TMT Science Operations Planning

Prepared by:
CASCA/ACURA TMT Advisory Committee (CATAC)

For:
Canadian TIO SAC Representatives
Canadian TIO Board Representatives
ACURA Board
CASCA Board

April 14, 2021

Introduction, assumptions and principles

The TMT Operations Plan was last updated in 2012, and there have been many important developments since then. For one, the construction delay has led to a different relationship between TMT and the other VLOTs: GMT and ELT. Secondly, the potential for NSF to engage with a 25% or greater share, in both the TMT and GMT, will also have an impact on operations planning. Finally, the social landscape has been changing, and Equity, Diversity and Inclusivity (EDI) concerns form an increasingly important role in funding, governance and scientific planning and execution. CATAC has therefore reviewed the latest Operations Plan Document (OPD, TMT.OPS.TEC.11.099.REL01, updated August 2012), as well as the Operations Requirements Document (ORD, TMT.OPS.MGT.07.002.CCR13, updated Dec 2018). This report summarizes our findings and makes several recommendations regarding Science Operations.

The European ELT project is now certainly going to see first light several years before TMT. The operations plan for ELT is the same as that for the VLT, with a dominant, adaptive queue mode component and a single Time Allocation Committee. While different from ESO in many ways, the TMT partnership shares some of the same challenges. In particular:

- There are significant language differences between TMT partner communities. Among other things, this could disadvantage proposals from non-native English speakers if TACs are not operated at the national level. ESO must also deal with this, and all ESO proposals are successfully reviewed and ranked by a single TAC.
- Just as is the case for ESO member states, there are significant differences in partner shares among the TMT partners. However, at ESO there is not a single dominant partner (Germany, the UK and France have the largest, and similar, shares). Whereas, if NSF joins at the 25% level, the US partners (NSF, Caltech and UC) will dominate the TIO partnership in a way that is not paralleled at ESO. The distribution of time pressure (oversubscription) among TMT partners may also differ from that among ESO members.

One aspect that significantly differentiates the TIO and ESO partnerships is that the US proposal to NSF is for US community access to both TMT and GMT, while none of the other TMT partners have access to GMT. Since NSF will hold the largest share, it is not unlikely that integrated or coordinated operations between GMT and TMT might drive the proposed operational model for TMT, with no immediate benefits to the non-US partners.

In forming our recommendations, we are guided by the following principles (in no particular order):

1. Maximizing science output for the Observatory as a whole, as well as for Canada specifically. The ORD quantifies this in terms of publications and citations.
2. Maximizing efficiency of telescope usage for the Observatory as a whole, as well as for Canadian astronomers specifically. By efficiency we consider primarily what the OPD refers to as Open Shutter Efficiency: “the fraction of time spent observing targets with the instrument shutter open as a percentage of Actual Science time (available minus weather and technical downtime).”
3. Providing flexibility; ensuring the Observatory and partnership have the opportunity to grow and change.

4. Encouraging inter-partner collaboration and cooperation, as a way to maximize science and also ensure a stable partnership.
5. Respecting partners' share, equitably to all observing conditions, instruments and modes.
6. Attention should be paid to Equity, Diversity and Inclusion (EDI) in all aspects of Operations
7. It is important that TMT provide training and development opportunities for students, postdocs and other highly qualified personnel (HQP)

We will refer to these principles by number, throughout this report. Most of these principles are already included in the ORD and OPD, with the notable exceptions of #6 and #7. If time allocation, for example, is left to individual partners, this could lead to telescope usage that reflects badly on the Observatory as a whole, in consideration of EDI. The Observatory should consider ways to mitigate this. For example, non-US time on the Astronomical Event Observatory Network (AEON) is allocated by individual (national or institutional) TACs. While each TAC is free to conduct the review any way they choose, AURA/NOIRLab provide guidelines and tools that each partner can use. These tools can be made to adhere to AURA/NOIRLab's own policies (e.g., regarding PI identification) if the partner so chooses.

Finally, we must consider operational practices that will be relevant in 2035 and beyond. During the past few decades there have been fundamental changes in the way that observatory operations are conducted. For example, the once popular classical observing model has been replaced at most major national facilities by queue-scheduled service observing. The degree of collaboration between astronomers in different countries has also increased, and there is a growing demand for conducting large programs, dedicated campaigns (sometimes across multiple facilities), and novel observing modes such as fast turn-around programs. We anticipate that these trends will continue, and the classical observing model will become even less popular than it is now. What is certain is that needs and expectations for operations will continue to change and evolve, leading up to construction and also throughout the lifetime of the Observatory. Appropriate statistics need to be collected, from TMT and other Observatories, to make informed decisions for continually updating the operations model (**Priority #3)**.

Recommendation #1: The ORD and OPD should be regularly revisited and updated as necessary, based on the most recent available information from TMT and other Observatories.

Recommendation #2: Specifically, the ORD and OPD should be updated to reflect the importance of EDI considerations in Operations planning. In particular, guidelines should be introduced to ensure that activities that may be left to individual partners, such as time allocation and software development, nonetheless conform to an Observatory standard in this respect.

Time Allocation

The twice-annual allocation of time to investigators on most ground-based telescopes is largely historical and not driven by scientific or efficiency considerations (**Principles #1 and #2)**. Most space-based telescopes allocate time on longer cycles, and with programs divided into categories depending on size and scope. Large and Long Programs (LLPs) have also been implemented on many ground-based telescopes, often later in their lifetime. At the other extreme, Gemini has introduced a Fast-turnaround (FT) process, with a monthly cycle that allows for more rapid science return. Recently, VLT has introduced a similar mode, called “fast-track” proposals. Canadians have made good use of these different modes. For Gemini, Canadian demand for LLP and FT time, relative to other partners, is comparable to its partner share (20%). Statistics from Gemini show that FT papers are published much more quickly than regular proposals, and that LLP papers have greater impact.

Much of forefront astronomy today is internationally collaborative, and draws upon multiple facilities. There are also likely to be efficiency and scientific benefits in encouraging inter-partner collaboration within TIO. Traditional, semester-based and partner-based time allocation processes can be a hindrance to this. In particular, the “double-jeopardy” proposers face when requiring proposal acceptance by multiple partners and/or telescopes, over multiple semesters, can both delay and compromise the science return. The fundamental assumption in the Operations Planning Document, that partners are responsible for allocating their own time and that a merged schedule is created every six months, has the potential to negatively impact our **Principles #1 - #4.** The double jeopardy faced by proposers can also impact **Principles #6 and #7.**

It is, of course, essential that TIO partners have access to TMT observations that are in line with their share, and with equitable access to all observing conditions (**Principle #5)**. In particular, tracking usage over suitably long time scales, relative to all relevant observing conditions, is important to ensure this principle is respected. Within that constraint, however, the Project should consider ways to make time allocation flexible, efficient and scientifically productive (**Principles #1 - #4)**. There is some scientific appeal of a single TAC model, as used by HST, JWST and ESO, for example. The valid concern is the impact this might have on partners with different shares, oversubscription pressures, and scientific interests. Care also needs to be taken so that proposers are not disadvantaged because they are non-native English speakers, or for other reasons that would contradict **Principle #6.** CATAC notes that ESO makes no effort to impose “quotas” on the amount of time each member country receives, yet the partner shares are still largely respected on average. For several reasons, including ones we list in the Introduction, this same model may not work naturally (i.e. without intervention) with TMT. However, with appropriate oversight, review and adjustment mechanisms, it is feasible to both rank proposals with a single TAC and still allocate time equitably. Hybrid models could also be considered where, for example, partners submit ranked lists of proposals evaluated by their national TACs, overfilling their allocated time by some factor (say, 30%). This subset would then be reviewed by a single TAC with representation from across the partnership for final relative ranking and allocation. In any case, it will be important to track collaboration-wide statistics regarding the proposed vs. accepted vs. available time, as well as (self-reported) statistics regarding the demographics of proposers, observers and archive users. This is important for our **Principles #1, #2, #6 and #7**.

The current Operations Plan explicitly notes that there will be no attempt to identify or remove requests for duplicate observations from different partners.   This violates our **Principles #1, #2, and #4, as well as possibly #6 and #7**. Proposals from different partners that are asking for identical setups on the same target(s) should be compared to each other, and duplications should be avoided. This could be done, for instance, by an "International TAC" (ITAC, as is the case for Gemini and CFHT). Some form of ITAC is likely to be necessary in any case when merging proposals from different partners.

Finally, CATAC recognizes that Director's Discretionary time is important for enabling exciting science opportunities. Additionally, DDT can be used to support small science investigations by the Observatory staff, ensuring that the staff remains engaged, invested, and directly involved in the Observatory operations (and supporting **Principles #1, #2, #4, #6, #7)**. The OPD describes a maximum of 30 hours per semester for this mode, which could also include time for staff astronomers. This seems both reasonable and important.

Recommendation #3: The infrastructure to support pan-partnership Large and Long Programs should be available at first light, even if it might be prudent to delay the first call for LLPs until the telescope/instrument performance is well characterized (e.g. one semester after first light). Even if regular time is allocated by National TACs, Large and Long Programs should be evaluated separately by a single, multi-partner, dedicated TAC. Such a TAC should be administered independently of any single partner, and with membership from all partners that is representative of their partner shares.

Recommendation #4: Fast Turnaround programs should be implemented at first light.

Recommendation #5: Proposal modes and time allocation should encourage and facilitate inter-partner collaboration. This should go beyond the proposal process described in REQ-1-OPSRD-4050 of the ORD.

Recommendation #6: There should be a capability to consider multi-facility proposals, including but not limited to other VLOTs. This is less urgent and need not be implemented at first light.

Recommendation #7: It is essential that time is allocated to partners according to their share, and spanning all relevant environmental factors (such as image quality and sky brightness), when averaged over a suitable time period (1-3 years). However, the Observatory should explore ways to maximize efficiency, science impact, collaboration and productivity in the time allocation process.

Recommendation #8: Even if time is allocated via national TACs, there must be a mechanism to identify and review proposals for duplicate observations. If competing teams are willing to collaborate and/or share data, this should be facilitated by the observatory.

Recommendation #9: The time tracking process (REQ-1-OPSRD-4010 in the ORD) should include more than the amount of bright, grey and dark time as described. Other factors, including image quality, water vapour content and sky transparency are relevant.

Recommendation #10: TIO should keep and publicly release anonymised, collaboration-wide statistics regarding the proposed vs. accepted vs. available time, as well as (self-reported) statistics regarding the demographics of proposers, observers and archive users. This information should be used to evolve and adapt TMT operations throughout the lifetime of the Observatory.

Observing modes and calibrations

We begin by summarizing and defining the observing modes considered for TMT, in the OPD and ORD:

- **Remote or physical presence observing**. Both modes are to be provided, though the normal mode of operation will be remote, either from headquarters or another location. Physical presence at the summit will only occur when it is essential (REQ-1-OPSRD-2055)
- **PI Directed (Classical).** During their assigned time, PI-directed classical observing users will have complete responsibility for how they use and configure the telescope and instruments and may modify their observing program during the night as they require (REQ-1-OPSRD-2050)
- **Service Observing.** This refers to execution of observing programs by Observatory staff. This can take two forms:
    - **Pre-planned Queue Service Observing:** A schedule is built from a combined list of observing blocks from all partners, in a given semester ([REQ-1-OPSRD-2605)
    - **Adaptive Queue Service Observing:** The queue is executed according to conditions. The Observatory shall implement adaptive queue scheduling, *within a single partner's allocated observing time*, if that is desired by the partner ([REQ-1-OPSRD-4005). The tools to do this are not part of the construction project and thus not planned to be available at first light. A fully adaptive queue, which dynamically adjusts to atmospheric conditions from a larger pool of partnership-wide observing programs will not be offered initially, however service queue processes and tools should not preclude this being offered later (REQ-1-OPSRD-2605).
- **Eavesdropping** (REQ-1-OPSRD-4020), **Target of Opportunity** (REQ-1-OPSRD-4035) and **synoptic/cadence** observations (REQ-1-OPSRD-4030) are provided.

Canadian astronomers have become accustomed to adaptive queue observing, with both Gemini and CFHT running primarily in this mode for many years. Many large observatories (including all space-based observatories) today operate at least partially in this mode; for instance VLT runs about 80% of its observing in “service” mode, and 20% in “visitor mode”, a model ESO is planning to adopt for the ELT. For this to be effective, however, the queue-scheduled time must be substantial enough to cover all environmental conditions and instrument combinations (**Principles #2 and #5**). Although a thorough experiment has never been performed, experience at Gemini, which operates almost exclusively in adaptive queue mode, suggests that removing more than 30% of observing from the queue starts to have a significant, negative impact on the queue efficiency. ESO and CFHT also estimate (although there is no hard data to back this claim) that queue efficiency would be negatively impacted if more than ~30-40% of the total available time were classically scheduled. For these reasons, an adaptive queue run within a limited block of time associated with just one partner is unlikely to be effective. Both ESO and Gemini report that the trend towards adaptive queue observing is “strong and relentless”; in the US, there are almost no requests for classical observing at Gemini. Moreover, a heavy reliance on classical mode observing can disadvantage many groups of people who are unable to travel, conflicting with our **Principles #6 and #7**.

Adaptive queue mode observing is demonstrably efficient (**Principle #2**) and scientifically productive (**Principle #1**). At Gemini, statistics now running over 30 semesters show that between 80% and 100% of highly ranked programs that run in adaptive queue mode reach 80% completeness or more. For programs running in block schedule (which approximates classical observing), the completeness fraction is essentially random and can be as low as 20%; this can be extremely detrimental, in particular for partners with small shares. Similar statistics have been reported for CFHT and ESO/VLT. Perhaps as a consequence, classical programs do not lead to more or faster publications (**Principle #1**). Adaptive queue mode observing also allows calibration data to be acquired and shared among programs in the most efficient way, maximizing open shutter efficiency (**Principles #2 and #4**).

Service observing might add somewhat to the operations cost since it requires a staff observer as well as a telescope operator to be present during observing. Additional, mostly front loaded, costs are needed to set up the service scheduling software, and these costs may be higher for a fully adaptive queue. However, ultimately these costs could be recovered through increased efficiency and science productivity. Service observing also reduces the opportunity for observatory staff to interact with the users, and for the users (in particular students and postdocs) to understand and feel ownership of their data. CATAC considers it very important for young astronomers to have ample opportunity to engage closely with the observing process (**Principle #7**). The opportunities for hands-on experience at Observatories are decreasing, as more telescopes move to service-mode observing (though, at the same time, remote observing capability may actually be improving access, for people who are unable to travel). These concerns can be mitigated through a hybrid mode, like Gemini's Priority Visitor mode. In this case, the PI visits the telescope (or remote observing location) during a time that is favourable for observing their program, but they execute the queue when conditions do not match those requested by their program. Observatory support (e.g. for students) is a good way to share the cost. "Eavesdropping" can also be an effective way for researchers to engage with the observations.

Our conclusion is that the large majority of TMT observations should be in a fully adaptive, cross-partner queue-mode. We also consider it important for a service mode to retain sufficient flexibility to allow PIs to engage with these observations, and also to adjust their programs in a way that maximizes science return, while remaining within the scope of the accepted proposal. This includes a mechanism for investigators to request modifications to their observing set up (including targets and instrument configuration) prior to execution; it is understood that all modifications must be well justified and approved by the observatory executive. When necessary, for instance in the case of ToO programs or programs targeting objects with unknown or uncertain properties, the Observatory must allow for program adjustments even as the program is being executed. For this, real-time coordination with the program PI/CoIs is essential and can be achieved, for instance, with the eavesdropping mode (REQ-1-OPSRD-4020).

Recommendation #11: The dominant mode of observing should be with adaptive queue scheduling. For this to be effective, it should represent a large majority of the allocated time, and be built from a merged list of programs from all partners. We urge the Project and partners to monitor the successes and failures of queue-run systems at other Observatories, including not only queue execution, but how the queue is filled, to identify best practices for use at the TMT.

Recommendation #12: The Observatory must allow PIs of accepted programs to request modifications, with the understanding that all changes to approved programs must be approved by the observatory executive. Under exceptional circumstances, changes should be allowed even while a program is being executed. There needs to be oversight to ensure that such modifications are within the scope of the original proposal, but this should not be onerous or a disincentive to users trying to maximize the impact of their observations.

Recommendation #13: Standard calibration data should be shared where possible among all programs and observing modes, not only service queue blocks as described in the OPD.

Data Archiving, planning and processing software

Both science and technical efficiency (open shutter efficiency) is improved with good software tools for planning, executing and reducing observations, and for accessing archival science and calibration data. Importantly, such tools cannot be simply created and left alone - they must be maintained, improved and adapted in order to remain useful. Moreover, in a diverse partnership like TIO, it is essential (**Principle #4**) that all partners have a role in specifying, monitoring, and benefiting from services that add substantial value to the Observatory.

The OPD describes a base level of support for these services, with extended capability provided, if at all, by the partners and instrument teams. There are three general areas of the OPD that were discussed by CATAC:

- **Proposal preparation and planning.** Only a preliminary “Phase I” tool will be provided as part of construction (REQ-1-OPSRD-2001). The OPD states that “the development of observing proposal and planning tools, exposure time calculators, and data processing software will be done in the first five years of operations by the Software and Information Technology (SWIT) Groups... Each instrument will be delivered with standard simulation software developed for the instrument performance modeling done during the instrument design and testing phases. TMT will take over and then maintain, improve and upgrade these software models.”
- **Data reduction and analysis tools.** Instrument teams will be responsible for delivering “standard data reduction software pipelines and observation simulators”. The OPD states that “TMT plans to provide and maintain these instrument pipelines running on observatory computers which can be accessed remotely by observers, along with disk space for data storage”.
- **Data archives.** While the Observatory will store all science and calibration data “indefinitely”, this will be in a database with modest search capabilities. The OPD notes that “Enhanced data archive features are not currently part of the baseline plan”.

As a concrete example, a recent development is that the NSF’s NOIRLab is developing the US ELT Program Platform (UPP), to support TMT and GMT use by the US community, through the full scientific data “life cycle”, from proposal to publication. This includes software tools and a data archive that are expected to be made available to all TIO partners.

CATAC is generally supportive of this model, where partners and instrument teams with the right expertise provide added value to the Observatory. It is potentially a good way to share cost, maximize productivity and efficiency, and build a collaborative partnership. However, CATAC has several specific concerns about the way this model is to be implemented, and the capabilities in the baseline plan.

- Partners should expect to be able to contribute to the Project in ways that draw on their expertise, interests and experience, and provide juste retour to their community, in approximate proportion to their share. Without careful coordination, leaving it entirely up to partners to develop tools compromises this, and can easily lead to an unhealthy competition. This is contrary to **Principle #4**.

- Tools developed by one partner will surely, and reasonably, be optimized for that partner's needs. As long as they are "good enough", and made generally available, other partners are unlikely to develop their own. This means in practice one partner having to accept the tools developed by another, even though they are not optimal, and this compromises **Principles #1, #4 and #6.**
- Reliance on a single partner to provide a critical component is risky. A change in priorities, funding levels, governance etc within that partner could significantly alter either development or support for these tools, upon which the other partners come to rely. Again this compromises **Principles #1, #4 and #6.**
- The absence of a searchable archive with high level functionality at first light is against **Principles #1 and #6.** An archive of this sort is fundamental to a modern observatory and must be part of the baseline plan.

These concerns can be mitigated, or avoided entirely, by having mechanisms for the partners to have input into defining, and setting the requirements for, tools that have the potential to benefit the whole partnership, as well as for monitoring their development. Any initiative to develop these tools needs to be resourced throughout the lifetime of the tool. Whose responsibility this is and where the resources come from should be agreed upon before any resources are allocated. In other words, software tools should be developed and deployed in analogy to instruments. While the OPD notes that "TMT will encourage and facilitate" collaboration between partners, and the sharing of expertise and techniques, it is CATAC's opinion that this is not enough. A steering committee, with participation from all partners, could be established to oversee software development that has the potential to benefit the whole partnership. Alternatively, permanent national/institutional offices could be designed to work closely together on these efforts. This model has had mixed success at Gemini, and it depends very much upon how these offices are established, staffed and run.

The OPD states that there will be an 18 month proprietary period on data obtained with TMT, but that partners may set shorter periods on their programs, if they wish. The 18 month period derives from the ORD REQ-1-OPSRD-2915, where it is listed as TBC. This seems reasonable to CATAC; certainly the default proprietary time should not be longer than 18 months, and an argument could be made to shorten it.

Recommendation #14: The Observatory should implement mechanisms to ensure that there is oversight, monitoring, and appropriate long-term maintenance of any software or data archive intended or expected for use by the broader TMT User community, even if developed within a single partner community. This could take the form of an advisory committee, drawn from across the partnership, that works closely with teams that are developing these tools.

Recommendation #15: A readily searchable and high-functioning archive, with equally good public access to non-proprietary data, is essential for maximizing science output and providing equal access to all members of the TMT community. This is especially true when the dominant observing mode is with an adaptive queue. Such an archive could be developed by a partner, but must be overseen and managed according to the previous recommendation.

Summary of Recommendations

1. The ORD and OPD should be regularly revisited and updated as necessary, based on the most recent available information from TMT and other Observatories.

2. Specifically, the ORD and OPD should be updated to reflect the importance of EDI considerations in Operations planning. In particular, guidelines should be introduced to ensure that activities that may be left to individual partners, such as time allocation and software development, nonetheless conform to an Observatory standard in this respect.

3. The infrastructure to support pan-partnership Large and Long Programs should be available at first light, even if it might be prudent to delay the first call for LLPs until the telescope/instrument performance is well characterized (e.g. one semester after first light). Even if regular time is allocated by National TACs, Large and Long Programs should be evaluated separately by a single, multi-partner, dedicated TAC. Such a TAC should be administered independently of any single partner, and with membership from all partners that is representative of their partner shares.

4. Fast Turnaround programs should be implemented at first light.

5. Proposal modes and time allocation should encourage and facilitate inter-partner collaboration. This should go beyond the proposal process described in REQ-1-OPSRD-4050 of the ORD.

6. There should be a capability to consider multi-facility proposals, including but not limited to other VLOTs. This is less urgent and need not be implemented at first light.

7. It is essential that time is allocated to partners according to their share, and spanning all relevant environmental factors (such as image quality and sky brightness), when averaged over a suitable time period (1-3 years). However, the Observatory should explore ways to maximize efficiency, science impact, collaboration and productivity in the time allocation process.

8. Even if time is allocated via national TACs, there must be a mechanism to identify and review proposals for duplicate observations. If competing teams are willing to collaborate and/or share data, this should be facilitated by the observatory.

9. The time tracking process (REQ-1-OPSRD-4010 in the ORD) should include more than the amount of bright, grey and dark time as described. Other factors, including image quality, water vapour content and sky transparency are relevant.

10. TIO should keep and publicly release anonymised, collaboration-wide statistics regarding the proposed vs. accepted vs. available time, as well as (self-reported) statistics regarding the demographics of proposers, observers and archive users.

This information should be used to evolve and adapt TMT operations throughout the lifetime of the Observatory.

11. The dominant mode of observing should be with adaptive queue scheduling. For this to be effective, it should represent a large majority of the allocated time, and be built from a merged list of programs from all partners. We urge the Project and partners to monitor the successes and failures of queue-run systems at other Observatories, including not only queue execution, but how the queue is filled, to identify best practices for use at the TMT.

12. The Observatory must allow PIs of accepted programs to request modifications, with the understanding that all changes to approved programs must be approved by the observatory executive. Under exceptional circumstances, changes should be allowed even while a program is being executed. There needs to be oversight to ensure that such modifications are within the scope of the original proposal, but this should not be onerous or a disincentive to users trying to maximize the impact of their observations.

13. Standard calibration data should be shared where possible among all programs and observing modes, not only service queue blocks as described in the OPD.

14. The Observatory should implement mechanisms to ensure that there is oversight, monitoring, and appropriate long-term maintenance of any software or data archive intended or expected for use by the broader TMT User community, even if developed within a single partner community. This could take the form of an advisory committee, drawn from across the partnership, that works closely with teams that are developing these tools.

15. A readily searchable and high-functioning archive, with equally good public access to non-proprietary data, is essential for maximizing science output and providing equal access to all members of the TMT community. This is especially true when the dominant observing mode is with an adaptive queue. Such an archive could be developed by a partner, but must be overseen and managed according to the previous recommendation.