Fifteen Years of Service Mode Operations: Closing the Loop with the Community

Francesca Primas[1]
Lowell Tacconi-Garman[1]
Stéphane Marteau[1]
Vincenzo Mainieri[1]
Marina Rejkuba[1]
Sangeeta Mysore[1]
Christophe Dumas[1]
Andreas Kaufer[1]
Ferdinando Patat[1]
Michael Sterzik[1]

[1] ESO

The first Service Mode (SM) observations with the VLT were made by ISAAC in April 1999. Since then, new instruments have become operational and first generation ones replaced, filling the 12 VLT foci and feeding the VLT Interferometer and its four Auxiliary Telescopes. Efficiently operating such a broad range of instruments, installed and available every night of each year, on four 8-metre telescopes offers many challenges. Although it may appear that little has changed since 1999, the underlying VLT operational model has evolved in order to accommodate different requirements from the user community and features of new instruments. As ESO and its Member States approach routine operations with ALMA, and at the same time prepare for the next challenge, the construction of the E-ELT, it seems timely to take a closer look at what SM has brought to the scientific arena, both in terms of science data and support. Did it fulfil its original goal, if so, how well, and what are the lessons learned? A careful analysis of statistics and trends in Phase 1 and Phase 2 are now being conducted in the DOME (Dashboard for Operational Metrics at ESO) project. We summarise the main findings, concentrating on the handling of Service Mode.

Introduction

When ESO's Very Large Telescope (VLT) opened its first dome at the beginning of Period 63 (P63; c.f., Table 1 for a summary of the milestones), it was only the second observatory in the world offering 8-metre-class telescope(s). It was also the first ground-based facility offered to the scientific community at large with both classical (also known as Visitor Mode [VM]) and queue observing. The latter, known as Service Mode observing in ESO terminology, was considered to be the most promising way to ensure the observing flexibility necessary to execute the most demanding scientific programmes under the required, usually very well-defined, conditions.

The official reference to the ESO VLT/Interferometer (VLT/I) operational model is the VLT/VLTI Science Operations document[1] approved by ESO Council (most recent update from 2004). The document touches upon all main areas of the ESO data flow system, from proposal submission, to execution and archiving of the data. It sets guidelines for the SM/VM ratio (at least 50 % in SM in order to achieve an optimal scientific return), describes what types of observing times exist (normal/open time, fraction of time reserved for Large Programmes, Guaranteed Time Observations [GTO], Director's Discretionary Time [DDT], Target of Opportunity [ToO]) and the Service Mode rank classes of the scheduled programmes (A, B and C) as a way to implement the scientific ranking and evaluations delivered by the Observing Programmes Committee [OPC]). Within the current framework, A class runs are the highest priority programmes for which ESO has committed to make all possible efforts to complete (within the requested observing Period or by granting them carryover status to the next useful visibility Period); B and C class runs however are executed only within the requested observing semester. The document mentions the set-up and maintenance of calibration plans, both for scientific and instrument health-checking purposes, as well as how the obtained data will then be distributed to the Principal Investigators (PIs) and later to the community at large. More details can be found in a number of SPIE and *Messenger* articles (e.g., Quinn et al., 2000; Silva et al., 2001; Comerón et al., 2003).

Table 1. Summary of VLT/I milestones.

Period (Start)	Milestone
P63 (April 1999)	VLT UT1 starts operations
P64 (October 1999)	VLT UT2 starts operations
P67 (April 2001)	VLT UT3 and UT4 start operations
P72 (October 2003)	VLTI starts operations

The "front-end" of the ESO data flow system begins when the period-based Calls for Proposals is issued and members of the user community submit their observing proposals. This, together with the scientific evaluation and ranking of the proposals by the OPC, defines the Phase 1 process, which formally ends with the delivery of the schedule for all available telescopes (c.f., Patat & Hussain [2013] for a review of this process). As soon as the time allocations are announced to the community, the Phase 2 process can start, i.e., the preparation of the SM observations by the successful PIs (or their delegates), followed by the verification and optimisation of each programme's observing strategy by the User Support Department (USD). Observing queues are then prepared per instrument and made available to the Observatory.

Night-time observations are carried out by the staff at the Observatory, following a complex ranking optimisation which takes into account the scientific priority set by the OPC, as well as the observability of the programme and the relative priorities within each programme set by the users (c.f., Bierwirth et al. [2010] for a description). If problems arise, the affected parties are informed via a ticketing system (directed to engineers if the problem is of a technical/instrumental nature, support astronomers when it concerns the science). As data are taken, these are transferred to ESO Garching and become available to the PI (and his/her delegate) after they reach the Science Archive. Smooth operations are guaranteed by real-time health checks on the instruments and careful monitoring of their performance (Hanuschik & Silva, 2002).

Despite being new to the SM concept of carrying out astronomical observations (except for the short period during which Service Mode was deployed at the New Technology Telescope [NTT] for trial purposes), the ESO community clearly adapted very quickly to this new mode. Figure 1 shows the relative demand (upper) and time allocations for VM and SM since the start of VLT operations in Period 63. The requested SM fraction consistently exceeds 50 %.

Figure 2 shows the global (i.e., SM and VM combined) oversubscription rates for

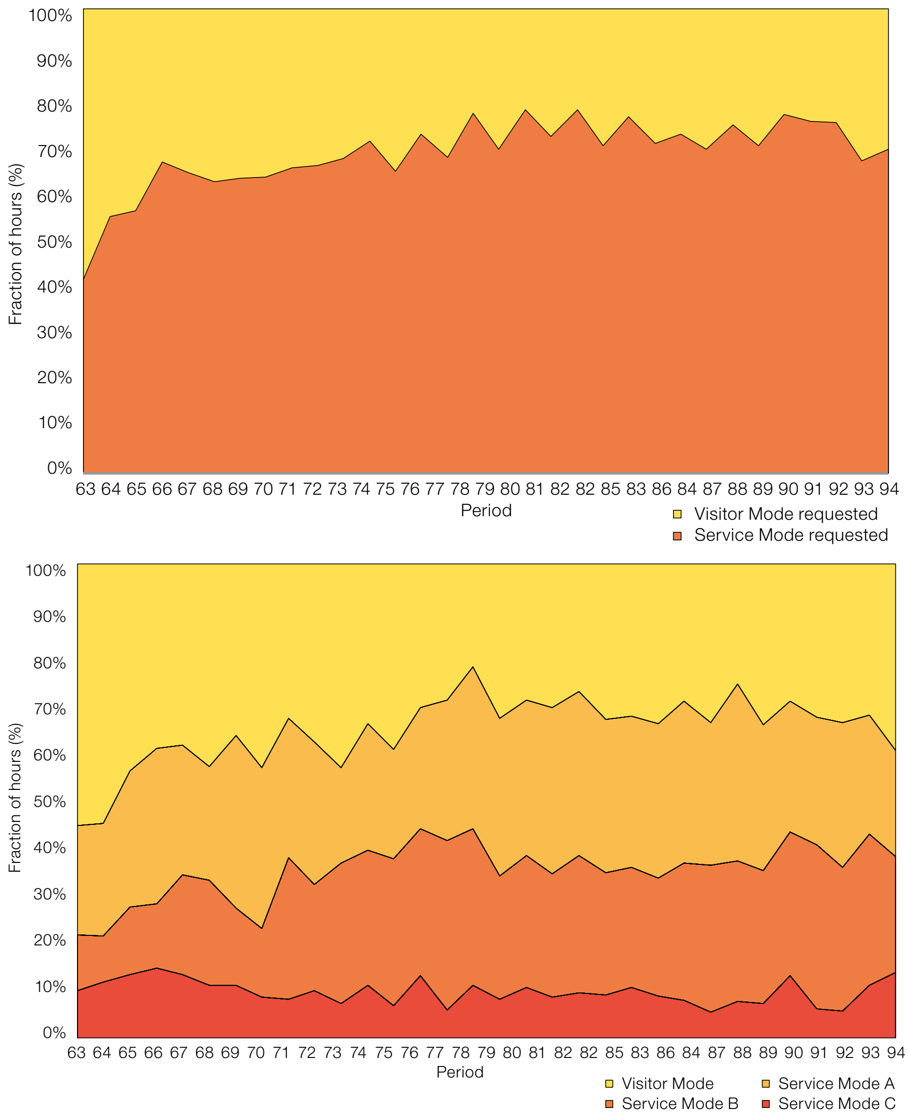

Figure 1. Upper: Relative percentage of the Service *vs.* Visitor Mode hours requested on all VLT/VLTI telescopes for each observing Period. Lower: Percentage of time allocated on all telescopes of the VLT/VLTI over the same time span. The Service Mode area is further split into rank class A, B and C allocations.

each Unit Telescope (UT) and for the VLT Interferometer. These rates are obtained by normalising the time requests to the time allocations. However, since we oversubscribe the time allocations in SM by ~ 10 % in order to ensure short-term scheduling flexibility, the true oversubscription rates would actually end up slightly higher. Moreover, if split between SM and VM, oversubscription rates tend to be slightly larger for Service Mode.

1999–2014: A steady evolution

Despite the fact that the backbone infrastructure of the VLT operational model has been in place since the start of VLT operations (Quinn et al., 2000), operational complexity has increased as new instruments, new operational and instrumental enhancements and new programme types were added. The request for SM (and its consequent allocation) has increased and is now stable at around 70 %, as is evident from Figure 1. The time invested in the Phase 2 verification process of SM observations has proved to be robust, as demonstrated by the very small number of problem reports related to the adequacy of the transmitted Phase 2 material.

The highest level change made to the system was probably the introduction of the User Portal, in November 2007. The goal of this effort was to provide the user community with a system in which account information (username, password, contact information) for all science- and observation-related web-based applications, and standalone software, is unified and can be controlled by the user. This change represented a big challenge, as it required the merging of a number of operational databases into a new one; but the end result makes the user experience with ESO web applications and

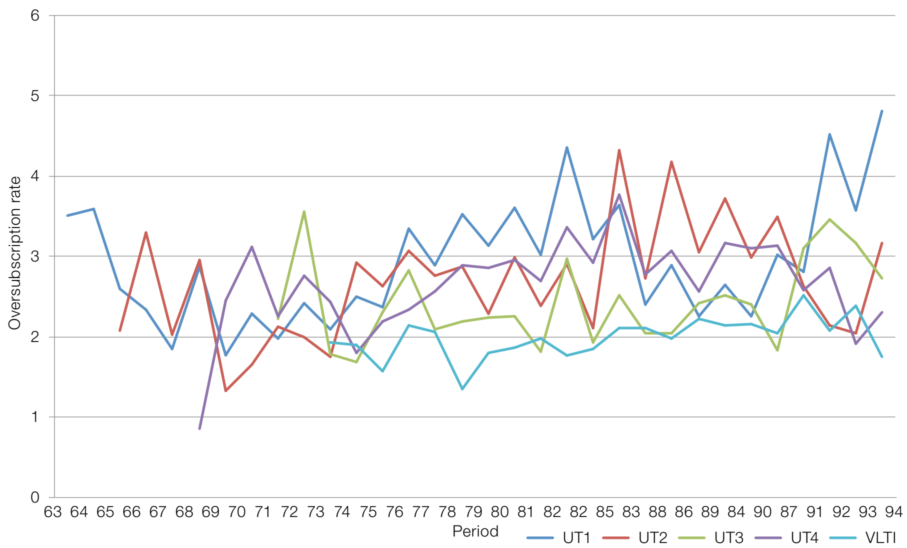

Figure 2. Overall (i.e., combining both SM and VM) oversubscription rates by Unit Telescope (UT) and for the VLTI are shown.

other software simpler and more manageable.

From a more observational point of view, the importance of fast follow-up of unexpected astronomical events was fully recognised. These science cases were implemented in the form of the Rapid Response Mode programme type. More recently, longer-term Monitoring of sources (spanning multiple semesters) was enabled through the monitoring programme type. Calibration proposals were also introduced, allowing the design of programmes that deliver products in support of a variety of science goals for the benefit of the entire user community.

The option to request changes with respect to the observations originally proposed in Phase 1 (as far as targets and instrument set-ups are concerned) was better formalised via a web interface to allow for proper tracking and protection of pre-approved targets. More recently, the upgrade of the Night Log Tool allowed users the option to subscribe to a tailored distribution of observing logs for their own programmes. On the operational side, the upgrade of the helpdesk ticketing system gathered all (operations-related) stakeholders under the same system, making the exchange of information easier, faster and more traceable.

Last, but not least, among the changes, the addition of the survey telescopes, the VLT Infrared Survey Telescope for Astronomy (VISTA) and the VLT Survey Telescope (VST), triggered a deep-seated review of the observing tools. The new operational dimension and challenge of carrying out long-term (five-year), massive surveys for which the fulfilment of the observing strategies, although simple, needed to be automated, demanded this change. A new version of the Phase 2 Proposal Preparation tool (known to users as P2PP) and the Observing Tool resulted, introducing the concept of scheduling containers of Observation Blocks (OBs), to better follow time links, concatenations of observations, etc., and a more robust ranking algorithm. The upgrade of the tools was also used to introduce the figure of the delegate into the system. The delegate is the person entrusted by the official programme PI to follow different (or all) phases of the project (Phase 2 and the execution phase and/or data access via the archive interface).

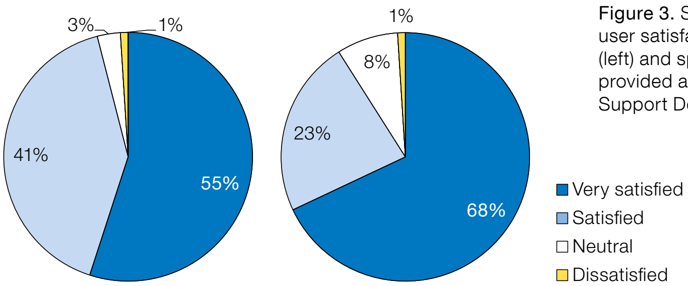

Figure 3. Summary of statistics for user satisfaction with overall support (left) and specific support/advice provided at Phase 2 (right) by the User Support Department.

Overall, the community seems to be appreciative of our system and of the improvements. Feedback from the community is sought regularly via dedicated feedback campaigns and is also received via the Users Committee. Figure 3 summarises the most recent feedback received in March 2014, based on a 31 % response rate (out of 502 targeted PIs)[2].

Closing the loop with the community

It is challenging to extract robust metrics from the operational databases, mostly because when the infrastructure was conceived and put in place (almost 20 years ago) operational statistics were obviously not considered as top priority. At that time ESO was operating its only facility, the La Silla Observatory, in Visitor Mode. The re-design and implementation of a completely new data flow system that combined both VM and SM observations was a major challenge in itself, without adding the extra complexity of by-products such as robust operational metrics. Hence, as will be mentioned in the examples below, even the most seemingly obvious number that could be retrieved from a database may be affected by subtle caveats, which we will describe.

This is not to imply that, until now, the operational efficiencies of the Paranal Observatory were never checked, but simply that a limited number of (more technical) benchmarks were used, such as shutter-open time, technical downtime, available science time, etc. All these numbers point to successful management of the Observatory in terms of "time made available to the Community for science observations". Nonetheless, with the occupation of all the VLT foci plus the VLTI, efficiency metrics regarding use of science time were not regularly collected, on account of the much broadened parameter space. At a minimum, having to deal with a very large number of observing programmes, use of three different instruments at the same Unit Telescope, divided into three priority classes (A, B and C) and with an additional division into different programme types, is challenging. At the same time there was an increasing number of requests for a variety of operational statistics used for reporting to various ESO committees. It thus became clear that there was a strong need to determine in a more systematic way how the time available for science observations is used, by identifying a set of robust benchmarks. Therefore the User Support Department has conceived and implemented a system that enables the production of reproducible, reliable queries and their graphical representation, internally called DOME (Dashboard for Operational Metrics at ESO; Primas et al., 2012).

Firstly we tackled the metrics at the front-end of our operational model, Phase 1, but limited consideration to the more operational metrics (pressure, requested *vs.* allocated time, number of unique PIs, etc). A few examples have already been shown (c.f., Figures 1 and 2). More statistics derived from Phase 1 (such as instrument demand, population of constraint space, statistics on users and teams) will be published in a series of dedicated papers.

Granularity is another important aspect of this exercise: most of the metrics can be extracted and displayed per telescope, per instrument, per rank class, per type of programme, etc. For the specific case of Phase 1 metrics, this often comes from simple database queries.

This ease of extraction does not always apply owing to the varied ways in which some information is encapsulated in the database(s). For instance, identifying adaptive optics (AO) assisted observations is complicated by the fact that this information is stored in the databases in different ways for the different instruments with AO capability. Nevertheless, no matter which criteria are used, Service Mode is clearly in high demand, as is evident from Figure 1.

The next, necessarily more complex, area to investigate is Phase 2, i.e., the execution of Service Mode programmes. Here, the two main questions are:

1. How many runs and/or how much scheduled observation time is completed in each rank class?
2. What levels of completion do terminated runs reach?

For A class runs, closer scrutiny of the first question reveals further details about the completion fraction within the first allocated Period, the amount of time needed to complete A class runs and the fraction of A class observations that is carried over into subsequent Periods.

In order to be able to retrieve the corresponding metrics reliably, it was decided to restrict the exercise to VLT/I only. Data from the VST, VISTA and La Silla are thus excluded, mainly because of their current usage (the VST and VISTA are basically survey telescopes; with the exception of a very few ToO cases [see below] La Silla has offered only Visitor Mode since Period 83). Within the VLT/I runs, we looked into the following types of programmes: normal, GTO, DDT, monitoring and calibration. We thus exclude ToO (and Rapid Response Mode) runs and Large Programmes. Exclusion of the former is justified since their completion rates strongly depend on availability of a suitable target for triggered/follow-up observations. This feature distinguishes ToO runs from all other runs, the execution of which depends on the realisation of the requested external conditions, combined with ESO's handling of the observation.

Exclusion of Large Programmes from early consideration is justified since they are a different "species": their completion timescales depend on a variety of factors, some of which are not under ESO's control (e.g., coordinated observations with other facilities and/or dependence on space-based observations). These aspects may affect the progress of a given Large Programme, introducing unavoidable delays that need to be carefully considered when investigating operational efficiencies. In turn, this aspect demands more complex and tailored database queries that can take account of all these different effects.

As is already apparent from the wording of the above two questions, the statistics can be examined in terms of runs (maybe more interesting for the users) and in terms of hours (requested/allocated/executed, possibly more interesting for the Observatory). In the following, we provide both views for some metrics, usually as a function of time (in the form of observing Period, on the *x*-axis).

The primary question is of course about completion fractions. Figure 4 represents the overall completion rates of A, B and C class runs/hours over more than seven years of VLT/I operations. These statistics are provided for Periods P78–P93 because this is the time over which we can guarantee homogeneous treatment of the data. The reason has to do with the insertion of new identifiers (like the run status in P78) so that database queries can extract the same information about these Periods; depending on the type of metrics, other graphs may cover slightly different time ranges. We note however that for the two most recent Periods, P92 and P93, overall completion

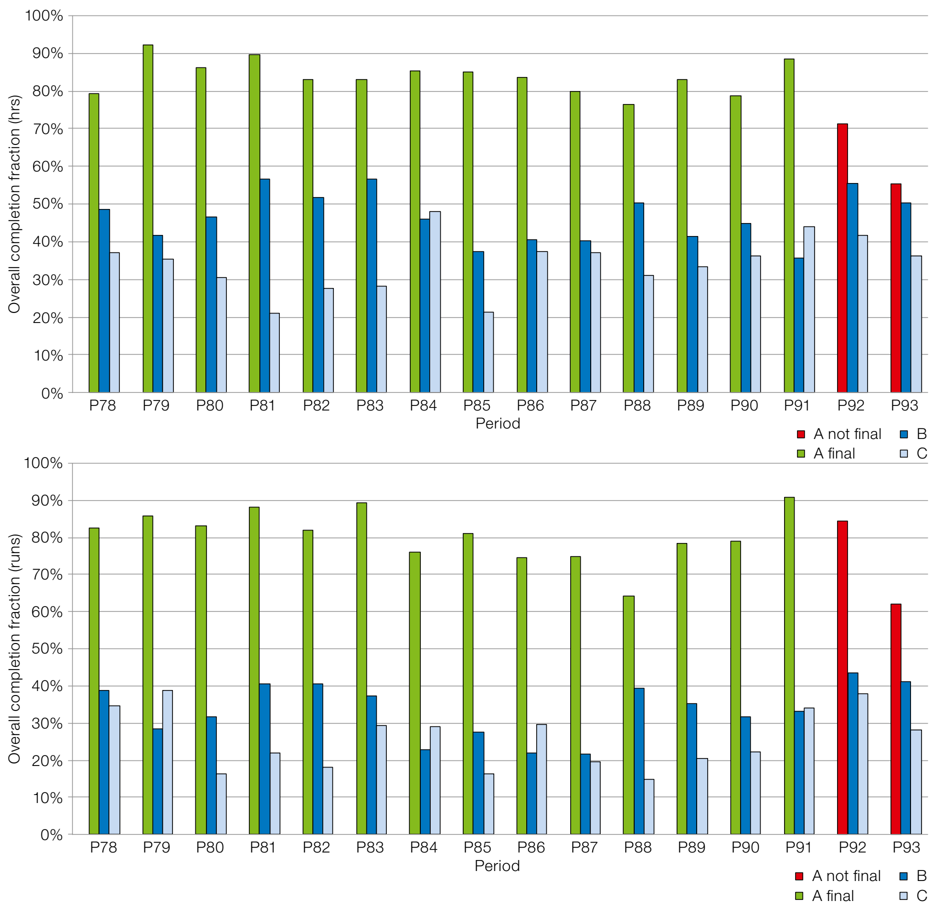

Figure 4. Overall fraction of completed programmes in terms of hours with respect to the total time that was allocated to A, B and C class runs (upper) and the fraction of completed programmes in terms of runs with respect to the total number of runs that were scheduled in A, B and C classes (lower). For the A class category, the fraction completed is irrespective of when the corresponding hours/runs were completed, i.e., hours/runs executed as carryover are included.

statistics are available only for B and C class runs, because some A class runs are still being carried over to future Periods and are not complete. This explains why the corresponding A class bars for P92 and P93 are coloured differently, and marked as “A not final”.

Over the past 7–8 years, on average, ~ 80 % of all A class runs (and time) could be completed, whereas for B and C class runs this completion rate decreases to 45 % and 35 % respectively; in terms of time, 34 % and 25 % of B and C class proposals could be completed, respectively. Of the ~ 20 % of A class runs terminated, approximately half reached a ≥ 50 % completion fraction. There seems to be no correlation between the terminated A class runs and the requested conditions (i.e., the terminated runs are not among the most demanding ones in terms of seeing/transparency). We note that among the terminated A class runs, there are some approved Director’s Discretionary Time programmes, that sometimes have higher chances of not being completed on account of the more challenging/pilot study nature of such programmes. In addition, some time-critical programmes, with one or very few visibility windows, are included among the terminated A class runs.

Figure 5 provides one extra dimension, i.e., the time needed to complete A class runs. It is evident that the red sections (the fraction of runs completed within the first period, defined as the Period in which the runs are first scheduled) dominate the height of the bars, followed by the P + 1 section (i.e., what is approved as carryover and remains visible also in the following semester, shown in green).

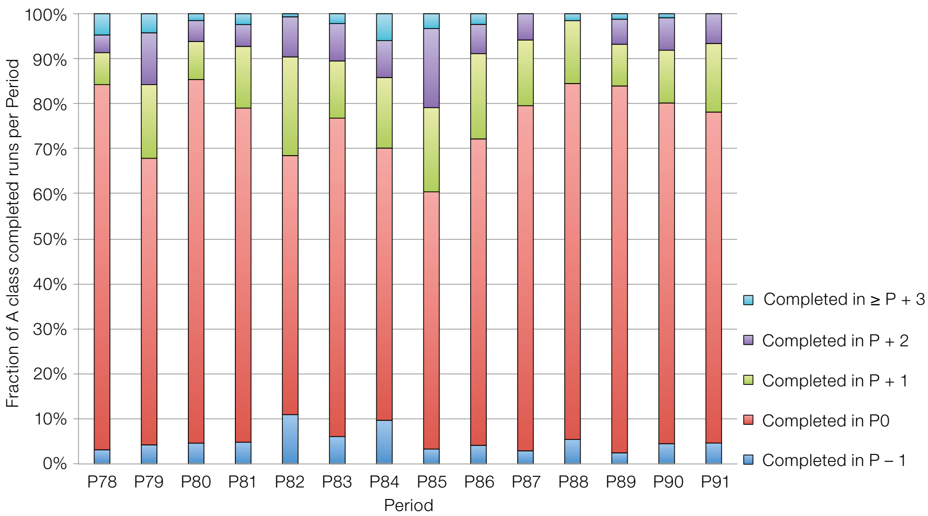

Figure 5. The timeliness of completion of A class runs is shown. This plot refers only to those A class runs that are indeed completed, i.e. the 100 % mark on the *y*-axis corresponds to the ~ 80 % completion fraction that has emerged from Figure 4. Legend: P0 (red) identifies the first Period in which a given run is scheduled; P – 1 is the preceding Period, during the last months of which some P0 runs may be executed as carryunders; P + 1, P + 2, P + 3, etc. refer to the first, second and third Periods following P0.

The most disappointing finding of this study is the rather low completion rate of B class observations. Irrespective of whether their completion is examined in terms of number of runs or number of hours, the picture remains unsatisfactory. On average, only ~ 30 % of the hours allocated to B class runs and ~ 45 % of the runs are completed. The difference between the completion rates in [hours] *vs.* [number of runs] indicates that we tend to preferentially complete shorter runs. It is also important to note that B class runs are as demanding – in terms of requested constraints — as A class runs, thus they end up in direct conflict with the higher priority category. Similarly, the distribution of the time allocations per B class run does not differ significantly from that of A class runs.

B class observations thus pay the highest penalty in the current system. This not only causes frustration to individual PIs, but also represents a concern for ESO because currently half of the available SM time is scheduled in the B rank class. For both reasons it is imperative to take a closer look at this rank class. For instance, we can look into the completion fractions of terminated B class runs (c.f., Figure 6).

From Figures 4 (lower panel) and 6 together, we find that, on average, 45 % of all B class runs are completed, 45 % are started, but terminated at the end of the Period, and 10 % do not even get

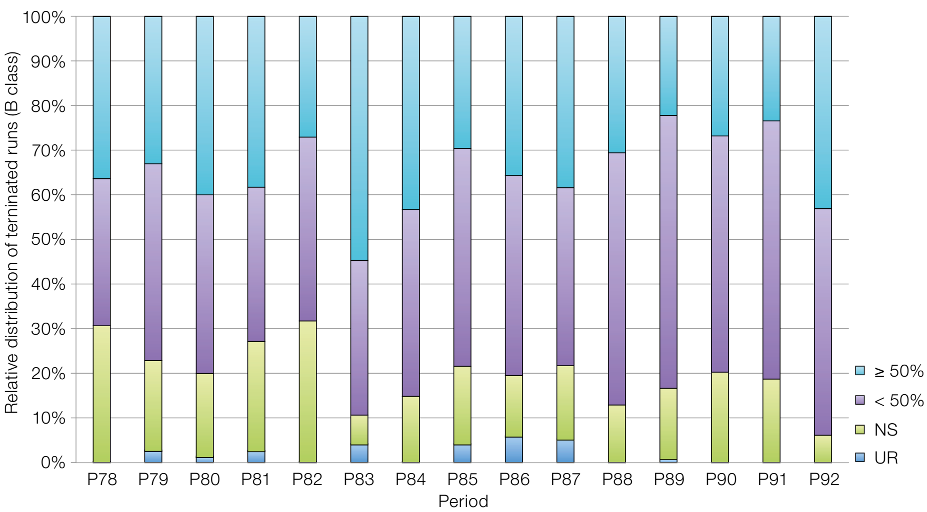

Figure 6. Relative distribution of all terminated B class runs, based on their completion fraction. Legend: NS (runs that were Not Started); ≥ 50 % or < 50 % refers to the completion fraction at the end of their validity Period. Runs terminated by user request are indicated as “UR”.

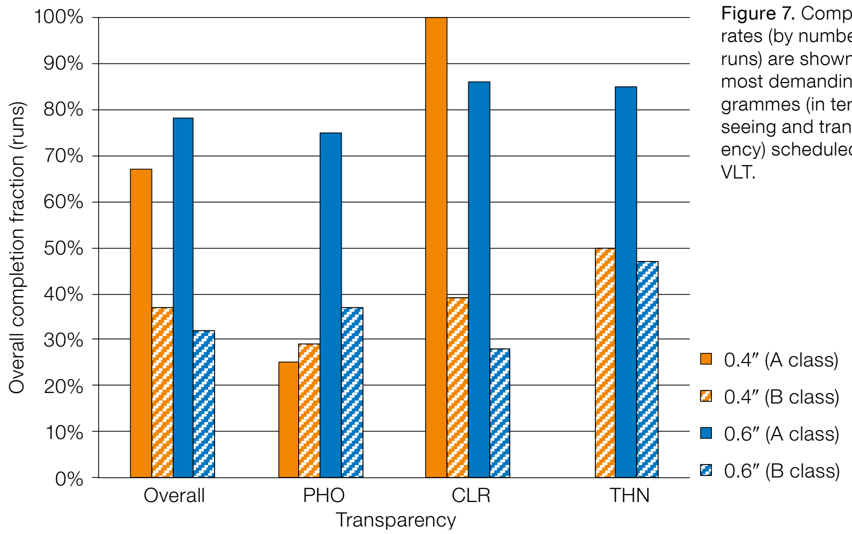

Figure 7. Completion rates (by number of runs) are shown for the most demanding programmes (in terms of seeing and transparency) scheduled at the VLT.

started. From Figure 6 alone, we determine that of the started but terminated runs, 40 % are completed at the level of ≥ 50 %, i.e., 18 % of all B class runs are terminated with completion fractions ≥ 50 %, whereas 27 % are terminated with a completion fraction < 50 %.

Service Mode: The door to very demanding conditions

One of the proclaimed goals for the implementation of a mixed (Visitor + Service Mode) operational model from the very start of VLT/I operations was to enable the successful execution of those programmes requiring very demanding observing conditions. While the definition of demanding observing conditions may include observations of targets widely spread in position, and thus not easy to observe in the few nights of a "classical" VM run, or the need for AO-friendly atmospheric conditions, here we take the seeing constraint as the demanding condition.

By looking at all SM runs that have requested very good seeing conditions (0.4 and 0.6 arcseconds) since P74 (noting, however, that the demand for these conditions has decreased significantly in recent years), we have traced their completion fractions as a whole, according to the requested atmospheric transparency (photometric [PHO], clear [CLR] or thin [THN]) and to their rank class (for this, only A and B rank classes are considered). Figure 7 summarises the results.

Statistics for 0.4-arcsecond seeing runs are very small (37 runs in total, from P74 onwards). After splitting the runs into rank class or transparency bins, the numbers are even smaller. Overall, one finds that approximately half of these runs have been completed, but it should be noted that the large majority of them (27 out of 37) were scheduled in B class (and one even in C class, thus impacting the chance of success from the very start). When runs are split between A and B rank classes, 67 % completion rates for the A class runs and 37 % for the B class runs are found. Within the A class, 25 % of those requesting 0.4 arcseconds + PHO and 100 % of those requesting 0.4 arcseconds + CLR conditions were completed. For the B class runs, the corresponding completion fractions are 29 % and 39 % respectively. Drawing firmer conclusions (e.g., why, under photometric conditions, class B runs achieved slightly higher completions rates than A class runs) requires a closer look at more parameters and/or correlating one parameter against another, e.g., the time of the year for the requested observations and the corresponding weather statistics.

Statistics for 0.6-arcsecond seeing runs are more robust with 559 runs in total. The overall completion fraction does not change much (now at 54 %), but some rates increase when grouping the runs according to their rank class or requested transparency. For A class runs, we reached a completion fraction of 78 %, whereas for the B class runs we remain at around 32 %. Within the A class, we find very satisfactory completion rates of 75 %, 86 % and 85 % for runs requesting PHO/CLR/THN conditions, respectively. For B class runs, as is to be expected, these same rates decrease to 37 %, 28 % and 47 % respectively.

Overall, we find these numbers reassuring because they match the global completion rates for all A and B class runs (c.f., Figure 4 and associated text). In order to determine whether Service Mode has indeed achieved one of its promised targets in completing very demanding programmes by implementing flexible scheduling, one would ideally need to compare them to the success rate of these same runs if they had been scheduled in Visitor Mode. This would have to be simulated, based on the weather statistics available for the past 15 years.

Figure 8 shows the probability of realisation of given sky transparency conditions on Paranal for a scale of seeing values, based on monthly weather loss, sky transparency and seeing statistics compiled between 2008 and 2012. Here we have assumed that the underlying seeing cumulative distribution applies equally to all months. From this figure, the probability of achieving photometrically clear conditions with a particular seeing value can be read off, e.g., for 0.6 arcseconds this is ~ 20–30 % in March and November. It is reassuring to see that these values are low, when compared to the completion rates obtained in SM (~ 50 % overall, ~ 67–78 % for A class runs only). This provides strong support for the ESO mixed SM–VM model, but admittedly our SM completion rates span more than one Period for the highest priority runs.

We note that among the top 20 most-cited scientific articles based on ESO data, 30 % of their respective SM runs requested and received photometric conditions. Among these same 20 articles, of those that involve a component of Service Mode, there are at least six SM runs that required 0.6-arcsecond seeing, ten that required 0.8 arcseconds and another six required 1-arcsecond seeing.

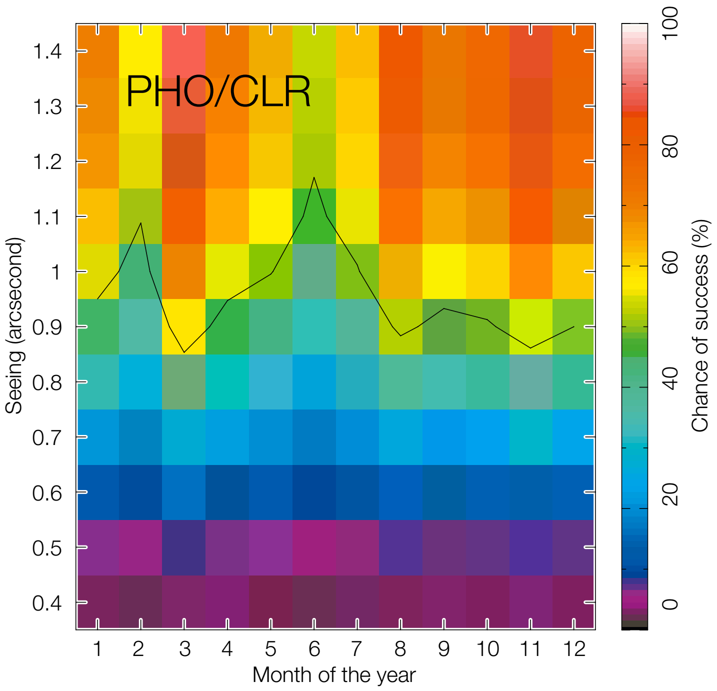

Figure 8. Percentage chance of success in the realisation of photometric/clear (PHO/CLR) conditions and different (Shack–Hartmann) seeing values, across the year at Paranal. The black line through the middle represents the 50 % chance of success.

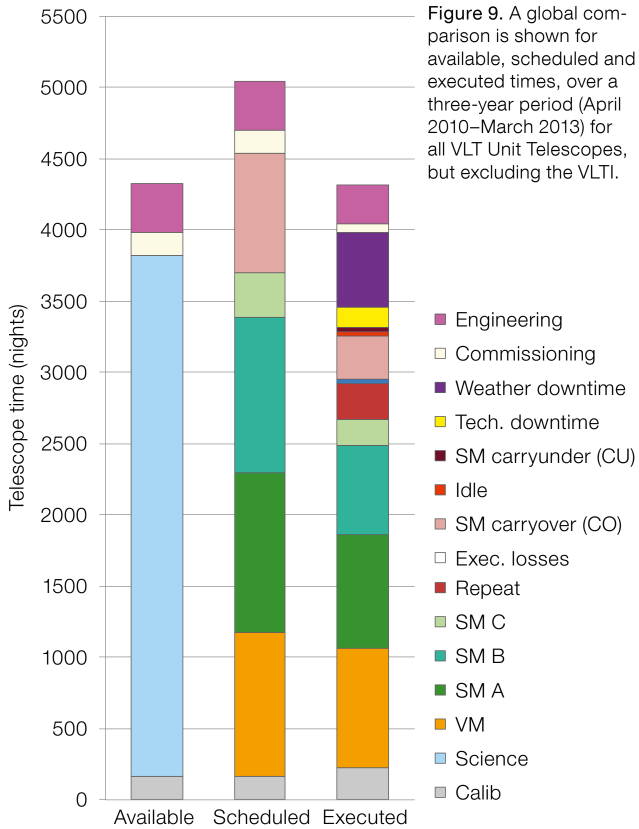

Figure 9. A global comparison is shown for available, scheduled and executed times, over a three-year period (April 2010–March 2013) for all VLT Unit Telescopes, but excluding the VLTI.

Putting it all together: The overall picture

Ideally, all this information could be captured in a single graph that takes all the relevant numbers into account. We have attempted this in Figure 9. Starting from the time available for scheduling, we looked at how this time was then scheduled and turned into successful execution. One can do this exercise per Period, but for the purposes of this article we decided to collapse all Periods between P85 and P90 into one graph — this possibly has the disadvantage of missing some details pertaining to specific Periods, but it clearly has the advantage of providing a global view over a recent three-year timespan. Figure 9 includes data for the four VLT Unit Telescopes, but omits statistics for the VLTI on account of its special features (e.g., block scheduling).

For a global perspective, one also needs to account for Visitor Mode; this appears under the only assumption that can be made, i.e., that everything scheduled was indeed executed, except for weather and technical losses. This downtime was proportionally distributed across the Periods, affecting the time assigned to SM and VM according to their allocation percentages. The result is apparent, for instance, in the orange strip of the second and third columns of Figure 9: the VM execution strip is indeed slightly shorter than the one that was scheduled, on account of this weather and technical downtime.

Figure 9 displays considerable granularity, especially in the two rightmost bars, corresponding respectively to the scheduled and executed times. Here “executed times” refer to what was executed within the available time; the numbers corresponding to Calib, SM A/B/C/CU/CO (carryunder/carryover) refer however to successfully executed observations. Before each OPC meeting, the Observatory releases a technical time schedule (shown in the first column). This includes the time reserved for technical activities in the next semester, split into engineering (shown in fuchsia in Figure 9) and commissioning (creamy-white) slots, as well as the time known (from previous Periods) likely to be invested for night-time calibrations (shown by the grey region at the bottom). The time remaining in Figure 9 is labelled as Science (shown as light blue) in the first column. This is the amount of time available for science observations and that needs to be filled by the schedulers (second column). VM slots are usually reserved first. What remains can be used for SM observations and is split evenly between A and B class runs. C class runs (sometimes referred to as fillers, because they typically have more relaxed observing constraints) come on top of the already 100 % full schedule, thus representing the protective buffer against idle time and/or bad weather conditions.

As is apparent from the third column of Figure 9, weather downtime can be significant in some Periods, making a 10 % provision insufficient, thus affecting completion efficiencies and creating larger carryovers. Looking then at the bar of the execution (third column), granularity increases because carryunders, carryovers, idle time, execution losses (time lost because of operational issues), repeats (i.e., time invested in SM observations that did not fulfil the requested conditions) can all be traced. It is only at this stage that we can properly close the loop on all the technical activities (i.e., the time that was truly invested in engineering and commissioning activities) and downtime (both weather and technical).

It should be noted that, until recently, no provision was made for weather downtime during scheduling, which had a major impact in some Periods. For example the weather downtime (purple area of the rightmost bar of Figure 9) is unusually large and dominated by the significant weather loss recorded for Period 85 — close to 20% of the available science time — and for P86–89 — close to 15%. Similarly, the time already committed to carryover runs was accounted for only in part, as becomes evident from the much longer bar in the second column.

From the Executed column in Figure 9, it also becomes apparent that: *i)* the UTs suffer from quite negligible idle time and very little execution loss; *ii)* carryovers and repeats (observations outside user constraints) instead have a large impact on the overall efficiency. The graph also confirms that B class runs are the most affected in terms of completion rates. As already mentioned, this is not really surprising: B class runs are usually highly demanding in terms of constraints (comparable, in fact, to A class runs) but in competition with higher priority (A class) runs.

The categories carryover and repeat, however, touch upon two different aspects of the ESO operational model: carryovers are the result of too aggressive or sub-optimal scheduling; the number of instances of repeats can in principle be reduced, but the reason(s) for them must first be properly understood. All OBs are started within constraints, but atmospheric conditions can change even within the first hour (the average OB length). From preliminary investigations, it appears that approximately 40% of repeated observations are due to varying seeing conditions and another 20% due to other variable meteorological conditions. This finding could indicate that the seeing stability that we assume at the time of scheduling is not realistic and/or that the first data quality evaluations (QC0) are too strict. The number of repeated observations varies strongly with instrument and correlates with operational complexity and sensitivity to good conditions (e.g., AO instruments tend to have larger numbers of repeated observations).

What's next

The Very Large Telescope (and Interferometer) is a complex facility. It is unique in the fact that it provides four 8-metre-class telescopes to its community 365 nights per year (except for VLTI with the Unit Telescopes, which has so far only been available in fixed time slots). Fifteen years after the start of VLT/I operations, all foci remain occupied with a suite of forefront instrumentation that allows a variety of scientific questions to be tackled. The facility remains in high demand, as seen by the number of new (unique) PIs who are attracted every Period (approximately 40% of the total).

But on account of this breadth and variety, the operation of the facility is challenging. Community feedback is generally very positive, the support interfaces are functioning well and highly valued. However, from the ESO side, it is also important to back up this positive attitude with facts and solid numbers. The effort invested so far in the DOME project is just starting to provide the baseline metrics required to analyse operational efficiency in a systematic way. Some of the graphs shown in this report seem to indicate areas that require closer scrutiny and that, if confirmed, may trigger optimisation processes.

Some actions have already been taken, such as more realistic provision for carryover and weather downtime at the time of scheduling, as well as an update in the scheduling tool for the underlying characterisation of the Paranal site in terms of seeing (distribution and stability), transparency conditions and photometric stability. On the side of operational procedures and/or science policies, ESO will now look into the low completion rates of B class runs, into the repeat category, into a renewed definition of filler programmes and ultimately into the scheduling tool, in order to make it more dynamic and able to react on shorter timescales.

In terms of the DOME project, none of the numbers and figures shown here include Large Programmes. They represent up to 15% of the total available time, secure large datasets and aim to tackle the top scientific questions of the day; thus it is definitely worthwhile to take a closer look at them. Similarly, in the near future, metrics will be further broken down (per instrument, per programme type, etc.) or the efficiency of the most complex observing modes (e.g., AO-assisted) may be examined. Part of the DOME project also consists of the monitoring of the metrics that have been extracted so far, in order to be able to uncover the trends and spot changes as quickly as possible.

Our final goal is to increase the scientific return and impact of the Observatory by optimising operations and its associated implementation. For this, we need to be able to retrieve a detailed view of parameters that measure the status and health of VLT/I operations, at any time, and be able to answer questions such as: is the current SM/VM ratio satisfactory; can any further optimisation help in improving our efficiencies and the satisfaction of the user community? The scientific return is very high, but there is clearly room for improvement. ESO will continue to look into means to improve the handling and execution of your programmes.

References

Bierwirth, T. et al. 2010, Proc. SPIE, 7737, 77370W
Comerón, F. E. et al. 2003, The Messenger, 113, 32
Hanuschik, R. & Silva, D. 2002, The Messenger, 108, 4
Patat, F. & Hussain, G. 2013, in *Organizations, People and Strategies in Astronomy 2* (OPSA 2), ed. Heck, A., 231
Primas, F. et al. 2012, Proc. SPIE, 8448, 84480L
Quinn, P. et al. 2000, Proc. SPIE, 4010, 56
Silva, D. 2001, The Messenger, 105, 18

Links

[1] VLT/VLTI Science Operations document: http://www.eso.org/sci/observing/policies/Cou996-rev.pdf
[2] Full report on the 2014 User Feedback campaign: http://www.eso.org/sci/observing/phase2/otherInfo/UserFeedback.html