The Scientific Return of VLT Programmes

Michael Sterzik[1]
Christophe Dumas[2]
Uta Grothkopf[1]
Andreas Kaufer[1]
Bruno Leibundgut[1]
Stephane Marteau[1]
Silvia Meakins[1]
Ferdinando Patat[1]
Francesca Primas[1]
Marina Rejkuba[1]
Martino Romaniello[1]
Felix Stoehr[1]
Lowell Tacconi-Garman[1]
Ignacio Vera[1]

[1] ESO
[2] TMT Observatory Corporation, Pasadena, USA

An in-depth analysis of the publications from 8414 distinct scheduled VLT observing programmes between April 1999 and March 2015 (Periods 63 to 94) is presented. The productivity by mode (Visitor or Service Mode) and type (Normal and Large, Guaranteed Time, Target of Opportunity, Director's Discretionary Time) are examined through their publication records. We investigate how Service Mode rank classes impact the scientific return. Several results derive from this study: Large Programmes result in the highest productivity, whereas only about half of all scheduled observing programmes produce a refereed publication. Programmes that result in a publication yield on average two refereed papers. B rank class Service Mode Programmes appear to be slightly less productive. Follow-up studies will investigate in more detail the parameters that influence the productivity of the Observatory.

Introduction

Very Large Telescope (VLT) observing programmes have been offered in Visitor Mode (VM) and Service Mode (SM) since the beginning of regular operations in 1999. Observing programmes* are implemented through a variety of different programme types such as Normal Programmes, Large Programmes (LP), Guaranteed Time Observations (GTO), Director's Discretionary Time (DDT), Target of Opportunity (ToO; including Rapid Response Mode), Calibration, and Monitoring Programmes, each attuned to the needs of proposers or operations. All programme types are described in detail in the ESO Call for Proposals[1]. GTO Programmes are almost exclusively executed in VM, while Monitoring, Calibration, DDT, and ToO Programmes are usually done in SM. Normal Programmes and LPs are pursued in both SM and VM.

Both Visitor Mode and Service Mode offer specific advantages. VM allows visiting astronomers to adapt their observing strategy in real time and ensures close relations between the visiting astronomers and the observatory staff. It also provides an opportunity for young researchers to gain hands-on observing experience at the Observatory. SM is designed to optimally use the range of observing conditions according to the observing constraints of the different science cases and enables execution of science programmes requiring particularly stringent conditions (Silva, 2001). For SM Programmes, different rank classes (A, B, C) are assigned by ESO, following the evaluations by the Observing Programmes Committee (OPC) according to the VLT/VLTI Science Operations Policy[2]. The top-ranked SM proposals are assigned an A rank, followed by B rank and the filler programmes, with more relaxed observing constraints, receive rank class C. Hence the rank classes set priorities for the execution of SM observations. It should be noted that VM Programmes are only scheduled if they are rated at the same level as SM A rank class Programmes.

The scientific productivity of the VLT (and other ESO facilities) compares very well with other observatories in terms of global bibliometrics[3]. Here we investigate the impact of science operations, in particular on how programme modes, types and rank classes compare to each other in terms of their scientific return.

Through the identification of specific strengths and weaknesses of the current operational model, as revealed through this analysis, we hope to enable improvements in the mid-term and to prepare for the future integrated VLT and European Extremely Large Telescope (E-ELT) operations scheme in the next decade.

Methodology

We utilise ESO-internal databases that collect information from Phase 1 (proposal submission) and Phase 2 (observation preparation and execution). Some of these databases were set up before the start of VLT operations, and over time, several modifications have been introduced to improve consistency and completeness. Among other studies, these databases enable the analysis of operational metrics and efficiency, which have been presented in Primas et al. (2014). We extracted the following information from the databases: scheduled observing proposal with its associated unique programme identification (ID, e.g., 089.A-0118), telescope, instrument, observing mode (VM or SM) and allocated time (nights for VM and hours for SM).

Observing programmes may consist of several runs, and these are identified through their run IDs (alphabetical letter added to the programme ID, e.g., 089.A-0118(B)). Each of these runs has its own time allocation, instrument, mode, etc. Observing runs are evaluated and ranked individually by the OPC.

A large fraction of programmes ask for only a few runs (68 % have one run and another 20 % two runs per programme). The number of "mixed mode" programmes (with SM and VM runs) is low (4 % by number and 9 % in time). Programmes with allocations on the VLT together with other ESO telescopes are more frequent (5 % by number and 12 % in time). Many multi-run programmes are LPs, while some Normal Programmes have multiple runs requesting, for example, different instruments to cover the same objects in one programme, different constraint sets or specification of several epochs to follow a variable object. The OPC rarely recommends different rank classes for individual runs within a programme. Other reasons for runs with different rank classes within an observing

* For the purposes of this study, we disregard VLTI programmes and spectroscopic surveys. Surveys on the VLT were introduced only in 2011. There are currently three such surveys: Gaia–ESO with FLAMES; Lega-C and VANDELS with VIMOS.

programme can be different pressures on telescopes or operational reasons, e.g., when a pre-imaging run receives a rank A while the corresponding spectroscopy run has B rank. The number of programmes with mixed rank classes amounts to 6 % by number and 9 % in time. Since the analysis is carried out at the programme level, we had to decide which rank class to assign to a programme. We added the time allocated to all runs and picked the rank class with the largest time allocation.

The ESO Telescope Bibliography (telbib[4]) provides the bibliometric information on the refereed publications for this study. The association of a publication with an observing programme is accomplished through the ESO programme ID. Cross-references between papers and programmes within the database are considered to be complete to over 95 %. We selected all papers associated with a programme including archival papers. Citations are drawn from the Astrophysics Data System (ADS) and the caveats on completeness of the ADS apply[5].

Table 1 summarises the 8414 distinct observing programmes scheduled between the start of Unit Telescope 1 (UT1) operations in April 1999 (ESO Period 63) and March 2015 (end of Period 94). The analysis was restricted to the programme types Normal, Large, DDT, GTO and ToO. Short Programmes, which were offered between Periods 80 and 86 have been grouped with Normal Programmes. Calibration Programmes and Monitoring Programmes were first implemented in Periods 82 and 92, respectively. Since both Calibration and Monitoring Programmes are very few in number, we refrained from a separate analysis.

The fractional time distribution between VM/SM Programmes is 32 %/68 % (and 26 %/74 % in terms of number of programmes). VM Programmes typically received a median (mean) time allocation per programme 64 % (30 %) higher than SM Programmes. This is due to a combination of VM Programmes applying for full or half nights and many SM Programmes requesting short observations of less than one night. Among the SM Programmes, C rank class Programmes contribute approximately 10 % to the total

Programme (Mode, Rank, Type)	No. of programmes	Median TTA per programme (hrs)	Mean TTA per programme (hrs)	TTA (nights)
Total	8414	12.2	17.1	16028
VM	2228	18	20.7	5130
SM	6186	11	15.9	10898
A rank	2672	9.6	16.2	4807
B rank	2841	11.5	14.3	4515
C rank	673	16	21.1	1576
Normal	6705	14.0	16.4	12216
Large	80	170	209.4	1862
GTO	498	9	15.8	872
DDT	689	3.5	4.8	371
ToO	416	12	14.6	672
Calibration	17	8	8.1	15
Monitoring	9	16.5	19.7	20

Table 1. The Total Time Allocation (TTA) statistics for the different observing programme modes, SM rank classes and types for Periods 63 to 94. A night is taken to be nine hours.

number and intentionally over-schedule the observing queues. This procedure ensures that the available observing parameter space is filled, statistically, with suitable programmes (Silva, 2001).

Although there are relatively few Large Programmes (less than 1 % by number), they account for a fair fraction of the observing time (12 %). The selection process for LPs occasionally resulted in the reduction of the allocation to only one period and hence to LPs with scheduled time below 100 hours. We kept all LPs, independent of their total allocation, in our sample. One VLT Imager and Spectrometer for the mid-InfraRed (VISIR) Large Programme, which was not started due to the delayed VISIR upgrade, was excluded. Twenty LPs had runs scheduled at the VLT and other ESO telescopes. We tried to assess the relative importance of the respective telescope allocations by examining the scheduled time per telescope. Fourteen LPs have a majority allocation on the VLT and are counted as VLT LPs. The remaining six LPs were excluded from the analysis. The number of LPs considered in our VLT analysis is 80 as given in Table 1.

It should be noted that LPs in Service Mode are normally allocated an A rank, but a few LPs requiring very loose observing constraints were scheduled as filler programmes with a C rank class.

Productivity

The number of refereed publications can be used as an indicator of the scientific productivity of a programme. This implicitly assumes that all programmes received the full requested observations and the corresponding data, which is not true in all cases. Weather or technical time losses during VM and non-completion of SM Programmes will affect the ability of the community to produce scientific results.

Effects of programme completeness

The productivity of programmes depends on the completeness level of observing programmes, in particular for B and C rank-class SM Programmes. As examples we mention here the fact that some ToO Programmes did not trigger observations as no suitable transient object appeared during the allocated time. About 50 % of the allocated ToO time is typically not used. This explains some of the low numbers of publications per allocated night. Conversely, the fraction of DDT publications is fairly high as the programmes often have a very direct scientific goal and the small allocation leads to fast publication.

An analysis of the completeness for SM runs has recently been presented (Primas et al., 2014). Since ESO Observing Period 78 (October 2006), the overall completeness distribution fractions in terms of number of runs and number of hours for the different SM rank classes have been systematically recorded. The completeness fraction is defined as the

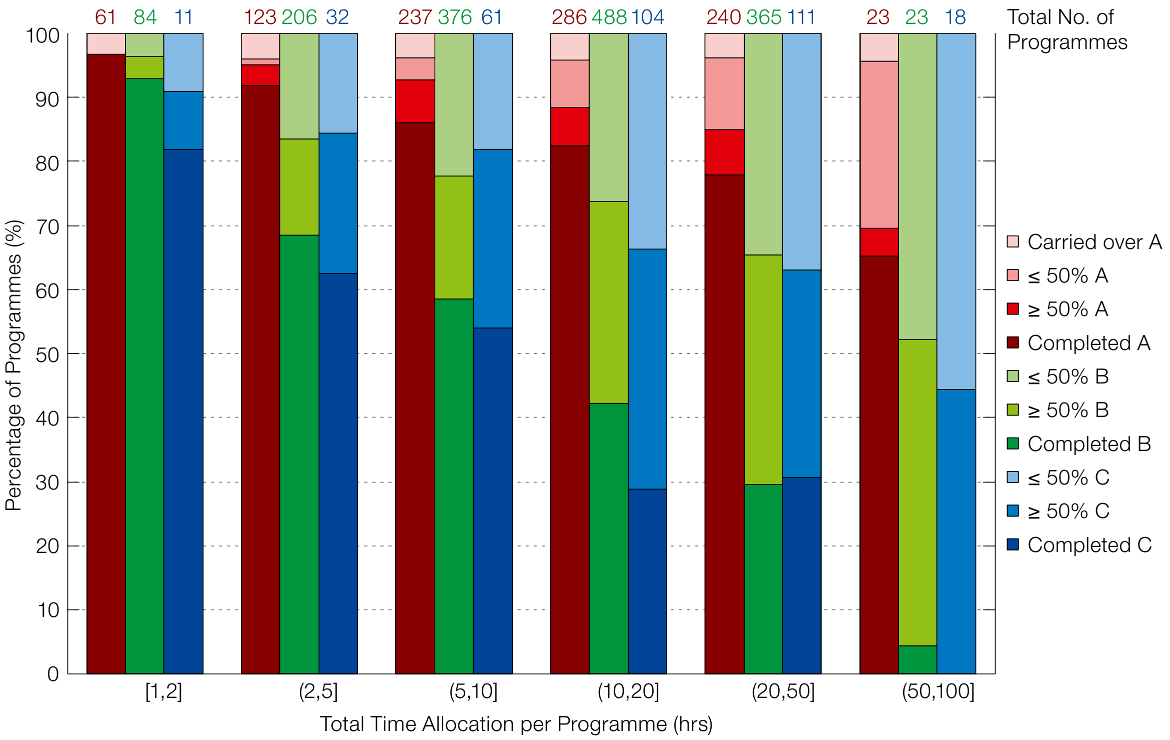

Figure 1. The fraction of observing programmes with different completeness levels separated by rank classes A, B, C, binned by total time allocation in hours. Each bar in a given allocation bin gives the completed, "≥ 50 %" and "≤ 50 %" fractions. For A rank class Programmes, the fraction of programmes with a carry-over status at April 2015 are included. The numbers at the top of the plot refer to the total number of programmes in each bin (= 100 %).

actual time successfully executed for a given run divided by the total time allocated to this run. At the end of every period, incomplete A rank class observing runs are judged on their science status and are either declared completed (in most cases the runs are nearly fully complete) or are carried over into the next period. Some runs may be terminated for external reasons (e.g., users request that the programme be stopped or technical feasibility issues require its curtailment). About 15 % of the B rank class runs were not started at all (Primas et al., 2014, Figure 6). The percentage of programmes that have not been started is even smaller. As long as there is only one run in a programme, the individual run completeness fraction applies. For multi-run programmes, we determined the completeness fraction as the average over all its runs weighted by the allocated time.

In the following we compare SM Programmes of the same type only. Therefore, we have selected a number of 968 A rank, 1550 B rank and 338 C rank class Programmes of type Normal and Short between Periods 78 and 94, for which we have reliable completeness information. We introduce three completeness groups: fully completed; more than half, but not fully, completed (designated "≥ 50 %" in the following); less than half (≤ 50 %) completed. Also a few (38) programmes that currently hold carry-over status are included. We do not include 282 programmes in the same period range that were never started, and consequently have a completeness fraction of 0 %. Most of these programmes were in B rank class.

Figure 1 displays how the completeness level of Normal and Short Programmes with different rank classes depends on the programme length. The fraction of completed A rank class Programmes is generally high as the Observatory commits to these programmes, e.g., through carry-over into future observing periods. More than 80 % of all A rank class Programmes with less than 50 hours allocated time were completed. We also explicitly show the fraction of A rank class Programmes being carried over in April 2015. Less than half of the B and C rank class Programmes with more than ten hours allocated time were fully completed; for the C rank class Programmes this is almost by design as they are filler programmes.

Table 2 compiles the productivities for subsamples with different completeness levels for these A, B and C rank class Programmes. It presents the fraction of programmes leading to a publication (4th column), the average fraction of publications over all programmes (6th column) and the average number of publications per programme that produced a publication (7th column). Analysis of publications is presented in the following section.

As expected, higher levels of completeness usually lead to higher productivity for the programmes. Completed programmes, and those with more than 50 % completeness, result in a significantly larger number of publications than those with a low observational return. This emphasises the importance of programme completion as an essential parameter to ensure the science return.

Amongst the fully completed programmes, the B rank class Programmes produce on average fewer publications per programme than A or C ranked Programmes. This effect is much smaller for the programmes that are at least half completed. Interestingly, the programmes that lead to publications typically produce more than one paper in all rank classes. This appears to be true even for programmes that obtained only very limited data.

Programme (Rank)	No. of programmes	No. of programmes with publications	Fraction of programmes producing a publication	No. of publications	Fraction of publications per programme (all programmes)	No. of publications per programme
Completed	1709	605	0.35	1023	0.60	1.69
A rank	820	303	0.37	553	0.67	1.83
B rank	762	248	0.33	368	0.48	1.48
C rank	127	54	0.43	102	0.80	1.89
≥ 50% completed	582	211	0.36	382	0.66	1.81
A rank	72	24	0.33	40	0.56	1.67
B rank	402	147	0.37	263	0.65	1.79
C rank	108	40	0.37	79	0.73	1.98
≤ 50% completed	565	108	0.19	162	0.29	1.50
A rank	76	11	0.14	19	0.25	1.73
B rank	386	76	0.20	111	0.29	1.46
C rank	103	21	0.20	32	0.31	1.52

Table 2. Average productivities of Normal Programmes in Service Mode (starting with Period 78) by completeness level for different rank classes. Three completeness categories are presented: fully completed (Completed); more than half, but not fully completed (≥ 50% completed); less than half completed (≤ 50%).

Further investigations considering the impact by actual observed time and including additional parameters — like programme length, instrument, observing conditions and observing requirements — need to be considered to interpret the findings.

Publication analysis

We have used the telbib database to correlate the ESO observing programmes with their refereed publications for all programme types and for the full length of VLT operations (Period 63 to 94). The results can be found in Table 3. There were 5907 publications based on VLT data to April 2015. Of the 8414 programmes allocated time, 3675 have contributed to publications. Many programmes contribute to more than one publication (compare columns 3 and 4 in Table 3) and conversely many papers are based on several programmes: the 3675 programmes are mentioned 11 291 times in refereed papers. During the last 16 years a VLT programme contributed on average to 1.34 publications (= 11291/8414). The publications are based on only 44 % (= 3675/8414) of all observing programmes. These are lower limits as there are a number of programmes that were not started (e.g., ToO, some SM Programmes) and hence the total number of programmes that received data is smaller. On average, one VLT Unit Telescope (UT) night resulted in 0.7 publications (= 11291/16028). Primas et al. (2014, Figure 9) have shown that typically only 75 % of the scheduled time results in useful science observations and hence the previous number can be corrected to 0.9 publications per useful UT night. A further increase is expected as we have also included recent programmes, and many of these may not yet have produced a publication.

There appears to be a significant fraction of observing programmes that do not result in publications and it is in the Observatory's interest to understand the reasons. It should be noted that the above statistics are based on allocated programmes and not on completed observations and hence the numbers are strict lower limits. The exact impact of completeness of the individual programmes is difficult to assess, as it may depend on many parameters (see discussion above) and we did not try to correct for it.

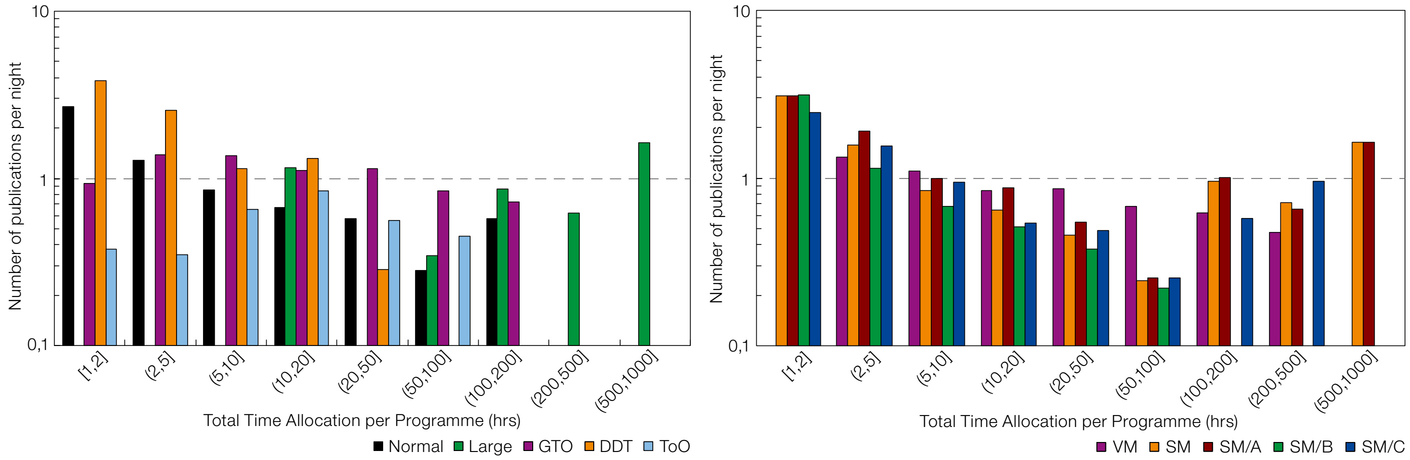

Figure 2. The dependence of productivity, by programme type (left panel) and observing mode and SM rank classes (right panel), in terms of publications per night is shown for the total time allocation per programme.

The productivity of the different observing modes and programme types shows some interesting features. It is noteworthy that Large Programmes produce many publications and are as productive per night as the other programme types. The GTO publication rate is high, which points towards a good match of the GTO science cases with the instrument and a high level of preparation of the GTO teams for the data analysis, e.g., through exquisite knowledge of the instrument and preparation of the instrument pipeline. Also the productivity of DDT Programmes is quite remarkable, profiting from their short duration and fast return. The least productive observing modes are SM Programmes in rank classes B and C.

Figure 2 shows the distribution of productivity per night for the different programme types (left panel), and modes and SM rank classes (right panel) as a function of allocated time. Programmes with a small allocation producing a publication are clearly favoured in such a comparison. Examination of Figure 2 demonstrates that:

- Large Programmes exhibit relatively constant normalised productivities. In general, LPs are allocated over 100 hours, but in rare cases the OPC recommended smaller allocations. These programmes are still listed as LPs in the database. This explains the small number of LPs with less than 100 hours in Figure 2. The number of publications per night is still rather high (mostly above 0.6) despite the large time allocations.
- DDT and Normal Programmes show decreasing normalised productivities with increasing allocated time, while GTO and ToO Programmes have a relatively flat distribution.
- VM Programmes also exhibit a rather flat normalised productivity distribution. They have the highest normalised productivities for programme lengths between 20 and 100 hours (approximately 2 to 11 nights).
- SM Programmes exhibit a high publication rate for short programmes and again for the LPs with > 100 hours allocations.

The trend for A, B and C rank class Programmes is very similar below 100 hours. A strong drop can be observed for programmes with more than 30 hours allocated. The fraction of completed programmes in this range is 30 % or less for B and C rank classes, while it drops below the 80 % mark for A rank class. However, given that non-completed programmes do publish, and that the fraction of programmes producing a publication is not dramatically lower for the programmes

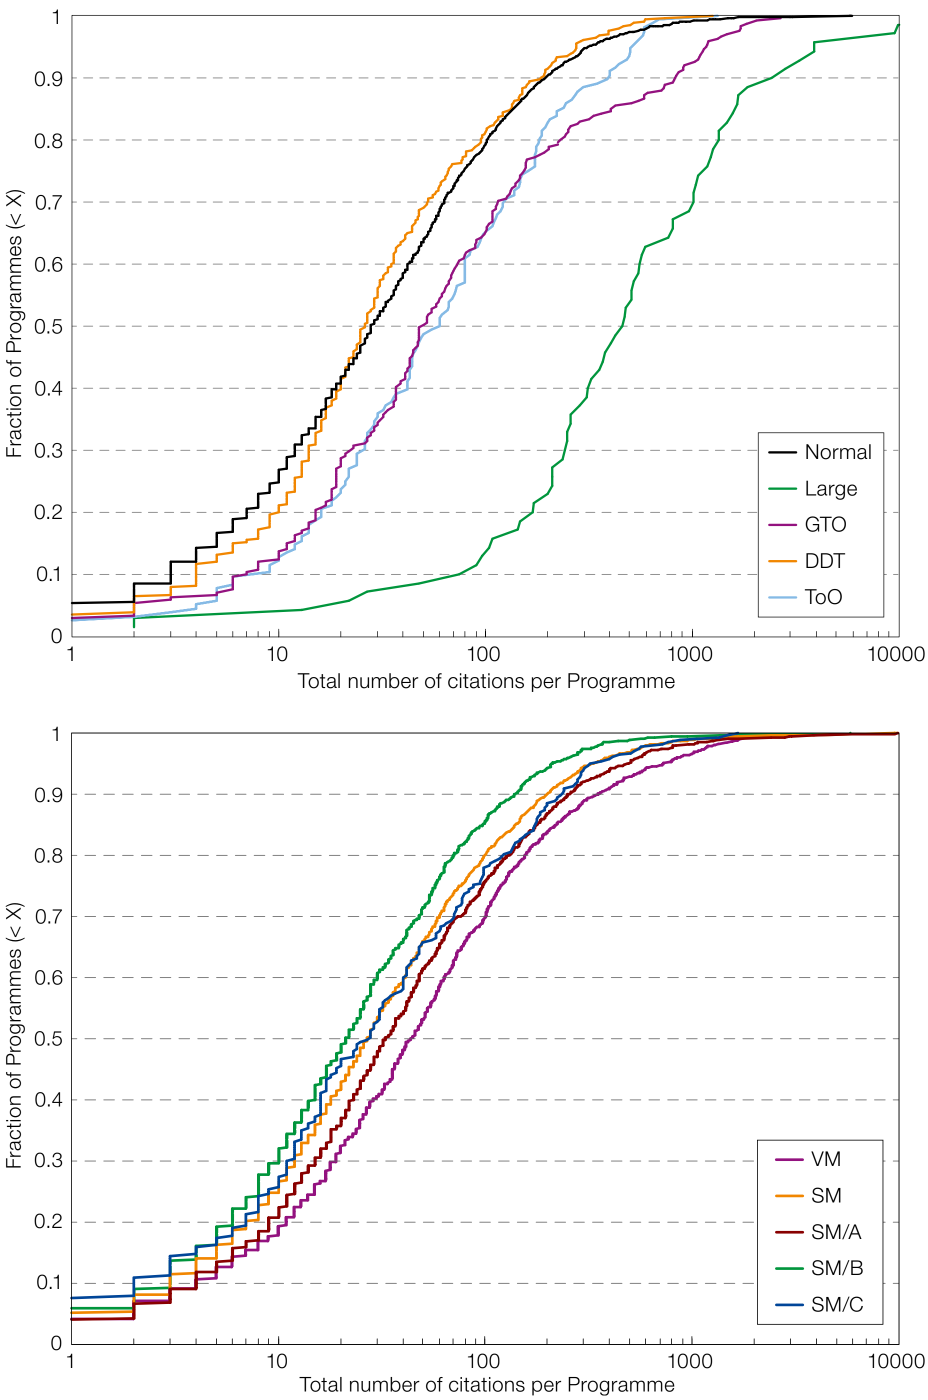

Figure 3. The cumulative distribution of the total number of citations for those programmes that produced at least one publication. The sum of all citation counts is used in case more than one publication contributes. The top panel shows the distribution by programme type and the lower panel by observing mode and SM rank class.

Programme (Mode, Rank, Type)	No. of programmes	No. of programmes with publications	Contributions to publications	No. of publications per programme	No. of publications per allocated night
Total	8414	3675	11291	1.34	0.70
VM	2228	1205	4211	1.89	0.82
SM	6186	2470	7080	1.14	0.65
A rank	2672	1186	3956	1.48	0.82
B rank	2841	1014	2292	0.81	0.51
C rank	673	270	832	1.24	0.53
Normal	6705	2863	7776	1.16	0.64
Large	80	70	1483	18.5	0.80
GTO	498	241	960	1.93	1.10
DDT	689	342	633	0.92	1.70
ToO	416	156	436	1.05	0.65
Calibration	17	3	3	0.18	0.20
Monitoring	9	0	0	0	0

Table 3. Publication statistics of VLT observing programmes by mode and SM rank class and type.

that are ≥ 50 % completed, the impact of completeness is not so straightforward to analyse. The cause for this distribution requires further investigation.

Impact

Citations can provide a measure of the scientific impact of observing programmes. We have obtained the number of citations for a given publication from the ADS.

The cumulative distributions of the total number of citations for the different programme types and the modes are shown in Figure 3. All citations of all publications for a given programme are summed. We chose to present the cumulative distributions to remove the effects of the sometimes very different sample sizes (cf., Table 1).

The median for Normal Programmes is approximately 20 citations, i.e., about 50 % of all Normal Programmes have produced more than 20 citations from their publications. The highest number of citations is reached by the Large Programmes: half of them resulted in more than 300 citations. The median citation count for GTO and ToO Programmes is around 50, while it is about 20 citations for DDT Programmes.

The median for all SM Programmes is around 25 citations and is the same for C rank class Programmes bracketed by A and B rank classes with 30 and 20 citations, respectively. VM Programmes have a median citation of about 40.

Nevertheless, all programme types produce a high science impact of more than 100 citations. Ninety percent of the LPs achieve more than 100 citations and some 20 % over 1000 citations. About 20 % of ToO and GTO Programmes also generated more than 100 citations.

Caveats

Some of the approximations and assumptions that have been made for this analysis have already been described in the relevant sections. In particular, we have not accounted for the fact that typically a delay of several years is seen between receipt of data and publication. Considering that it takes roughly five years on average for a publication to appear and hence we have not accounted for almost a third of the lifetime of the VLT, we expect the number of publications to increase by about 30 % in the future.

The programme implementation and classification policies have changed relatively little, but new instruments have been offered and the underlying model has evolved. Therefore subtle biases in the programme distribution and their performance over time cannot be excluded.

The records contained in the telbib database do not consider or evaluate the contributions of individual datasets of programmes to the scientific goal of a publication. On the one hand, some data (and programmes) contribute critically and others just provide complementary or ancillary data. Effectively, this may overestimate the scientific return for some programmes in publications contributing multiple datasets. On the other hand, the total time required to contribute specific data to a publication may not use the entire time allocation of the full programme. The assumption of accounting for the total time of an entire programme therefore overestimates the actual observing time attributed to a publication. There is no simple methodology to weight the different contributions objectively and we have not attempted to do so.

The incompleteness of SM Programmes of different rank classes introduces various biases. We have investigated its effect on programme productivity, but not on its associated citation statistics. Incomplete programmes lack specific observations and likely reasons are the (expected or actual) violation of observing constraints or some statistical overscheduling of constraint parameter space. The dependency of observing constraints and conditions for a given programme on its subsequent science return is another important performance indicator of SM Programmes and their implementation, and deserves further analysis.

Results and conclusions

The main findings of this study of the scientific return of VLT observing programmes are as follows:

1. Large Programmes (LPs) have by far the highest scientific productivity and impact. As expected, they fulfil their role of providing major scientific

advances, breakthroughs and often have a high legacy value. Normalised by the allocation time, their productivity per night is at least as high as for Normal Programmes. Thus LPs have proven to be a highly valuable asset in the strategic distribution and implementation of VLT programmes.
2. Most of the telescope time available at the VLT is allocated for the execution of Normal Programmes. They produce most of the VLT publications. However, their impact is relatively small when compared to the other programme types (cf., Figure 3). Still, the community prefers Normal Programmes as shown by the results of the ESO2020+ Users' Poll (Primas et al., 2015).
3. GTO Programmes have on average a higher impact than Normal Programmes, supporting their role as pathfinders using novel instrumentation for cutting-edge science cases. Some GTO Programmes are coordinated observations over several periods and resemble LPs in this respect.
4. ToO Programmes have on average a higher impact than Normal Programmes. Considering that many ToO Programmes are not triggered, the impact has to be regarded as even higher. This reflects the growing interest in the astrophysics of the variable sky.
5. DDT Programmes typically constitute small investments of telescope time and are targeted at specific, "hot" scientific questions that can lead to quick publication. Often they complement existing data to confirm or strengthen a result and represent less than 2 % of the scheduled observing time. They are productive in terms of number of publications per telescope night, but their absolute impact (in terms of citations) remains limited, as shown by Figure 3.
6. VM Programmes exhibit high productivity and impact, in particular for Normal Programmes with telescope time allocation of a few nights. In this parameter range, the specific strengths of VM allocations pay off: the visiting astronomer can optimise the observing strategy and implement back-up programmes to adjust in quasi real-time to changing observing conditions, thus possibly securing a higher data return.
7. A and C rank class SM Programmes yield, on average, nearly two refereed publications for programmes that produced a publication. B rank class Programmes do not produce as many publications. If one compares the number of publications per night, B and C rank class Programmes show lower publications than A rank class and Visitor Mode Programmes.
8. SM A rank class and VM Programmes have similar completion fractions, if we assume that statistically Visitor Mode Programmes receive about 85 % of the allocated time (allowing for ~ 15 % weather and technical downtime; Primas et al. [2014], Figure 9) while the B and C rank class Programmes are more incomplete. The number of publications per programme, and per allocated time, increases with completeness fraction for all ranks. It is quite possible that statistically most B rank class Programmes lose their programme execution competition to A rank class Programmes in the same, demanding, observing constraint conditions. We plan to investigate this aspect in the future in more detail.
9. The publication ratio for all VLT programmes (i.e., the number of programmes that published at least one refereed paper, divided by the number of all VLT scheduled programmes) is 44 %. This is a strict lower limit as the time between observations and publication is typically five years and we can expect an increase of about 30 % in this number over the next few years.

It is clear that attention must be given to the implementation of B rank class Programmes, and in particular to facilitate, if possible, an increase of their completion rates. C rank class Programmes exhibit a relatively strong performance.

Together with the community we should also try to better understand the reasons why a significant fraction of VLT programmes do not lead to results published in refereed journals. ESO plans to poll the principal investigators for the reasons for not publishing.

An important scientific return comes from archival research. We have included archival papers in our analysis, but did not separate them out for study, as we were primarily interested in the operational aspects. The impact of archive research will be investigated in a separate article (Romaniello et al., in preparation).

Several aspects of this study require follow-up analysis. The statistical investigation presented here is only one step towards a better global understanding of the complexity of the various processes that lead to advancement of scientific knowledge. The individual science cases, but also sociological factors within the science teams, may influence the scientific return, an aspect which is well beyond the scope of the present analysis.

Acknowledgements

We thank Piero Rosati for sharing his unpublished study with a similar scope made in 2011. Discussions with Ralf Siebenmorgen about the methodology and with Jeremy Walsh and Olivier Hainaut about the presentation of this article are highly appreciated. This research made use of the NASA Astrophysics Data System.

References

Grothkopf, U. et al. 2007, The Messenger, 128, 62
Grothkopf, U. & Meakins, S. 2012, The Messenger, 147, 41
Primas, F. et al. 2014, The Messenger, 158, 8
Primas, F. et al. 2015, The Messenger, 161, 6
Silva, D. 2001, The Messenger, 105, 18

Links

[1] ESO Call for Proposals, e.g., the most recent: http://www.eso.org/sci/observing/phase1/p97/CfP97.pdf
[2] ESO Council resolution on VLT/VLTI Science Operations Policy: http://www.eso.org/sci/observing/policies/Cou996-rev.pdf
[3] telbib statistical summary: http://www.eso.org/sci/libraries/edocs/ESO/ESOstats.pdf
[4] ESO telescope bibliography, telbib: http://telbib.eso.org
[5] ADS citations: http://doc.adsabs.harvard.edu/abs_doc/faq.html#citations