projects have taught the lesson that highly redundant safety and diagnostic systems are necessary to have smooth operations. Therefore the LGSF becomes a rather complex and elaborated system, especially to fulfil the requirements of automatic operation with moderate operator assistance.

In order to ensure the timely completion of the project, we have separated the design and installation phases of the Laser Clean Room, which requires heavy infrastructure work, from the remainder of the LGSF systems. The LCR has been placed on fast-track, and will be erected in February 2002, to minimise the impact on the UT4 telescope operations.

The critical items to be procured are the fast Launch Telescope and the PARSEC laser. The R&D activities related to the LGSF project are the PARSEC laser (MPE), the fibre lasers for MCAO and the single mode fibre relay (ESO).

The project status at the time of this writing is:

• LGSF Preliminary Design passed, identified perceived risk areas, identified back-up paths.

• Placed the contract of the Laser Clean Room and its support structure.

• Specialty fibre contract issued, 1st prototype received. Photonic Crystal Fibres received. Fibre relay tests on the way.

• *Launch Telescope:* feasibility assessed for SiC substrates and structure, other composite or lightweight optical materials are being explored. LT is out for enquiry, together with mechanics.

• Breadboard of the Fibre input subsystem assembled and under test.

Operation plan and LGS light-pollution policy for the Paranal observatory drafted, under discussion.

References

1 Bonaccini, D., Rigaut, F., Dudziak, G. and Monnet, G.: Curvature Adaptive Optics at ESO, in SPIE Proceedings of the International Symposium on "Astronomical Telescopes and Instrumentation", SPIE Vol. 3353. Paper no. 131, 1998.

2 Bonaccini, D., Rigaut, F., Glindemann, A., Dudziak, G., Mariotti J.-M. and Paresce, F.: Adaptive Optics for ESO VLT-Interferometer, in SPIE Proceedings of the International Symposium on "Astronomical Telescopes and Instrumentation", SPIE Vol 3353. paper no. 98, 1998.

3 Bonaccini, D.: The Paranal Model Atmosphere for Adaptive Optics, VLT-TRE-ESO-11630-1137, ESO Technical Report*, June 1996.1

4 Bonaccini, D., Hackenberg, W., Davies, R., Rabien, S. and Ott, T.: VLT Laser Guide Star Facility: First Successful Test of the Baseline Laser Scheme, *The Messenger* No. 100, Dec. 2000 – in http://www.eso.org/gen-fac/pubs/messenger

5 Hackenberg, W., Bonaccini, D. and Avila, G.: "LGSF Subsystems Design Part I: Fibre Relay Module", *The Messenger* No. 98, Dec. 1999 – in http://www.eso.org/gen-fac/pubs/messenger

6 Bonaccini, D., Hackenberg, W., Cullum, M., Quattri, M., Brunetto, E., Quentin, J., Koch, F., Allaert, E. and Van Kersteren, A., "Laser Guide Star Facility for the ESO VLT", The Messenger No. 98, Dec. 1999 – in http://www.eso.org/gen-fac/pubs/ messenger

7 Carsten Egedal: Hazard Analysis of the Laser Beam (open air part), ESO Technical Report VLT-TRE-ESO-11850-2435, Issue 1.0, January 2001.

*ESO technical reports may be requested from the authors or from vltarchive@eso.org

Service Mode Scheduling: A Primer for Users

D. SILVA, User Support Group/DMD

Introduction

The execution of observations in Service Mode is an option at many ESO telescopes, especially at the VLT telescopes. In this operations mode, observations are not scheduled for specific nights, they are scheduled flexibly. Each night observations are selected from a pool of possible observations based on Observing Programme Committee (OPC) priority and the current observing conditions. Ideally, the pool of possible observations contains a range of observations that exactly match the real range of conditions and the real number of available hours, so that all observations are completed in a timely manner. Since this ideal case never occurs, constructing the pool of observations must be done carefully, with the goals of maximising scientific return and operational efficiency.

In this article, basic Service Mode scheduling concepts are presented. The goal is to provide users with the information they need to better estimate and perhaps improve the likelihood that their observations will be completed. A specific VLT focus is maintained for most of this article, but the general principles are true for all ESO facilities executing Service Mode runs.

In the Beginning: Proposals, Programmes, and Runs

In general, users submit observing proposals twice a year for Observing Programme Committee (OPC) review. Each proposal describes a scientifically unified **observing programme** which is composed of one or more **observing runs**. A run provides the high-level technical specifications for a set of observations: operations mode (Visitor or Service), targets, telescope, instrument, total execution time, and required observing conditions (e.g., seeing, lunar phase, and transparency).

Pre-OPC: Determining the Available Time

Before each OPC meeting, ESO determines the total **available time**, i.e. how much time will be available for scientific observations. For example, for a normal Period, each VLT telescope will have about 140 nights available for scientific observations. The other 42 nights are used for the ESO Calibration Plan, the Director's Discretionary Time programme, and regular technical maintenance of the instruments and telescopes (e.g. pointing maps, multi-day technical interventions). Some Periods or telescopes have less available time, either due to major technical activity (e.g. instrument commissioning periods) or because the time has been preallocated to Large Programmes. As a guideline, the OPC will allocate up to 30% of available time to Large Prorammes. For any given Period, the time allocated to Large Programmes in previous Periods must be deducted before new time can be allocated.

Over-Subscription and Relative Visitor/Service Mode Demand

Once the available time is determined, the ratio between total requested time and available time (**global over-subscription**) can be calculated. The Paranal global over-subscription ratio is shown in Figure 1 (left axis = Mode Over-subscription) for both Visitor and Service Mode as a function of Period. Over-subscription has been falling steadily over time. Figure 1 also shows the requested time ratio between Service and Visitor Mode (**mode demand**). The demand for Service Mode has been climbing. Note that the allocated mode demand can be larger than the requested mode demand because the OPC may select more Service Mode runs than Visitor Mode runs. But in the end, the scheduled mode ratio is enforced by ESO to be close to 1, i.e. an equal split between Service and Visitor. The issue of Service/Visitor Mode balance is discussed in more detail below.

Users should note, however, that **local over-subscription** (over-subscription as a function of RA) can be much larger, and is typically highest in the RA ranges 0–4 and 10–14. In these ranges, the over-subscription ratio regularly exceeds 5 and has approached 10, especially during dark-time. Demand in these RA ranges is highest not only because they provide access to the prime extragalactic co-ordinate space, but also because they straddle Period boundaries.

The OPC: Scientific Prioritisation and Time Allocation

The main task of the OPC is to produce a scientifically prioritised list of runs and to allocate a total execution time (i.e. integration time plus operations overheads) to each run. Although all the runs within a given programme are usually given the same grade, the OPC does have the option of assigning each run a different grade, or even rejecting individual runs within a single programme. (The OPC is subject of an article in *The Messenger* No. 101, p. 37.) The details of the OPC process are not discussed here; suffice it to say that a lot of time and effort goes into this generally thankless task!

While making these decisions, the OPC does not typically consider technical feasibility (unless a proposed run contains an obvious error) or requested observation mode (Visitor or Service). An OPC grade is based primarily on scientific merit. The OPC also does not generally consider the final distributions in RA or observing conditions. In principle, it is possible for the OPC to allocate all available time to runs requiring excellent seeing and photometric conditions within a narrow RA range. Of course, in practice this extreme case does not occur, but a post-OPC technical and scheduling review is necessary before the final schedule can be constructed.

Post-OPC: Technical Review and Preliminary Long-Term Schedule

Once the OPC review is completed, it is the responsibility of ESO to produce the **Long-Term Schedule (LTS)**, i.e. the list of runs scheduled for a given Period. The goal is to schedule (and execute) all runs above the so-called OPC cut-off line, i.e. the line defined by the available time at each telescope and/or instrument.

This process starts with a technical feasibility review. Each telescope team is given the opportunity to provide technical feedback on runs above the OPC cut-off line.

The technical review evaluates whether or not the technical goals (e.g. signal-to-noise, observation execution concept) of each run are achievable. Technically infeasible runs are rejected, no matter what their scientific priority was. This may seem wasteful – why ask the OPC to review a technically infeasible run? Consider the over-subscription rate: a pre-OPC technical review would take 3–4 times as much effort as a post-OPC review. Furthermore, the number of runs rejected for technical reasons is very small, e.g. approximately 2% per Period at the VLT.

The technical review also evaluates whether or not a run is suitable for Service Mode. Runs which requested Service Mode can be switched to Visitor Mode if the telescope team judges that successful completion of the observations cannot be guaranteed in Service Mode. This decision is usually taken when a run requires a complex, unusual observing strategy and/or a less common or non-standard observing mode. More rarely, runs that requested Service Mode are switched to Visitor Mode, typically to reduce the number of Service Mode runs per Period to a level that ESO can support within available operational resources.

In parallel to this technical review period, the preliminary LTS is constructed. Pre-allocated, continuing Large Programme and newly approved Visitor Mode runs above the OPC cut-off line are assigned specific nights. The remaining available time is assigned to Service Mode. The split between Visitor and Service Mode varies by telescope and Period. For the NTT and 3.6-m, approximately 10% of the available nights are assigned to Service Mode. Approximately 50% of the available time is assigned to Service Mode at the VLT telescopes. Starting with Period 68, at a large fraction of the available time will be assigned to Service Mode at the 2.2-m/WFI, and eventually this may climb to close to 100%.

After the technical review is completed, the preliminary LTS is adjusted to reflect the outcome. Runs are moved from Visitor Mode to Service Mode, or visa versa, as necessary. Technically infeasible runs are removed from the LTS. This revised LTS, particularly the revised list of Service Mode nights, is one of the inputs to the Service Mode LTS construction process.

Building the Service Mode LTS

In the classic Visitor Mode style of operations, users are assigned specific nights. Sometimes these nights are not scheduled exactly when the user wanted. During the actual nights, the observing conditions may not be exactly what was desired. In combination, these two things force the user to adapt their observing programme (and often their science goals) to the actual situation.

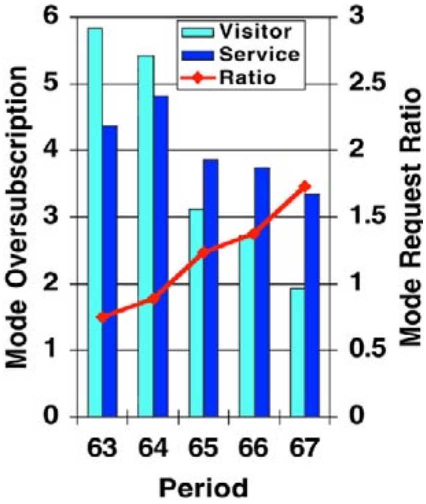

Figure 1: Paranal Global Over-subscription and Mode Demand. *The bars show the oversubscription (left axis, total requested time over available time) across all available instruments. For Periods 63 and 64, only ISAAC and FORS1 were operational. For latter Periods, the instruments were FORS1, FORS2, ISAAC, and UVES. Requested and available time per instrument as made available to the OPC are used. Actual telescope used is ignored. The red line illustrates the mode demand (right-axis, time request ratio between Service and Visitor Mode).*

One of the goals of Service Mode is to execute the observations exactly as described in the approved observing proposal. At the end of the OPC meeting, however, there is no guarantee that this is possible, even for the highest ranked runs. It is necessary therefore to determine if a Service Mode run is executable or not within the context of the actual OPC ranked list of runs and the nights allocated to Service Mode in the LTS.

Basic Principles

Due to statistical fluctuations in observing conditions and down-time, it is highly unlikely that all runs above the OPC cut-off line can be completed. For example, it is known that 15% of available time will be lost to downtime randomly over a long enough time baseline. Initially, any LTS assumes ideal conditions (clear skies, good seeing) but reality is never so kind. Thus, ESO has adopted the following high-level principles:

(1) In general, the scientific objectives of an observing run are not achieved unless all observations are completed.

(2) Therefore, a run should not be scheduled unless it has a high probability of completion.

(3) It is better that a smaller number of runs are totally completed than that all runs are incomplete.

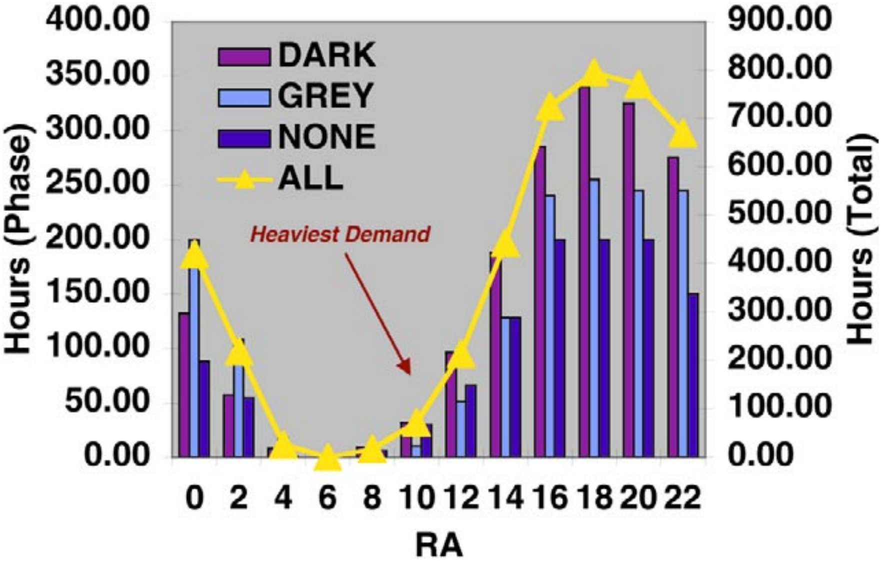

Figure 2: RA/MOON Accessibility Example. *The computed number of hours per RA bin and lunar phase bin are shown for Period 67 and Kueyen (UT2). The shape of this function is driven by the specific nights assigned to Service Mode and the finite length of the Period. Without Period boundaries, each RA bin would be roughly equally accessible over time.*

(4) Observing conditions permitting, runs with higher scientific priority as defined by the OPC should be completed preferentially to lower-priority runs.

These principles are conservative and have been discussed in many forums. Those discussions will not be repeated here.

An important technical principle is that the Service Mode LTS process should manage RA space, not nights. A range of RA is available on any given night. Conversely, any given RA is observable on many nights. For Service Mode, it is more appropriate to manage co-ordinate space than calendar space. This facilitates one of the key advantages of Service Mode: the time-averaged observing conditions for any given target will be better than the conditions on any random night. However, it is also true that this will only be true if a large enough fraction of time is made available to Service Mode operations.

Describing Schedule Parameter Space

The Service Mode LTS review process is driven primarily by principle 2 above. Many parameters determine whether an observing run is likely to be completed or not. *Since some of these parameters are under control of the user, it is possible for the user to fine-tune them at the time of proposal submission to maximise the likelihood that their observations will be completed.*

The most important parameters are requested target distribution and lunar illumination. Within a given sequence of nights, any given point on the sky is observable for a finite number of hours under specific lunar conditions. This **accessibility function** (number of hours available per RA and Dec at a given lunar illumination) is determinate, i.e. it is fixed by the specific sequence of Service Mode nights. An example is shown in Figure 2. Each Service Mode run can be de-composed into specific targets at specific positions with specific total times and lunar conditions. If too many runs want to observe in the same region of the sky under similar lunar conditions (e.g. HDF-S and dark-time), the number of possible hours can be exceeded. This is one example of **local over-subscription**.

To simplify the construction of accessibility function, some assumptions are made. First, RA is binned into 2-hour intervals. Second, for each RA bin, the mean visibility (hours per night above 1.5 air masses) is assumed to be the maximum visibility for that RA that night minus one (1) hour. Objects at different declinations will have different visibilities, but simulations show that this assumption is reasonable. Finally, the lunar illumination distribution is parameterised as dark (moon below horizon or FLI < 0.3), grey (moon above horizon and 0.3 ≤ FLI < 0.6), or bright (moon above horizon and FLI > 0.6).

The next most important factors are the fractional seeing and transparency distributions. Unlike the accessibility function, these distributions are statistical and only valid over long enough time baselines. Each site has a known statistical free-air R-band seeing distribution. For scheduling purposes, it is assumed that in the mean, the R-band delivered image quality distribution in the focal plane follows the R-band free-air distribution, at least for the VLT telescopes. This distribution shifts as a function of wavelength – how this affects scheduling is discussed below. This distribution has been calculated from historical DIMM data in 0.2 arcsec bins averaged over 30 minutes. Analogous to seeing, the statistical site transparency distribution is also available. Here, a conservative assumption is made. The reported photometric fraction (e.g. 78% for Paranal) is split into photometric (PHO) and clear (CLR) bins. It is likely that much of the CLR time is truly PHO, based on available trends in FORS zeropoints. All other usable time is called thin cirrus (THN), i.e. non-photometric. Of course, some time is completely lost to bad weather (e.g. clouds, high humidity), as discussed further below.

The fractional seeing and transparency distributions can be combined into a single **cumulative seeing and transparency distribution**. The Paranal distribution used for Period 67 scheduling is shown in Figure 3. *Users should consider this figure carefully.* Many users regard a seeing of 0.8 and a CLR transparency to be conservative. In fact, this combination of conditions arises only 42% of the time on Paranal. Moreover, while it is exciting to consider that the VLT can deliver 0.4 arcsec images under CLR conditions, Figure 3 shows this happens less than 5% of the time in the R-band. It will happen more frequently at near-IR wavelengths but not by many factors. Finally, users should also keep in mind that these percentages are valid over long-time baselines. On a night-to-night basis, seeing and transparency can vary on short-time scales. Such short-term variability is obvious from the astroclimatology information linked from the Paranal Observatory Web home page.

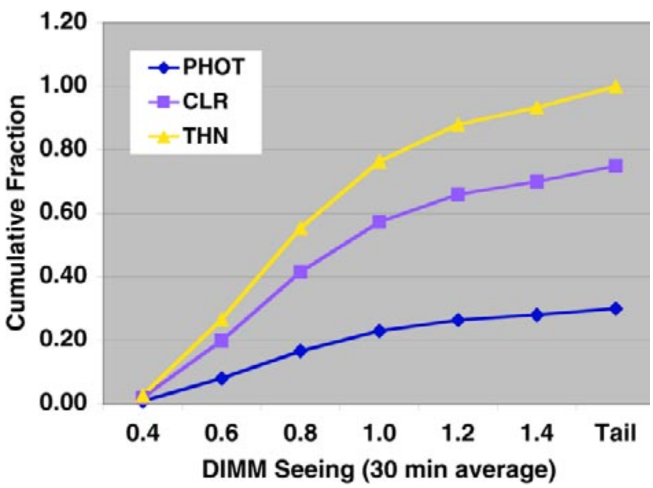

Figure 3: Adopted Paranal SEE/TRANS Cumulative Function. *See text for description.*

The combination of the accessibility function and the cumulative seeing/transparency distribution produces a four-dimensional matrix called the **RA/MOON/SEE/TRANS (RMST) accessibility matrix**. Each element of this matrix is an estimate of cumulative time available in a given RA bin for a specified lunar illumination bin, seeing value, and transparency value. Since this is a cumulative distribution, these are estimates of the number of hours that the seeing and transparency will have these values or better.

The final important factor is downtime, i.e. time schedule for science operations but lost to technical run weather problems. As technical downtime is laudably low (2–5%), total downtime is driven by weather and is a function of month. For Paranal, the total down-time since the start of science operations has been 10–15% per Period. Of course, downtime is statistical.

The Review Algorithm

The review algorithm is best presented in pseudo-code.

```
INPUT
   prioritised list of runs
   precise list of Service Mode nights
      for these Service Mode nights,
      computed:
         total Available Time
         RMST matrix

FOR each newRun

   CREATE newSchedule
   NewSchedule  = currentSchedule
   + newRun
   newRMST = currentRMST +
   newRunRMST

TEST newSchedule
   NewScheduleTotalTime  ≤
   availableTotalTime?
   newRMST ≤ accessibleRMST?

IF both TRUE:
   ACCEPT newRun
   currentSchedule = newSchedule
   currentRMST = newRMST
ELSE
   REJECT newRun
```

Input runs are prioritised in the following order: Guaranteed Time Observations (GTO), Large Programmes (LP), selected Chilean runs (RCH), and Normal runs. High priority Chilean runs are selected in accordance with the principles established in the Chile/ESO operations agreement. There are two additional special cases. Target-of-Opportunity runs are de facto high priority unless the OPC recommends otherwise, in recognition of their time-critical nature. Runs that require a rigid time sequence of observations are also de facto high priority – such observations have to be done on a fixed schedule or it is not possible to achieve the science objective.

If a run has multiple targets, the RMST test must be done for each target or group of targets. If the test is violated for one target, the entire run is rejected.

Although lunar illumination is strictly not cumulative, runs with high priority that can be executed in bright-time are allowed to consume grey and dark-time if necessary.

Likewise, high priority grey-time runs can consume dark-time. This concept is illustrated by the UT1/Antu situation over the last few periods: the OPC has allocated significantly more time to ISAAC runs than to FORS1 runs and consistently given ISAAC runs higher priority. To schedule and execute these runs, ISAAC runs have been allowed to consume dark-time.

So far, downtime has not been explicitly accounted for in this review process. In general, technical downtime has been negligible, except for some early problems with ISAAC. Fractional weather down-time is somewhat dependent on time of year, but this is difficult to model, even in a statistical sense. It has been ignored for now.

Each rejected run is reviewed. Based on this review, a run can be:

Accepted without priority change: run was only marginally in violation of RMST boundary conditions;

Accepted at reduced priority: run significantly violated one run more RMST boundary conditions, but the amount of available time has not been exceeded;

Rejected*:* run significantly violated one or more local boundary conditions and/or total available time has been exhausted.

Because the Service Mode review assumes a specific, preliminary Visitor Mode LTS, it is possible that low priority Visitor Mode runs consume enough time within a specific RA range that a higher priority Service Mode run cannot be scheduled. In this special case, the lower priority Visitor Mode runs may be removed (rejected) from the preliminary Visitor Mode LTS to allow the higher priority Service Mode run to be scheduled, and the Service Mode review is repeated. This is an iterative (and manual) process.

This rejected run review can be illustrated by two real-world examples. The runs used in both examples were highly ranked by the OPC. Table 1 shows the first example. Here, a highly ranked run requested more excellent seeing time than was statistically available. It was automatically rejected. However, this was an ISAAC run specifying observations in the K-band where the delivered image quality distribution is known to be shifted to better seeing. This run was accepted at reduced priority (Rank B), and ultimately more than 70% of the run was completed. Table 2 shows a more complicated example. This run did not request very strenuous conditions. However, the single target was located in a part of RA space demanded by other higher priority runs. This run was rejected. In fact, it has proven difficult to finish the scheduled runs in this RA range due to weather downtime – the rejected run, if accepted at lower priority, would never have been started.

Above the OPC Line: Rank A and B

Once the Service Mode review is completed, each accepted Service

	RA Bin = 2 hrs, Moon = NONE, Trans = CLR						
	Seeing						
	0.4	**0.6**	**0.8**	**1.0**	**1.2**	**1.4**	**Any**
Accessible Hours	3.8	38.4	80.6	109.4	126.7	134.4	144.0
Scheduled Hours (All Runs)	0.0	0.0	0.0	6.0	0.0	0.0	5.0
Requested (Run X)	7.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 1: First Run Review Example. *This run had one target in the RA = 2 hours bin. No lunar restrictions were specified but clear conditions were desired. The requested seeing was 0.4 arcsec. Three vectors are shown. Top: the cumulative accessible hours for this RA, lunar, and transparency bin as a function of seeing. Middle: the scheduled hours for higher priority runs. This row is not cumulative. Bottom: the user requested hours. As described in the text, although the user request is statistically infeasible, the proposed observations were in the K-band, where the delivered image quality is known to be on average better. Furthermore, this run could be executed under any lunar condition. This run was accepted with lower priority.*

	RA Bin = 10 hrs, Moon = DARK, Trans = CLR						
	Seeing						
	0.4	**0.6**	**0.8**	**1.0**	**1.2**	**1.4**	**Any**
Accessible Hours	0.6	6.4	13.4	18.2	21.1	22.4	24.0
Scheduled Hours (All Runs)	0.0	0.0	13.2	4.5	0.0	0.0	0.0
Requested (Run X)	0.0	0.0	0.0	14.0	0.0	0.0	0.0

Table 2: Second Run Review Example. *This run had one target in the RA = 10 hours bin. Dark-time under clear conditions and 1.0 arcsec seeing were requested. Three vectors are shown. Top: the cumulative accessible hours for this RA, lunar, and transparency bin as a function of seeing. Middle: the scheduled hours for higher priority runs. This row is not cumulative. Bottom: the user request. The accessible hours are 18.2. The cumulative scheduled hours are 13.2 + 4.5 = 17.7. Thus, the sum of the Run X requested hours plus the cumulative scheduled hours exceeded the accessible hours. Scheduling this run was made more difficult by the request for dark-time and the knowledge that this RA is often negatively affected by weather downtime. This run was rejected.*

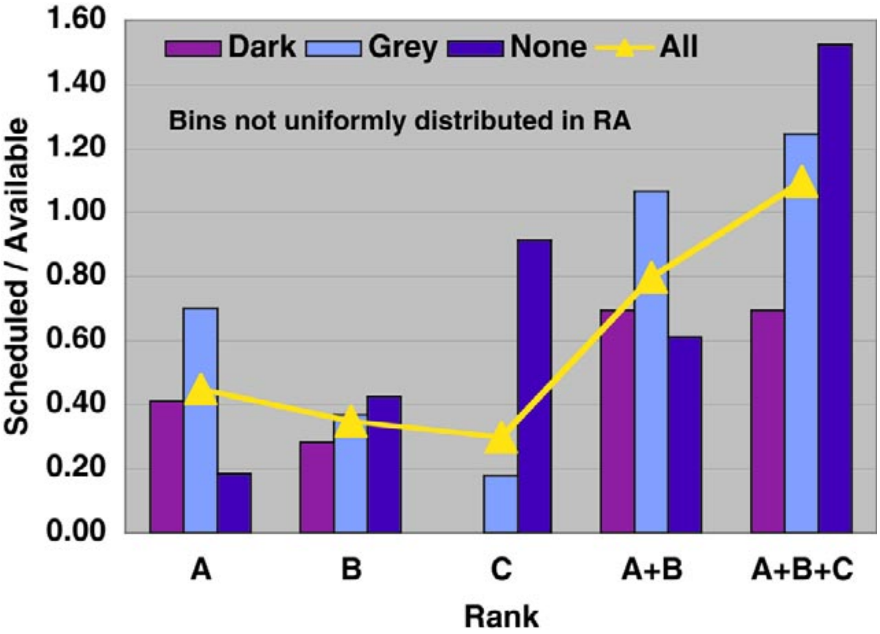

Figure 4: Scheduled vs. Available Time by Rank. *For Period 67 at Kueyen/UT2, the ratio of scheduled vs. available time is given as a function of rank and lunar phase bin. As discussed in the text, scheduled time is split roughly equally between Rank A and B, but the sum of Rank A and B is less than the available time, except for grey-time. This implies that runs that requested grey time received a high enough OPC grade to be scheduled in dark-time. As expected, the filler queue (Rank C) is heavily weighted to runs with no lunar restriction.*

Mode run above the OPC cut-off line can be assigned a priority rank. Such runs are assigned either Rank A ("high priority") or Rank B ("medium priority"). Rank is assigned primarily based on OPC priority. The available Service Mode time is split roughly evenly between Rank A and B. This is illustrated by an example in Figure 4. In principle, this means that statistical fluctuations in down-time or observing conditions will not have a significant impact on Rank A runs. The exception is when a Rank A run has a specific target or time-constraint which is unachievable due to a prolonged period of downtime.

ESO commits to completing Rank A runs whenever possible, even if it takes multiple Periods. ESO does not commit to complete Rank B runs – these have lower scientific priority. If a Rank B run is incomplete at the end of a Period, it is terminated. This strategy ensures that the highest priority runs from each OPC meeting are eventually completed, while preventing too large a fraction of the LTS from being filled with runs carried forward from previous Periods.

The Filler Queue: Rank C

In an ideal world (from the scheduling perspective!), the runs allocated time by the OPC would request targets and observing conditions that span RA and expected observing condition space uniformly. Furthermore, delivered observing conditions would follow their statistical trends and there would be no down-time. In this situation, Rank A and B run completion rate would approach 100%. Reality is not so kind. In particular, the list of Rank A and B runs tends not to include enough runs for the inevitable periods of seeing worse than 1.0 arcsec and/or non-photometric transparency (see Fig. 3). Without LTS modification, unnecessary telescope idle time becomes inevitable.

To deal with this situation, the so-called filler queue (Rank C) is created. Candidate runs for this queue are selected from below the OPC cut-off line but with an OPC grade better than 3. Only runs requesting seeing worse than 1 arcsec and non-photometric conditions (CLR or THN) are selected. Preference is given to runs with no lunar restriction. Runs which have strict timing constraints (i.e. target-of-opportunity projects, time series observations) are not considered. The ideal filler run contains a sample of targets that span RA space but does not require that all targets are observed to produce a sound scientific result. Candidate runs are reviewed by the telescope team for technical suitability. They are also reviewed by the OPC chairman to obtain formal OPC approval. Runs which pass these reviews are inserted into the LTS with Rank C ("low priority"). One realisation of this process is illustrated in Figure 4.

Closing the Loop: Phase 2, Run Execution, and Run Completion

Phase 2

Users awarded Service Mode time are required to submit a Phase 2 package, which includes a description of their observations in the form of Observation Blocks. The scheduling constraints included in these OBs are checked against the observing requirements requested in the original observing proposal, i.e. what was used to build the LTS. Users are not allowed to specify better conditions at Phase 2 than they requested at Phase 1. By enforcing the original requests, it is assured that the executed schedule is in close agreement with the constructed schedule.

However, users are not prevented from requesting more lenient conditions. Some users take advantage of this to relax their Phase 2 scheduling constraints to increase the likelihood that their OBs will be executed. From the ESO (and OPC) perspective, this is an acceptable strategy, as long as this relaxation is not extreme. For example, relaxing the seeing constraint by 0.1 arcsec is not extreme, but relaxing it by 0.5 would be! Such a change would call into question the original justification for the observing run. In truly extreme situations, the OPC would be asked to review the situation.

Some relaxation is actively encouraged. In particular, users who requested photometric (PHO) conditions at Phase 1 are encouraged to submit enough OBs to obtain a proper photometric calibration but to request clear (CLR) conditions for most of their OBs.

Users are also not prevented from redistributing their allocated time between a sub-set of targets. This is necessary in cases where, for example, the OPC only approved a sub-set of targets, time required for operational overheads was underestimated in the original observing proposal, or higher signal-to-noise is desired for a smaller set of targets. From a scheduling perspective, such changes can be problematic. In the worse case, the Phase 1 run had many targets over a range of RA, but the Phase 2 proposal is to use all the allocated time on a single target, creating an unexpected case of over-subscription. During the Phase 2 review, users are contacted when their time re-distribution creates a potential over-subscription situation.

Run Execution

Finally, Phase 2 packages are delivered to the telescope teams for execution. The whole operational process of run/OB management and execution could be the subject of another lengthy *Messenger* article. Only a few key points are mentioned here.

Naturally, OBs are selected for execution primarily by OPC priority, as parameterised by Rank. The next most important criteria is lunar illumination followed by seeing, transparency, and air mass. It is sometimes necessary to override other considerations to execute time-critical OBs (e.g. ToO, time series). OB scheduling is always more complicated in situations where the instrument configuration must be changed manually before the start of the night (e.g. to insert a special filter or a MOS mask). To maintain operational efficiency, it is sometimes necessary to continue executing OBs requiring these manually inserted elements, even if the (improved) conditions would allow the execution of different OBs.

Finally, recall that ESO is trying to complete entire runs, under the principle that the entire run is needed to achieve the desired scientific goal. When selecting OBs for execution, emphasis is placed on completing runs before starting runs. This becomes more important as the Period progresses.

Run Completion Statistics

The fundamental goal of Service Mode is to complete the highest ranked runs first under their requested observing conditions. Is this goal being achieved? Figure 5 illustrates that the answer is "yes".

In Figure 5, the Service Mode completion status for Periods 63–66 are presented. Run Status is explained in the figure caption.

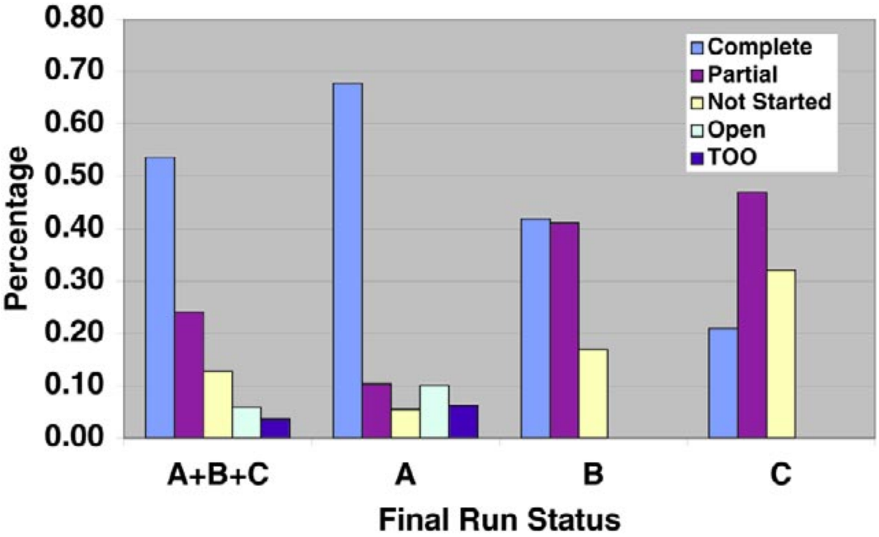

Figure 5: Period 63–66 Run Completion Summary. *Run completion status percentages are given for Period 63–66 VLT Service Mode runs.* ***Completed:*** *all user observations executed with specifications;* ***Partial:*** *run started, not completed;* ***Not Started:*** *run not started;* ***Open:*** *on-going Large Programmes, incomplete Rank A runs;* ***TOO:*** *Target of Opportunity Runs.*

All **Rank A** Open runs will eventually be completed, driving the Rank A Completed fraction to above 80%. It will never be 100% for several reasons. First, Rank A Target of Opportunity runs depend on random events – if the events do not occur, the runs cannot be completed. Second, some Rank A runs turn out to be impossible to complete due to post facto impossible combinations of target, date, instrument configuration, and/or observing conditions. Consider a real situation. Titan observations were requested with 0.4 arcsec seeing on specific dates in February 2001 – since the seeing was never that good on the specified days, the observation was impossible. Finally, there can always be unforeseen technical difficulties. After consultation with the users, such runs are abandoned. Incomplete Rank A runs eventually end in the Partial or Not Started categories.

The **Rank B** situation is more complicated. Although runs in this queue are not guaranteed to be completed, it is perhaps disappointing that the Rank B Completed fraction is only approximately 40%. However within the Partial category there are many runs which are more than 50% completed. Those runs probably produced a scientifically useful data-set as well, but of course that must be evaluated by the users, not by ESO. On the other hand, most runs in the Partial category which are less than 25% complete probably do not produce a scientifically useful data-set and might as well be considered Not Started. This is a hidden scientific inefficiency – it would be far better to produce fewer runs with scientifically useful data-sets than many runs with marginal data-sets. By putting emphasis on run completion, not just OB completion, ESO is trying to avoid the latter outcome.

Rank B completion rate is ultimately limited by the combined technical and weather downtime fraction. To date, Paranal is suffering 10–15% downtime (mostly due to weather) per Period. By design, Rank B runs absorb the impact of this downtime. Since such downtime occurs semi-randomly (some months are statistically worse than others), more than 10–15% of the Rank B runs may be affected.

As expected, the **Rank C** filler runs have the lowest completion rate and the highest Not Started rate. However, the relatively high Partial fraction (40–50%) is consistent with the filler queue concepts discussed above.

Summary: Lessons for Users

When writing observing proposals or preparing Phase 2 packages, users should consider the following key points.

The most critical consideration is a strong observing proposal which results in a high OPC grade. No matter what else is needed or wanted, a high grade increases probability of execution success. Suggestions from the OPC for writing a successful proposal can be found on the ESO Web site.

The local over-subscription is highest in the RA ranges 0–4 and 10–14 hours. If possible, select targets at other RA ranges. It is also recommended to propose specific targets, not a range of targets to be reduced at Phase 2. Finally, targets at widely separated RA (e.g. 2 and 11) should be split into separate runs.

The VLT is capable of delivering truly excellent image quality in the focal plane. Nevertheless, such excellent seeing occurs relatively infrequently (see Fig. 3). Keep in mind that runs with lower priority (Rank B or C) which require better than median seeing are unlikely to be completed and may not even be started. To achieve success with rare conditions, a high OPC grade is necessary.

Consider Figure 3 and the filler queue description carefully. It is much easier to schedule and execute runs which require less stringent conditions (upper right corner). Furthermore, these runs are candidates for the filler queue (Rank C), so they have an increased chance of being scheduled (although not necessarily executed).

Also remember that seeing and transparency varies on short-time scales. This makes it difficult-to-impossible to obtain *continuous* conditions (especially seeing) over many hours within a single night. This is the main reason that ESO requires individual OB execution times to be less than one (1) hour and discourages the submission of tightly linked sequences of OBs.

At telescopes where the fraction of time devoted to Service Mode is low (3.6-m, NTT), it is unrealistic for users to expect excellent observing conditions (e.g. better than median seeing) during Service Mode nights. Service Mode proposals for these telescopes should plan accordingly.

Users are reminded to read the Phase 2 instructions and the specific User's Manuals carefully when preparing their Phase 2 packages. These documents provide more hints and suggestions about maximising the success of your run.

Acknowledgements

The ESO scheduling process is a delicate balance of scientific and technical considerations, as well as human-relations management. The ESO community has benefited immensely in this area from the dedication and careful work of Jacques Breysacher over the last 25 years.

Scheduling concepts related specifically to Service Mode have been discussed within ESO for many years. Roberto Gilmozzi, Gautier Mathys, Palle Møller and, of course, Jacques Breysacher have made significant contributions to the refinement of these discussions into the current operational concepts.

Interesting and probing discussions with the ESO Observing Programme and User's Committees are gratefully acknowledged.

The actual operational execution of these concepts would not be possible without the hard work and dedication of the members of the Visiting Astronomer Section (VISAS), the User Support Group (USG), and Paranal Science Operations (PSO).

Systems and tools that support these activities have been provided by the User Support System (USS) team.

Jacques Breysacher, Fernando Comerón, Peter Jakobsen, Gautier Mathys, Francesca Primas, and Martino Romaniello provided thoughtful and thorough feedback on earlier drafts.

Peter Quinn provided the original version of Figure 1.

Hunting the Southern Skies with SIMBA

(Taken from ESO Press Release 20/01 – 30 August 2001)

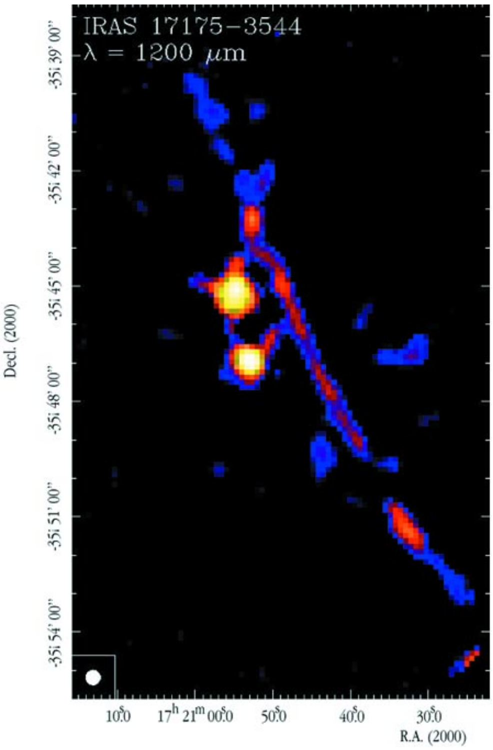

Figure 1: This intensity-coded, false-colour SIMBA image is centred on the infrared source IRAS 17175-3544 and covers the well-known high-mass star formation complex NGC 6334, at a distance of 5500 light-years. The southern bright source is an ultra-compact region of ionised hydrogen ("HII region") created by a star or several stars already formed. The northern bright source has not yet developed an HII region and may be a star or a cluster of stars that are presently forming. A remarkable, narrow, linear dust filament extends over the image; it was known to exist before, but the SIMBA image now shows it to a much larger extent and much more clearly.

A new instrument, SIMBA ("SEST IMaging Bolometer Array"), was installed at the Swedish-ESO Submillimetre Telescope (SEST) at the ESO La Silla Observatory in July 2001. In order to achieve the best possible sensitivity, SIMBA is cooled to only 0.3 deg above the absolute zero on the temperature scale.

The SIMBA ("Lion" in Swahili) instrument detects radiation at a wavelength of 1.2 mm. It has 37 "horns" and acts like a camera with 37 picture elements (pixels). By changing the pointing direction of the telescope, relatively large sky fields can be imaged.

SIMBA was built and installed at the SEST within an international collaboration between the University of Bochum and the Max Planck Institute for Radio Astronomy in Germany, the Swedish National Facility for Radio Astronomy and ESO.

SIMBA is the first imaging millimetre instrument in the southern hemisphere. Radiation at this wavelength is mostly emitted from cold dust and ionised gas in a variety of objects in the Universe. Among others, SIMBA now opens exciting prospects for in-depth studies of the "hidden" sites of star formation, deep inside dense interstellar nebulae. While such clouds are impenetrable to optical light, they are transparent to millimetre radiation and SIMBA can therefore observe the associated phenomena, in particular the dust around nascent stars.

This sophisticated instrument can also search for disks of cold dust around nearby stars in which planets are being formed or which may be leftovers of this basic process. Equally important, SIMBA may observe extremely distant galaxies in the early universe, recording them while they were still in the formation stage.

During the first observations, SIMBA was used to study the gas and dust content of star-forming regions in our own Milky Way Galaxy, as well as in the Magellanic Clouds and more distant galaxies. It was also used to record emission from planetary nebulae, clouds of matter ejected by dying stars. Moreover, attempts were made to detect distant galaxies and quasars radiating at mm-wavelengths and located in two well-studied sky fields, the "Hubble Deep Field South" and the "Chandra Deep Field" .

Various SIMBA images have been obtained during the first tests of the new instrument. The first observations confirm the great promise for unique astronomical studies of the southern sky in the millimetre wavelength region.

These results also pave the way towards the Atacama Large Millimetre Array (ALMA), the giant, joint research project that is now under study in Europe, the USA and Japan.

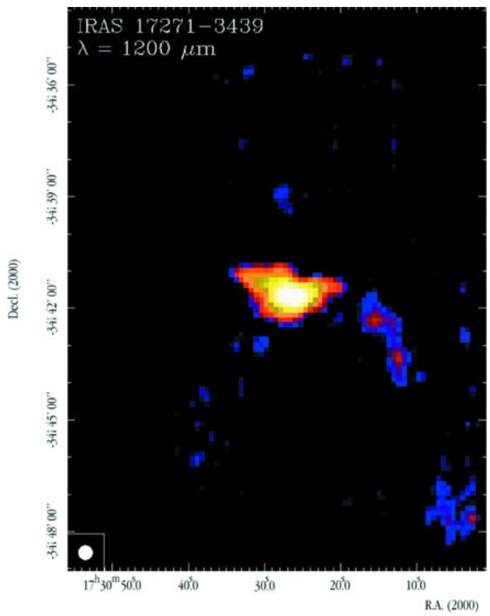

Figure 2: This SIMBA image is centred on IRAS 17271-3439 and includes an extended bright source that is associated with several compact HII regions as well as a cluster of weaker sources.