Telex:	5282822 eso d, attn. ESO-MIDAS
Telefax:	+49-89-3202362, attn. ESO-MIDAS
EARN:	MIDAS@DGAESO51
Internet:	midas@eso.org
Anonymous ftp:	ftphost.hq.eso.org (134.171.40.2)
WWW	http://http.hq.eso.org/ midas-info/midas.html