

The Time Allocation Working Group Report

F. Patat¹

¹ ESO

A Time Allocation Working Group was charged with the task of reviewing the telescope time allocation process at ESO. The working group submitted a report to the Director for Science, including a set of recommendations and suggestions for an implementation plan. This paper gives a general overview of the recommendations and a status report on their implementation.

The Time Allocation Working Group

The ESO 2020 prioritisation initiative prompted a series of actions (Primas et al., 2015). One of these was to constitute a Time Allocation Working Group (TAWG). The TAWG included experts in the field of scientific resource allocation and representatives from the Users Committee (UC) and the Scientific Technical Committee (STC): Almudena Alonso Herrero (STC delegate), Antonio Chrysostomou (SKA, Science Operations Planning), Stefan Janssen (Paul Scherrer Institute, User Office), Arvind Parmar (ESA, Science Operations), Rachel Mason (Gemini Observatory), Neill Reid (STScI, Science Mission Office), Stephen Smartt (UC representative), and Ferdinando Patat (ESO, Observing Programmes Office) — the Chair of the TAWG and author of this article.

The TAWG reviewed the entire time allocation process as currently implemented at ESO, examining the procedures and analysing a number of statistical indicators. The recommendations in the report¹ were formulated with the aim of achieving two goals:

1. to maximise the scientific return of ESO facilities by selecting proposals that promise to result in significant advancements;
2. to improve the level of feedback provided to the community by the Observing Programmes Committee (OPC).

Given the current workload placed on the panel members, the TAWG reached the conclusion that the most urgent measure

ESO needs to take in order to achieve these goals is a drastic reduction in the number of proposals that each of the panel members have to review.

Although there is no clear indication from the literature of the optimal number of proposals per referee, the current number (over 70) is certainly above what is reasonable when trying to perform a thorough and consistent review. Most of the recommendations discussed in the TAWG report attempt to address this matter and aim at: a) reducing the total number of submitted proposals; b) increasing the average telescope time requested per proposal; and c) enhancing the quality of the review process and the feedback provided to the applicants — this last being the most recurrent complaint ESO receives about the whole process.

The TAWG considered possible radical changes to the way proposals are discussed at length, also considering substantial departures from the peer review schema. Notwithstanding the known limitations of this paradigm, the TAWG concluded that peer review still remains the most satisfactory way of selecting time applications. The choice of any time allocation committee is unavoidably subjective, and a different set of referees would provide a different list of top-ranked proposals, with only a fraction in common. However, in a scheme in which the community nominates panels, this type of selection still represents the priorities of the underlying population. Based on these considerations, the TAWG did not recommend changes in the overall process, although it advocated some substantial modifications. In general, the TAWG is in favour of deploying a distributed review, which is considered a valid means of sustaining the proposed fast-track channel (see below for more details on this channel). In this article, the main advantages and disadvantages seen by the TAWG for each recommendation are discussed, and suggestions for a possible implementation plan for the proposed recommendations are presented. The report is complemented by statistical studies probing various aspects of the process, some of which have already been published (Patat, 2016; 2017a; 2018).

The recommendations: an overview

The interested reader is encouraged to read the report; only a brief summary will be provided here. Perhaps not too surprisingly, some of the TAWG recommendations are very similar to those contained in the ESO OPC Working Group report (Brinks et al., 2012).

Recommendation 1: Decrease the frequency of the Call for Proposals to once per year

Given the artificial pressure that a semester-based call puts on the community, the typical time scales for publication return (see Patat et al., 2017) and the large number of resubmissions, there is no compelling scientific reason behind this frequency. A decreased frequency of regular calls could be compensated for by the creation of a new, fast submission and review channel (see next recommendation).

This change would bring a number of advantages; we mention only two of them here. The first is the increase of distinct referees who can be recruited within the same annual budget^a which allows the panels to cover more scientific areas. The second is that the number of submitted proposals will be less than a factor of two larger than the current value (~ 900 proposals/period), simply because it will remove all the resubmissions generated by the artificial right ascension limitation introduced by a semester-based system. Obviously, an annual call would introduce some loss of flexibility and response time (but see Recommendation 2) and it would also increase the difficulty of adding or changing plans regarding technical activities. This could be mitigated by the creation of a devoted group in charge of more dynamic short-term scheduling, which would in any case be required if a fast-track channel (FTC) were introduced (see Recommendation 2).

Recommendation 2: Create a fast-track channel for the VLT

This new peer-reviewed channel would be in addition to the existing Director's Discretionary Time (DDT) and is not meant to replace it. It would rather provide a quick duty cycle — on a timescale of months — allowing users to obtain data for amounts of time below some

threshold value. As in the case of similar channels that already exist (for example, the Fast Turnaround observing mode at Gemini²), specific criteria would be applied to eligible programmes to establish why they could not be submitted through the regular channel. To ease the review and scheduling processes, the deadline would be periodic (for example, once a month) and not continuous. ESO should also consider a maximum amount of time to be allocated to these programmes, which may be adjusted depending on how the community reacts to it.

The introduction of the FTC would have two beneficial effects: a) to improve the response time for science cases that do require a short duty-cycle; and b) to decrease the total number of proposals. The latter stems from the fact that applicants work to meet deadlines; the possibility of submitting short programmes at the next fast-track deadline will naturally lower the pressure, as users will not feel compelled to submit at every deadline.

While the FTC would have clear benefits for the community, it would also impact ESO operations. The review of FTC submissions would have to take place within a relatively short amount of time (depending on the final cadence of the FTC deadlines). In addition, it would require a strategy for dynamic scheduling in order to accommodate the newly approved programmes on top of those already allocated through the annual cycle described in Recommendation 1. The implementation plan proposed by the TAWG suggests ways to address these potential issues.

Recommendation 3: Radically change the proposal review procedure for the annual Call

In the current system, all proposals assigned to a panel are reviewed by all of the non-conflicted members of that panel. This results in 85% of the proposals being reviewed by 5 or 6 referees. The reviewers are required to read and grade all proposals; this corresponds to a load of more than 70 proposals per panel member and can reach up to about 100 proposals for each member of the “OPC proper”. It is widely acknowledged that this number is too large to allow a thorough review and can adversely affect the

quality of the feedback. This recommendation — which is in line with practices at other facilities (also outside astronomy) — is to stipulate that only a fraction of the panel should read and grade each proposal in the pre-meeting phase.

Once the pre-meeting grades were defined, a substantial triage (~ 30%) could be applied to the weakest proposals; triage is already standard practice in the current system. In the schema proposed by the TAWG, after triage was applied, the members of a given panel would subsequently have to read the non-triaged, surviving proposals that they had not read in the first round. The time freed by the load reduction would allow the referees to dedicate more time to assessing more meritorious proposals and providing feedback to all of the applicants.

This recommendation is based on the observed fact that the grade dispersion for higher-graded runs is lower than that for the lower-graded runs (Patat, 2018), so it is unlikely that a meritorious proposal would be triaged because of low number statistics. The TAWG did discuss the possibility of applying different selection criteria or figures of merit to select from the mid-range runs which have a higher dispersion, including semi-random selection, or schedule optimisation, but it did not make any specific recommendation in that direction. Obviously, having a reduced number N of reviewers increases the uncertainties related to the subjectivity of the process; however, the accuracy of the review increases only as \sqrt{N} . Therefore, from a statistical point of view, the difference between, $N = 6$ and $N = 3$ is not very significant (Patat, 2018).

Recommendation 4: VLT Large Programmes should have more massive allocations and shorter completion times

The current VLT Science policy includes a 30% upper limit for Large Programmes (LP) at the VLT. However, the time allocated to LPs is on average around 17%. LPs at the VLT can span up to four semesters. However, in a scheme in which LPs can only be requested on a yearly basis, there is no compelling reason that all the time should not be requested and allocated in one single cycle (barring time monitoring needs). As the statistics show

(Sterzik et al., 2016), LPs are the most productive programmes (even when normalised by the allocated time), and should therefore be promoted and completed in the shortest amount of time compatible with their scientific requirements. Therefore, the TAWG recommended that LPs should normally span only one year, and larger time spans would have to be justified on scientific grounds. In addition, ESO should consider increasing/removing the 30% limit if the demand goes in that direction as a consequence of the proposed changes. This goes along the lines of encouraging the community to cluster around larger projects, which are expected to produce a higher scientific impact (as opposed to an increasing fragmentation into small programmes — see below).

Recommendation 5: Redistribute the time allocated in favour of larger programmes

The purpose behind this recommendation is to encourage the submission of larger, more comprehensive requests. At the same time, it aims to limit the number of submissions of short proposals during the annual cycle by increasing their rejection rate. One could argue that the community would react by artificially increasing the requested times in order to get above the threshold. However, this depends on the threshold — if it were set to 30 hours, about 25% of the current proposals would have to triple their requested time, at least, to elude the restriction. Although this is not impossible, it would require a well-substantiated scientific justification.

The distribution of the time request on the VLT for normal programmes submitted in the last five years (Periods 93–102) is shown in Figure 1; 50% of VLT normal programmes request less than 12 hours, 75% less than 20 hours, and 95% less than 41 hours. Although normal programmes are allowed to request up to 100 hours, there is a “desert” above 50 hours; in general, proposals for more than 30 hours are rare (11%).

In this respect, it is instructive to view the trend in requests since the start of VLT operations. The evolution of the distribution of the time request (per proposal) is presented in Figure 2, which






Figure 1. (Above) Distribution function (upper panel) and cumulative distribution function (lower panel) of the time request per proposal in the last five years (Periods 93 to 102; 6540 proposals). The vertical dashed lines in the bottom panel indicate the 5th, 25th, 50th, 75th and 95th percentiles. The peaks observed at 20, 24, 30–32 and 40 nights are generated by the visitor mode requests for 2, 3, 4 and 5 nights and their conversion to hours (8/10 hours per night in even/odd semesters). Only normal VLT proposals are included.

Figure 2. (Below left) Time evolution of the time request per proposal from the start of VLT operations (Periods 63 to 102; 22800 proposals). The grey-shaded area indicates the inter-quartile region of the time distributions in each semester. The connected symbols trace the average (blue) and the median (orange) of the distribution. Only normal VLT proposals are included.

Figure 3. (Below right) Fraction of submitted proposals for five time request bins. The time intervals (in hours) are shown in the legend. Only normal VLT proposals are included.



shows different statistical indicators. This clearly demonstrates that there is a systematic and significant tendency to submit shorter (and more numerous) requests. Within a few semesters of the start of operations, the median VLT proposal stabilised at around 20 hours; since then it has steadily decreased to reach around 12 hours. The decrease is reflected in the entire distribution, which has shifted towards smaller time requests.

A closer look at the time evolution within time bins (Figure 3) reveals that the observed trend is related to a proliferation of programmes asking for < 10 hours; while in Period 70 (P70) this concerned only 18% of the proposals, in P101 the fraction reached 50%. There are likely to be various reasons behind this trend, relating to various demographic and sociological aspects. One factor may be linked to a general perception that it is easier to get small requests approved, and another factor is related to the conservative and risk-averse attitude of some reviewers who, for highly oversubscribed facilities, may prefer to allocate the same total amount of time to many small programmes rather than to fewer relatively large projects. These two factors lead to a negative feedback mechanism, which may explain the trends in Figures 2 and 3.

Although it is certainly true that ESO serves a wide community (about 3500 distinct scientists in the recent semesters; Patat et al., 2017a), time allocation is fragmented, and most proposals are comparatively small, both in terms of their time requests and their team sizes. The TAWG argues that, after twenty years of VLT operations, it is time to reconsider the way time is distributed. Setting a maximum fraction of accepted short programmes will encourage the community to propose larger, more ambitious projects. As shown by the ESO Survey of Non-Publishing Programmes (SNPP; Patat et al., 2017b), increasing the number of programmes requesting more than 20–25 hours would improve the scientific return of the VLT (at least in bibliometric terms). Also, applicants would be incentivised to complete their programmes more quickly (for example, completing samples with no right ascension limitations is easier with an annual cycle than with a semester-based cycle). Finally, since larger time requests require more thorough justifications and stronger teams, the total number of submitted proposals should decrease.

Recommendations 6–10

Whilst the first five recommendations affect the core of the time allocation process, the remaining five are more general and are discussed together here for the sake of conciseness.

Recommendation 6 is to create a proper filler channel. This is dedicated to proposals with loose observing constraints, which are allocated time at very low priority. Their review should be light, with a single, streamlined scientific evaluation that ensures a minimum standard, and a technical assessment to check that candidate projects conform to the filler criteria. Of course, nothing would prevent users submitting a proposal that required loose observing conditions through the regular channel for strong scientific cases. At the moment there is a deficiency in the request, leading to idle time, especially at some telescopes. The implementation of a filler channel should decrease the number of proposals requiring a full review and provide a sufficient number of reasonable projects to be executed under poor conditions.

Recommendation 7 encourages ESO to create a joint ALMA–VLT channel. In the current implementation, scientists applying for projects that require data from both facilities need to submit proposals to two different committees that are out of sync and are currently not coordinated. This can create discordant outcomes, with projects being judged very promising by one committee but rejected by the other.

Recommendation 8 proposes the introduction of a high-risk channel for the submission of high-risk/high-gain projects that require significant amounts of time. In the TAWG proposition, the high-risk channel would be kept separate from the DDT channel, and reviewed by an external board. Only large requests of time for very risky, but potentially highly rewarding proposals would be considered.

Recommendation 9 stems from the study of systematic effects in proposal selection carried out during the TAWG activities (Patat, 2016) and similar analyses performed at other large astronomical facilities (Reid, 2014; Lonsdale, Schwab & Hunt, 2016). Specifically, it is to consider limiting the level of information about the proposing team that is accessible to the reviewers (for example, the identity of the Principal Investigator [PI], the affiliation and the team composition). The TAWG acknowledges that this is potentially sensitive, and needs to be treated with care. Nevertheless, it recommends that ESO address this topic in coordination with other large scientific facilities worldwide.

The list of tasks assigned to the TAWG included the following item: “Examine the foreseeable evolution of the proposal selection and time allocation processes into the ELT era”. The TAWG advised that it is too soon to be providing recommendations on this matter without more information on the strategy envisaged for VLT operations in the era of the ELT. The TAWG would like to see the impact of the proposed changes to the VLT model before making any recommendations. It therefore preferred that ESO consider the development of a strategic plan for the use of the various ESO facilities at the start of ELT operations. When this is ready, a new working group should be constituted and tasked with the review of the recommendations presented in

the TAWG report, how they were implemented by ESO, and the lessons learnt, before looking at their application to the ELT.

Implementation status and outlook

The TAWG report was presented to the ESO Scientific Technical Committee and to the Users Committee, which both provided very detailed feedback about the implementation of the proposed changes. Following these discussions and after an internal analysis, ESO decided to take a gradual approach, in which each change is reversible and subject to revision. The first steps undertaken by ESO are as follows:

1. Move to a yearly cycle for the submission of LPs. This was announced in the Call for Proposals (CfP) for Period 102 and in science announcements issued in February, 2018^{3,4}. From Period 104 onwards, LPs will be allowed in only even semesters. Starting from that period LPs are required to span the minimum number of cycles compatible with the science case.
2. Encourage the submission of larger requests to improve the project diversity. This was solicited in the CfP for Period 102³. While it will probably take some time before the community reacts in a significant way to this call, ESO will continue to work actively to ensure that the distribution of requested time remains unchanged after the allocation and scheduling processes, guaranteeing that proposals of all lengths will have equal chance of success. For the time being there are no plans to implement a quota for short programmes (Recommendation 5).
3. Reduce the number of reviewers from six to three in the pre-meeting phase. This was experimentally introduced in Period 102, and the lessons learnt are being considered for future semesters. A poll to the OPC and Panel members for Period 102 is being prepared.
4. The obfuscation of information about the proposing team (Recommendation 9) has been thoroughly discussed, and the deployment of initial measures within the current proposal submission system has begun. As of Period 103, all information about the PI and co-investigator (col) affiliation (country,

institute, email address) will be removed from the version of the proposal distributed to the reviewers. The PI will be listed together with the cols, and the list displayed in pure alphabetical order on the last page of the proposal. More sophisticated measures (for example, different obfuscation levels for different proposal review phases), or more radical measures — see the recent change introduced at HST⁵ — will be considered for the new system.

The most substantial changes, i.e., the global move to a yearly cycle and the introduction of a fast-track channel (Recommendations 1 and 2) require structural modifications to the proposal handling software and underlying databases, and will therefore have to wait for the deployment of the new Phase 1 system. This particularly concerns the review of the FTC proposals, which involves automatic detection of conflicts and optimisation of the matching between proposals and reviewers. For the latter step — which will also be extended to the proposals in the regular cycle — a machine-learning approach is being investigated, along the lines described by Strolger et al. (2017).

The current implementation of proposal peer review at ESO involves a face-to-face meeting, which places strong constraints

on the number of reviewers. Although there are certainly good reasons for having such a meeting, including educational and social aspects, the limitations and costs involved are sufficiently important to warrant changes^a. Various approaches are being discussed, along lines that are also being considered for ALMA⁶. In the long term, other options are being considered, such as the abolishment of the panel meetings while keeping the OPC proper meeting for Large Programmes alone, and the deployment of distributed review for the remaining applications. Given the scale of proposal submissions at ESO (more than 700 distinct PIs in a semester), this would allow a very substantial increase in the statistical robustness of the review (by decreasing the effects of subjectivity; see Patat, 2018). This represents a major change, marking an epochal turning point in the 50+ years of the OPC tradition at ESO. It will therefore require thorough discussions with the governing bodies, the advisory committees and the community at large.

Acknowledgments

The author is grateful to all the members of the TAWG for their very constructive approach. The feedback provided by the Users Committee and the Science and Technical Committee is also acknowledged.

References

- Brinks, E. et al. 2012, *The Messenger*, 150, 20
- Lonsdale, C. J., Schwab, F. R. & Hunt, G. 2016, arXiv:1611.04795
- Patat, F. 2016, *The Messenger*, 165, 2
- Patat, F. et al. 2017a, *The Messenger*, 169, 5
- Patat, F. et al. 2017b, *The Messenger*, 170, 51
- Patat, F. 2018, *PASP*, 130, 4501
- Reid, I. N. 2014, *PASP*, 126, 923
- Sterzik, M. et al. 2016, *SPIE*, 9910, 03S
- Strolger, L.-G. et al. 2017, arXiv:1702.03324

Links

- ¹ The TAWG report to the ESO Users Committee: http://www.eso.org/public/about-eso/committees/uc/uc-41st/TAWG_REPORT.pdf
- ² The Fast Turnaround channel at Gemini Observatory: <http://www.gemini.edu/sciops/observing-gemini/proposal-routes-and-observing-modes/fast-turnaround>
- ³ ESO science announcement for the P102 Call for Proposals: <http://www.eso.org/sci/publications/announcements/sciann17098.html>
- ⁴ ESO P102 Call for Proposals document: <http://www.eso.org/sci/observing/phase1/p102/CfP102.pdf>
- ⁵ HST Cycle 26 Anonymous Proposal Reviews: <https://hst-docs.stsci.edu/display/HSP/HST+Cycle+26+Anonymous+Proposal+Reviews>
- ⁶ Report of the February 2018 ALMA ESAC F2F Meeting, 19 February 2018

Notes

- ^a The current logistical cost of the OPC and Panel meetings is approximately 240 000 euros per year. This includes travel, accommodation and full board for all participants.



The VLT at sunset.

G. Hudepohl (atacamaphoto.com)/ESO