Four years of Service Mode observing at the VLT

performance and user feedback

On the night of 3rd April 1999 the ESO Very Large Telescope (VLT) started regular science operations with Service Mode (SM) observations using ISAAC at Antu, the first Unit Telescope. Nineteen days later, SM observations were executed for the first time with FORS1. Over the four years following that date, three more Unit Telescopes have joined Antu, five more instruments (FORS2, UVES, NACO, FLAMES and VIMOS) have joined ISAAC and FORS1, and SM observations are now being carried out during more than 50% of the observing time available on Paranal.

It thus seems timely to present an overview of the performance of SM at the VLT, and how it is judged by its users. In this article we provide such an overview, discuss some lessons learned during the last four years, and present a summary of the main results collected from the VLT SM users community through the extensive questionnaire that ESO released in September 2002, where a broad range of the aspects of SM observing, from the Phase 1 proposal preparation to the scientific evaluation of the results obtained, is covered.

F. Comerón[1], M. Romaniello[1], J. Breysacher[2], D. Silva[3], G. Mathys[4]

[1]ESO, User Support Group, Data Management Division;
[2]ESO, Visiting Astronomers Section, Office of the Director General;
[3]ESO, Data Flow Operations Group, Data Management Division;
[4]ESO, Paranal Science Operations

Service Mode (SM) observing has been perceived by ESO since the early days of VLT operations planning as a key component in optimizing the scientific return and the operational efficiency of the VLT. It is a concept deeply embedded in the VLT end-to-end science operations model (Quinn et al. 1998, 2002; Mathys et al. 2002). Experience with SM was already gained in February 1997 at the NTT, as soon as that telescope entered into operations again after the "*Big Bang*" (Wallander and Spyromilio 1997) that provided a field testing both for the VLT control system and for much of the VLT end-to-end system.

The current effort invested in SM operations at the VLT can be illustrated by some recent operational statistics. A total of 2909 SM Observation Blocks (OBs), corresponding to 200 observing runs submitted for execution by the five instruments in operation during Period 70 (1st October 2002 - 31st March 2003), were completed within or nearly within the user-specified constraints, amounting to 1685 hours of execution time, without counting the time spent in calibration observations provided by the observatory or the time used in the execution of OBs that turned out to be outside specifications. In the ongoing Period 71, with VIMOS and FLAMES now also in operations, 2043 OBs for 1559 hours of execution time have already been completed under or near user-specified conditions at the time of this writing (mid July 2003).

SM is the most requested observing mode at the VLT, as Figure 1 illustrates. After a steady increase in the SM vs. Visitor Mode (VM) demand ratio over the first two years of VLT operations, the pressure has now stabilized at a ratio of about 15.5 hours requested in SM for each night in VM, or a ratio of 1.7:1 if we take an average duration of 9 hours per night. Although this would directly translate into a community demand of approximately 63% in SM vs. 37% in VM, constraints and limitations resulting from the current Garching-based staff available for support of the front and back ends of the operations has forced ESO to move towards a 50%-50% share, which is mostly achieved by moving selected SM programs to VM. Such a share also ensures that Paranal staff astronomers and fellows keep in direct contact with the astronomical community, that ESO receives external feedback based on first-hand experience of visiting astronomers about the Paranal instrumentation and operations, and that a certain level of know-how about the actual observation process and about observatory operations is maintained in the community.

How Service Mode works: Long-, Medium-, and Short-Term Schedules

The Long-Term Schedule combines the VM observing runs and the SM periods in the best possible way for every observing semester. The main goal of SM planning at the VLT is to make possible the realization of the rationale behind flexible scheduling: to adjust to the prevailing external conditions by executing those programmes that can make the best use of them, and to ensure that each programme is carried out under the conditions that best suit it following a priority scheme that gives precedence to those programmes that received the highest scientific rating by the Observing Programmes Committee (OPC). This leads to the separation of SM runs into the three priority classes A, B, and C, the last one corresponding to low priority runs that can be executed under relatively poor conditions. The allocation of the priority classes is made by the Visiting Astronomers Section in strict accordance with the OPC scientific ratings, taking also into account technical feasibility, target distribution on the sky, and user-specified constraints on the execution conditions, as has been described in detail by Silva (2001).

The outcome of this process is the SM Long Term Schedule (hereafter LTS). Once it is ready and the list of scheduled runs is finalized, the Visiting Astronomers Section notifies the Principal Investigators via e-mail, giving them access to password-protected webpages where detailed information on the time allocated to each of their runs, as well as possible OPC comments, are given. This marks the beginning of the Phase 2 process. Users have to provide at Phase 2 the set of OBs fully defining their observations, prepared with the Phase 2 Preparation Program (P2PP) produced by the Data Flow Systems group of the Data Management Division, and with instrument-specific preparation software maintained by Paranal Observatory. Ancillary information, such as specific execution instructions, internal priorities, and finding charts, must also be submitted at this time. Detailed information on the Phase 2 Preparation Process can be found at *http://www.eso.org/observing/p2pp/ServiceMode.html*. The Phase 2 packages are reviewed by the User Support Group for compliance with SM policies, technical correctness, and consistency with the Phase 1 information as approved by the OPC. Once certified, the runs are included in the Medium Term Schedule queues that are provided daily to Paranal Science Operations, and which form the basis for the Short Term Schedule (STS). The STS is the actual sequence of observations carried out on a given night, and is prepared by the astronomer in charge of SM observing based on target visibility, external conditions, run priority, and possibly other factors such as instrument mode availability or timing constraints.

Once an observing run is completed, the Data Flow Operations group prepares a data package containing all the science data obtained for the run, the corresponding calibration data and, for most instrument modes, also pipeline-reduced data useful for a quality assessment of the science data and, to a limited extent, for their scientific analysis. The Science Archive Facility produces the media (normally, CD-ROMs or DVDs) containing the data package and sends it to the Principal Investigator. Data packages are also produced for non-completed runs at the end of the period. Fast-track procedures for the early delivery of data have been set up to deal with Target of Opportunity runs and for pre-imaging runs to be followed by multi-object spectroscopy. Under special circumstances, it is also possible for Principal Investigators to retrieve raw data from the archive while the observing run is still being carried out if strong scientific reasons require it.

Some recent run completion statistics

All the time in principle available for Service Mode observations during a given semester is distributed among the priority A and B runs. This time would be really available for scientific observations only under ideal conditions, including no technical and weather downtime. The unavoidable deviations between such ideal conditions and the reality, which we describe in more detail below, naturally lead to a certain level of oversubscription of the actually available time. The goal of ESO at the time of executing the priority A and B runs is thus that all the runs in class A, and a large fraction of those in class B, are completed at the end of an observing semester. Indeed, as far as class A runs are concerned the completion fraction is virtually 100%, both due to their higher scheduling priority and to the fact that non-completed priority A runs are eligible to be carried over to the next period unless they can be considered as essentially completed. In the recently completed Period 70, 17 runs out of 71 obtained carryover status. While this is 24% in number of runs, it is actually less than 5% of the time initially allocated to them, since a large fraction of the observations for the runs to be carried over had been already completed by the end of the Period.

The lower scheduling priority of class B runs, designed to make them absorb the impact of the deviations between the ideal assumptions used in preparing the LTS and the reality encountered during the observing semester, implies that a certain number of runs in this class cannot be completed the end of that period, as we have said above. On the other hand, the occurrence of less-than-ideal observing conditions, or the occasional completion of all class A and B runs with targets at a given right ascension interval, gives a chance for the execution of observations of priority C runs.

The statistics on the completion of class B and C runs for the most recently completed Period, given in Figure 2, show that the goal of completing a substantial fraction of priority B runs is being met. An interesting feature visible in Figure 2 is the large fraction of completed class B runs followed by a tail of incomplete runs, whose number actually increases towards the lowest completion fractions. The rea-

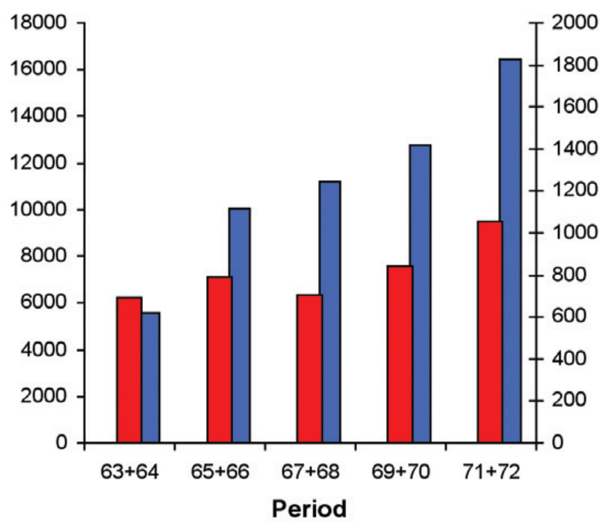

Figure 1: User demand of Visitor Mode (in nights, red bars) and Service Mode (in hours, blue bars) since the beginning of VLT operations. The periods have been binned in couples to remove seasonal differences in the hours/nights equivalence. The right and left vertical scales are directly comparable if we assume a night length of nine hours averaged over the year. After a steady increase of the SM/VM ratio over the first four periods, the ratio has kept roughly constant at 15.5 hours in SM requested for each night in VM since Period 67.

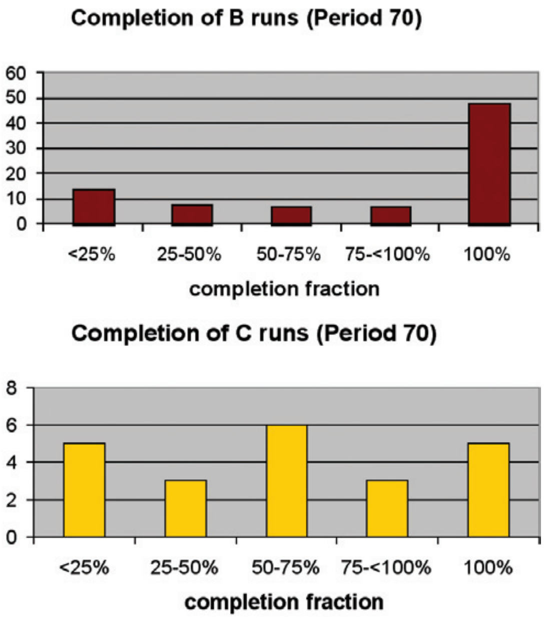

Figure 2: *Degree of completion of class B and C runs at the end of Period 70. The priority given in the STS to completing entire runs as opposed to uniformly distributing time over all runs results in a large number of completed runs at the expense of a fairly large number of runs with completion fractions below 50%. Such priority cannot be applied to the "filler" runs of class C*

son lies in the practical application of the principle that the scientific objectives of an observing run are most likely to be achieved only if all its observations are completed. Therefore, at the time of building up the STS, priority is given to completing as many class B runs as possible, rather than to obtaining observations for runs likely to be left incomplete, thus explaining the existence of class B runs that were either not started or obtained less than 25% of their observations in Period 70 (such priority is harder to apply to class C runs given their "filler" character, explaining why such a trend is not visible in the lower panel of Figure 2). In terms of the time devoted to priority B runs, 642 hours were allocated in priority B; 421 hours (66%) were actually spent on priority B OBs executed within constraints; 251 h (39%) corresponded to runs that were completed at the end of the Period (since short runs are easier to complete, the fraction of completed B runs is 57%).

Let us recall that the LTS is a forecast on the scheduling of the observing semester based on the Phase 1 information given in the proposals, coupled with a model of the observing conditions during the semester. In practice, a variety of factors contribute to create differences between the LTS and the actual outcome at the end of the period. Most of these factors tend to decrease the time actually available for the execution of scientifically valuable SM observations:

A) Weather and technical downtime are not factored into the LTS preparation process. On Paranal, the total amount of downtime ranges typically between 10 and 15% of the available time, and is largely dominated by weather losses.

B) The actual pressure on each right ascension interval is known only at the end of Phase 2, i.e., after the LTS has been prepared. Deviations from the right ascension distribution assumed at the time of preparing the LTS can be due to a variety of reasons:

① The time allocated is sometimes reduced at the time of reviewing the proposal, leaving to the user the choice of a subset of the proposed targets to observe out of those listed in the Phase 1 proposal.

② The actual time that users planned to spend on each target is often not known in detail at Phase 1 time. Furthermore, users occasionally underestimate overheads or overestimate instrument performance at Phase 1, and thus have to drop targets at Phase 2 to keep within the allocated time and intended S/N.

③ Target change requests at Phase 2 may be approved, as long as they are based on convincing scientific arguments, are consistent with the goals of the project, no conflicts with other runs exist, and the impact on the schedule is small or moderate.

C) Director's Discretionary Time (DDT) and Target of Opportunity (ToO) runs are needed to give the VLT the capacity to react to sudden or unpredictable astronomical events or to very recent scientific developments, and have a random impact on the schedule.

D) The amount of time that will need to be devoted to carryover of runs into the upcoming semester, and the distribution on the sky of the targets of these runs, are not known when preparing the LTS, as this happens long before the current period concludes.

E) The completion of highly rated runs with very demanding observing conditions constraints is usually very expensive in terms of excellent weather conditions, and involves a "hidden" overhead that can greatly increase the actual time needed to complete a highly rated run within constraints. Observations started within constraints are sometimes finished outside, for example because the seeing or the transparency worsened during the execution. According to the current ESO policy of considering an OB as completed only if the user-specified constraints were fulfilled, such OBs need to be repeated, but the time that was spent on them (generally still in reasonably good conditions) is lost to other runs. The difficulty in satisfying demanding constraints over a long period of time is one of the main motivations for imposing a maximum duration of SM OBs to one hour, a limitation that often increases the execution overheads for individual programmes but that greatly increases the overall efficiency of VLT science operations.

F) Finally, and related to the previous item, extremely good observing conditions occur rather rarely, and therefore the relative fluctuations on the amount of time in which they occur are large. When we enter the range of conditions that occurs in only a very small number of nights on a given semester, small number statistics come into play and the risk that a large fraction of the allocated time under those conditions may actually be unavailable becomes very real.

User Feedback

As we have noted above, the preferential choice of SM over VM at the VLT can be taken as a direct indicator of user satisfaction with current ESO Service Mode execution. Feedback from SM users over these last years, either by direct communication with ESO or via the Users Committee, has been essential to prioritize improvements in the system, such as the major reengineering of the P2PP software or the constant upgrading of the user support tools. In addition, the extensive VLT SM web-based questionnaire made available to the community in September 2002 provides a way to collect comprehensive feedback on the rating that the community makes of the whole SM mode process, covering the proposal preparation, the Phase 2 preparation, the programme execution, the evaluation of the data obtained, and the fulfilment of the science

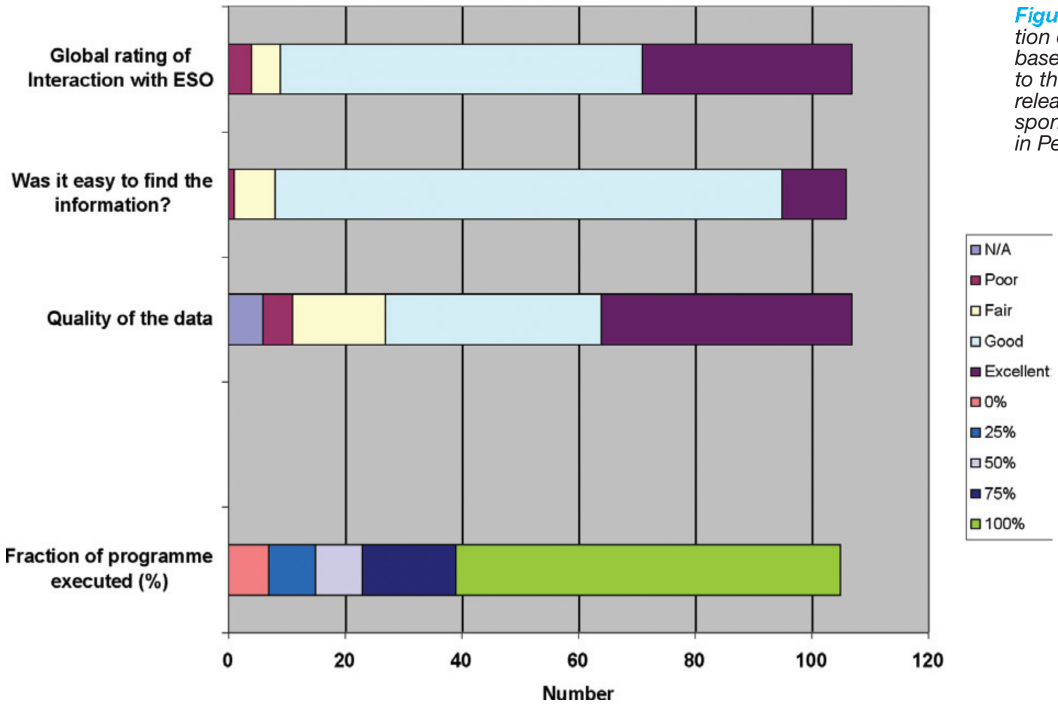

Figure 3: Summary of the evaluation of Service Mode performance based on the 117 answers received to the Service Mode questionnaire released in September 2002 corresponding to runs that obtained time in Periods 63 through 68.

goals. Feedback via this questionnaire was directly requested from the Principal Investigators of runs having obtained SM time in Periods 63 to 68, i.e., covering the entire first three years of VLT operations. The questionnaire can be found at *http://www.eso.org/org/dmd/usg/survey/sm_questionnaire.htm*, and all SM users are encouraged to fill it in. In particular, the note sent to the PIs of SM runs near the end of their proprietary data period reminds them to fill in the questionnaire, as this is the time when a complete scientific evaluation of the data obtained is most likely to have taken place, and when feedback is most valuable to ESO.

A complete report on the results collected from users who had approved SM runs in Periods 63-68 and completed the questionnaire can be downloaded from *http://www.eso.org/org/dmd/usg/reading/smquest_report.doc*. Here we summarize the main points of interests of those results, based on the responses that have been received from 74 unique users on 117 runs, out of the 371 users to which time was awarded in 886 runs between Periods 63 through 68. Most of the responses received (55%) concern runs that obtained time in Periods 67 and 68 (April 2001 to March 2002), and only 14% refer to runs corresponding to the first year of VLT operations. Thus, the results presented here mostly represent a recent evaluation of VLT SM.

The results providing an overall rating of the process (Figure 3) yield a percentage of 92% of runs in which the interaction with ESO is qualified as good or excellent. Concerning data quality, the most frequent rating of the data is "excellent" (40% of the runs), with a rating of "good" in another 35% of the runs, although room for improvement is demonstrated by the 20% of runs for which data were deemed to be of poor or fair quality.

It is interesting to review also the responses given on specific areas of the SM process:

A) Phase 1 receives a very positive global consideration. The Call for Proposals was found to be good or excellent by all the users who replied with one single exception, and the Proposal submission process was considered as good or excellent in 86% of the responses. The web-based documentation provided at Phase 1 on the call for proposals, on the ESOFORM package preparation, and on the instruments (by means of the instrument User Manual) receives over 85% of good-to-excellent marks. At the time of preparing their proposals, 72% of the users find it clear how to compute the overheads on the basis of the documentation provided, and 80% consider the Exposure Time Calculators to be good or excellent.

B) On Phase 2 preparation, 30% of the users found the time between the notification of the time allocation and the Phase 2 package submission deadline too short. Although this time has generally been four to five weeks, it has been possible to extend it to six weeks in Period 72, and we thus believe that this point has been mostly addressed now. The Phase 2 instruction webpages, both general and instrument-specific, and the instrument manuals, receive similar approval rates with 86%–90% of responses giving ratings of good to excellent. However, the still relatively low percentage (8%–12%) of the responses giving a rating of "excellent" are a good reason to continue the sustained effort to improve the documentation.

C) P2PP, one of the most visible software products currently provided by ESO to its user community, is now rated as good or excellent in between 75% and 78% of the cases regarding its installation, user manual, usability, and functionality. Its installation receives the highest marks, being considered "excellent" by 22% of the response. Improvements in both functionality and documentation introduced over the last year will hopefully increase these marks in the near future.

Among other tools used for Phase 2 preparation, *Skycat* receives a very positive consideration: 58% of the users consider its functionality as good, and 35% as excellent. Over half of users (55%) use it to produce their finding charts.

At the time when the survey was released, the FORS Instrumental Mask Simulator (FIMS) was the only auxiliary, instrument-specific tool that had been released for use in the preparation of Phase 2. Its usability and functionality were considered as good or excellent by 55% and 58% of the FORS users, respectively. Several other auxiliary preparation tools

(*NAOS PS* for NACO, *VMMPS* and *Guidecam* for VIMOS, and *FPOSS* for FLAMES) have been released in the last year.

D) The assistance of the User Support Group astronomers in preparing the Phase 2 package has been considered good in 42% of the cases and excellent in another 37%. The Phase 2 review process, also carried out by the User Support Group, is rated as excellent in 36% of the responses and good in another 53%. Overall, the phase 2 process is considered as good or excellent in 89% of the answers.

E) Most of the users (88%) check the progress of their SM observations during the period through the webpages, and a similar percentage find the information clear, up-to-date, and complete. However, almost half (46%) of the users who answered the survey complained that the run progress information was not easy to find. This is one of the cases in which the questionnaire allowed us to identify a shortcoming perceived by many users that had passed unnoticed to us! We have tried to make this information more visible now by including links to it from more ESO webpages.

F) The SM data package, prepared once a run has been completed or terminated, is globally considered as good in a wide majority (81%) of the cases, with an additional 11% that rate it as excellent. The amount of data seems adequate to 92% of users, and the data volume is unanimously considered to be manageable. The typical delay of four weeks between the completion and the delivery of the data is found to be acceptable for 88% of users, although ESO is studying ways to speed up the process for users who need earlier access to the data.

G) The quick-look science data are mostly found to be of good (55%) or excellent (8%) quality. The pipeline products are generally considered as useful (71% of responses), but were directly used for science in only 18% of cases. Shortcomings identified in this area thanks to the SM questionnaire are a perceived insufficiency of information on the reduction process (72% of responses) and the limited usability of the quality control parameters. Slightly more than half of the users (53%) visited the Quality Control webpages, where they generally found the information to be useful (94%) and up-to-date (79%). Also in this area, the Data Flow Operations group has made a considerable effort in making data reduction and quality control documentation available through its Web pages (*http://www.eso.org/qc*; see also Hanuschik and Silva, 2002, and Hanuschik et al., 2002) in order to increase the usefulness of processed data products to the end users.

H) Finally, most of the answers received to the questionnaire (53%) corresponded to runs for which the data analysis was completed, while it was still in progress in another 38% of the cases. A major concern in this respect is that only for 10% of the runs are the calibration plan data considered as good: by far the dominant rating is fair, with 78% of the answers. Also, the information on the calibration plan is rated as only fair by 91% of the answers, and none of them gives the rating "good" or "excellent". Nevertheless, the data quality was rated as excellent regarding the fulfilment of the scientific goals in 49% of the cases, and good in another 28% (note that these percentages include runs that were only partially completed, thus preventing the full achievement of their scientific goals).

Concluding remarks

In summary, although there is still room for improvement in a number of areas, the users' satisfaction with the large majority of the services that ESO provides in Service Mode observing is high. Needless to say, a fair measurement of the performance of Service Mode and its perception among the users community relies on a continued and abundant feedback to ESO; we thus encourage SM users to continue providing ESO with such feedback through the Service Mode questionnaire at *http://www.eso.org/dmd/usg/survey/ sm_questionnaire.htm.*

On ESO's side, much has changed and improved since Service Mode observations were started in 1997 at the NTT and in 1999 at the VLT: new and better tools are available, more experienced and skilled staff are in charge of planning and executing the observations, policies and procedures leading to an efficient use of the Service Mode time have been defined and evolved over time, a clearer picture of the advantages and limitations of Service Mode observing has emerged, and many lessons have been learned, both obvious and subtle. On the users community side, the principles underlying Service Mode observing have become better appreciated and increasingly used to the advantage of the scientific goals of the projects, while new projects with new demands keep pushing the boundaries of its possibilities. The high standards expected by the community, its sustained high demand for Service Mode, and the valuable feedback received from it, together with the challenges set by new and more complex instruments recently entered into operations or soon to do it, drive Service Mode as an evolving process at the VLT and as an essential ingredient in maintaining its high scientific productivity, at the same time that it explores operations paradigms that will be essential to the success of ALMA and OWL in the coming decades.

Acknowledgements

The success of Service Mode operations at the VLT owes much to many individuals who are or have been involved in multiple aspects on both sides of the Ocean: the software engineers who develop tools; the staff astronomers and fellows who carry out the observations and sometimes provide lessons learned the hard way; the telescope and instrument operators; the user support astronomers; the scientists in charge of scheduling, maintenance of reporting tools, quality control and data distribution; the people who have developed and refined its underlying concepts with time; and, last but not least, the broad community of Service Mode users.

References

Hanuschik, R., Silva, D., 2002, *ESO Messenger*, 108, 4.

Hanuschik, R. W., Hummel, W., Sartoretti, P., Silva, D.R., Proc. *SPIE*, 4844, 139.

Mathys, G., Gilmozzi, R., Hurtado, N., Kaufer, A., Lidman, C., Parra, J., 2002, Proc. *SPIE*, 4844, 25.

Quinn, P.J., Albrecht, M.A., Ballester, P., Banse, K., Chavan, A.M., GrosbǸl, P., Péron, M., Silva, D.R., 1998, Proc. *SPIE*, 3349, 2.

Quinn, P.J., Comerón, F., Gilmozzi, R., Mathys, G., Péron, M., Pirenne, B., Silva, D.R., 2002, Proc. *SPIE*, 4844, 1.

Silva, D., 2001, *ESO Messenger*, 105, 18.

Wallander, A., Spyromilio, J., 1997, Proc. *SPIE*, 3112, 9.