HST Rapid Response Observations White Paper

John W. MacKenty
Last Revised: 18 July 2018

Executive Summary:
HST conducts ~30-40 observations per year with a rapid (<21 days) response to unplanned events. These are conducted either as Target of Opportunity (ToO) observations selected via the normal Time Allocation Committee process or as Directors Discretionary (DD) observations. This white paper summarizes the policies, processes, and recent history of such observations.

HST policy establishes a 21 day definition of a "disruptive" ToO observation to accommodate the complex implementation and scheduling process for an observation. The fastest possible turn-around time, with significant impact on the schedule and staff resources, is 36-48 hours. Due to the nature of ToO (and time critical DD observations), the implementation and scheduling of these observations requires timely activities by the PI and STScI Program Coordinators, Contact Scientists, Schedulers, Science Policy, and Management personnel.

An examination of the past five years of ToO and seven years of DD observations shows a trend towards increasing usage of shorter turn-around observations in both categories, which the system is handling successfully. In fact, even visits classified as "non-disruptive" most frequently are observed within 18 days of request and are usually scheduled early in the week. Very rapid observations (T<11 days) occur approximately every 6 weeks and arise mostly from TAC approved ToO observations.

Current practices and policy leave some opportunities to streamlining rapid observations. Four recommendations are provided in Section 5: (1) DD proposals requesting observations in less than three weeks should be submitted as Phase 2 proposals, (2) STScI Contact Scientists should be granted increased authority to modify the proposal without recourse to the PI to speed the implementation process, (3) rapid proposals should use SCHD=100 and allow a small pad at the each of each orbit to provide greater flexibility in scheduling, and (4) an advance determination by the PI should be made as the acceptability of scheduling with a single guide star.

We find that HST is currently providing the fastest realistic response time to triggering events. In principle, HST has the capacity to conduct rapid-response observations more frequently, but in practice, there would be a cost in observing efficiency and staffing resources.

1. Introduction

1.1 Rapid Observations with HST

While the great majority of HST observations are planned well in advance, our policies and scheduling system support observations requiring a relatively rapid response to an astronomical event. These observations may be approved via two different paths: Target of Opportunity (ToO) or Director's Discretionary (DD) observations. Within this white paper, we will refer to these collectively as "Rapid" observations. As discussed in Section 1.2, rapid observations may be considered "disruptive" or "non-disruptive" on the observation implementation and scheduling process. Here we consider "rapid" observations to include both when they result from a time critical request not placed in the long range plan.

1.2 Target of Opportunity Policies

Per the Cycle 26 HST Call for Proposals, ToO observations are defined as:

> *A target for HST observations is called a 'Target-of-Opportunity' (ToO) if the observations are linked to an event that may occur at an unknown time. ToO targets include objects that can be identified in advance but which undergo unpredictable changes (e.g., specific dwarf novae), as well as objects that can only be identified in advance as a class (e.g., novae, supernovae, gamma ray bursts, newly discovered comets, etc.). ToO Proposals must present a detailed plan for the observations to be performed if the triggering event occurs.*

The Call for Proposals defines two distinct classes of ToOs:

> *ToOs are therefore classified into two categories: disruptive ToOs that require observations on a rapid timescale and therefore revisions of HST observing schedules that are either active or in preparation; and non-disruptive ToOs that can be incorporated within the standard scheduling process. Disruptive ToOs are defined as those having turn-around times of less than three weeks. Non-disruptive ToOs have turn-around times longer than three weeks.*

The formal HST policy definition for disruptive ToOs is:

> *Disruptive ToOs: The minimum turn-around time for ToO activation is normally 2-5 days; this can be achieved only if all details of the proposal (except possibly the precise target position) are available in advance. Any required bright object screening (COS, STIS/MAMA, or ACS/SBC) must be completed before a ToO can be placed on the schedule. The ability to perform any bright-object check will depend on the quality of the flux information provided by the observer, the complexity of the field, and the availability of suitable expertise at STScI to evaluate that information on a short time scale. Under exceptional circumstances, it may be possible to achieve shorter turn-*

around times, but only at the expense of significant loss of observing efficiency. Ultra-rapid (<2 day turn-around) ToOs therefore require an extremely strong scientific justification, and may only be requested for instruments that do not require bright object checking (ACS/WFC, WFC3, STIS/CCD, FGS). Because of the significant effect disruptive ToO observations have on the HST schedule, the number of activations will be limited to eight in Cycle 26; this allocation will include no more than one Ultra-rapid ToO.

As will be explained in Section 2, these policies are based upon the HST scheduling process. The great majority of HST time is awarded via an annual TAC process. The selected proposal PIs then develop Phase 2 proposals with the observations grouped into "visits" of one or more orbits.

A Long Range Planning (LRP) process identifies times of the year during which each visit could be optimally scheduled (e.g. all constraints are satisfied and the observations are reasonably efficient) and attempts to create the best overall balance and efficiency for the entire HST observing program over the next 12+ months. The LRP is frequently updated and typically contains a significant "tail" of proposals that will be scheduled more than one year into the future. This provides both flexibility against schedule disruptions and failed observations and a more efficient overall program by somewhat oversubscribing the near-term schedule.

A detailed schedule is prepared weekly for actual execution on the spacecraft known as a Science Mission Specification (SMS). This schedule consists of visits from the LRP suitable for that week with attention to various priorities while attempting to maximize the efficiency for that week.

1.3 DD Policies

Many rapid observations are implemented via the TAC process (i.e. they are defined well in advance as ToO observations and are executed following an activation request from the PI and, if generic, accurate coordinates). However, a significant fraction (see Section 3) is received via the Director's Discretionary (DD) time process.

The Cycle 26 Call for Proposals contains three relevant statements regarding DD proposals:

A proposal for DD time might be appropriate in cases where an unexpected transient phenomenon occurs or when developments since the last proposal cycle make a time-critical observation necessary.

Recognizing the limited lifetimes for major space facilities such as HST and Chandra, DD Proposals for timely follow-up of new discoveries will also be considered even if the astrophysics of the phenomena do not require such rapid follow-up.

Weekly HST Command Loads are uplinked to the telescope on Sunday evenings; for nominal operations, the observing schedule is determined eleven days in advance of the uplink date. Although it is technically feasible to interrupt the schedule and initiate observations of a new target, short-notice interruptions place severe demands on the planning and scheduling process, decreasing overall observing efficiency and delaying other programs. Hence, requests for DD time must be submitted at least two months before the date of the requested observations, if possible. Requests for shorter turn-around times must be exceedingly well justified. In the case that a DD Program with a turn-around time of less than one month is accepted, the PI or his/her designee is required to be reachable by STScI personnel on a 24 hour basis between the submission and the implementation of the program, for Phase II preparation.

1.4 Motivations for this White Paper

Two concerns motivate the development of this white paper: (1) the expected decline in HST resources in the coming years, and (2) the growing scientific interest in time domain science. Current best estimates of the reliability of HST's hardware systems predict more than five years of future science observations. With flat, at best, budgets this implies a decrease of ~3 percent per year in staffing. By their very nature (as discussed below), ToO observations place significant burdens on HST operations and science staff members. Recognizing these costs – and identifying potential mitigations – is one motivation for this white paper.

The second motivation for this white paper is future science opportunities. HST's ability to respond rapidly has supported many of its outstanding scientific accomplishments including studies of supernovae, gamma ray bursts, and transient solar systems events (e.g. comets and impacts on Jupiter). With the recent emergence of gravitation wave astronomy as a source of transient targets and the growing numbers of ground based facilities devoted to time domain observations, it is likely that demand for HST observations of transient events will increase in the coming years. This white paper attempts to delineate the capabilities of HST to support such observations and to indicate which aspects of our policies and operations might be modified to increase these capabilities and mitigate resource limitations.

This white paper reviews our policies and processes for conducting Rapid observations with HST in Section 2. In Section 3 we examine the history of ToO observations for Cycles 20-24 and DD observations since January 2010. In Section 4 we consider opportunities for enhancing rapid response observations and in Section 5 list some potential opportunities for streamlining the implementation processes. For convenience our Conclusions are summarized in Section 6.

2. Proposal Preparation and Implementation Processes

2.1 Submission and Review Process

HST proposals are submitted via the annual TAC or the Director's Discretionary review processes (ToO programs are not permitted within the mid-Cycle proposal opportunities). In each instance, a Phase 1 proposal is provided for review and if accepted the PI submits a Phase 2 proposal containing the detailed description of the desired observations. A unique aspect of ToO observations submitted via the TAC process is that full details of the observations are not included in the Phase 2 submission. Missing information typically is target position but may also include exposure times, number of exposures, and various timing requirements. This necessitates interaction between the PI and STScI immediately prior to moving the proposal into a weekly schedule and thus ToO visits are never meaningfully included in the LRP.

Conclusion #1: ToO observations always require Phase 2 (re-)submission at the time of activation.

Our current review processes frequently require the involvement of several senior people often including the Director's office, SMO, HSTMO, and the people directly tasked with maintaining the LRP and building the weekly schedules. Because the implications for other programs are not easily known in advance, the decision to execute a Rapid observation may require iteration in the days (or hours) prior to submission of an SMS to GSFC. Thus, in some cases the final decision to execute a Rapid observation is not reached until its priority is assessed relative to other displaced programs. This may have significant consequences for the coordination of HST observations with other observatories.

Conclusion #2: Rapid observations frequently require resolution of competing priorities late in the process of generating the observing schedule.

2.2 Contact Scientist Reviews

HST proposals are reviewed by science staff members (Contact Scientists) in INS to help assure the success of the planned observations. While the ultimate responsibility for the scientific outcome resides with the PI, these CS reviews are often both a resource for the PI and a valuable check on the Phase 2 proposal design. In the case of the ACS SBC, COS FUV and NUV, and STIS MAMA channels, these CS checks are the primary protection against overlight of their sensitive high-voltage detectors. As the consequences of overlight events range from temporarily suspending instrument operations up to irreversible damage to the detectors, STScI retains responsibility for checking every such observation during the CS review process.

Conclusion #3: Timely Contact Scientist review of Rapid observations is required for all proposals.

Conclusion #4: Successful and timely CS reviews, especially for proposals using the high-voltage detectors, require that the PI provide sufficient information or respond quickly to requests for clarification or additional information.

2.3 Program Coordinator Activities

The Program Coordinator (PC) has the overall responsibility for preparing and optimizing the Phase 2 proposal while working closely with the PI and CS. A key aspect of this process is to avoid iterations with the PI or the scheduler at the last minute.

However, as noted above, final work on preparing the proposal for implementation cannot start until the updated Phase 2 is received. A critical next step is the identification of suitable guide stars. If acceptable guide stars exist, this step requires 1-2 hours of effort. Otherwise, iteration with the PI may be required. Such iteration may change orient restrictions, move target within the field of view of an imager, or accept single guide star observations. The PC also has the option to use FGS3 (a limited lifetime resource for HST), to expand the FGS search radii, or to accept a higher risk guide star.

Conclusion #5: Guide star selection is fundamentally a risk activity during the activation of a Rapid observation that may either delay the workflow or result in the observation being un-schedulable within the desired timeframe.

2.4 Long Range Planning Process

Most HST observations are selected via the TAC process and optimized within the Long Range Planning (LRP) process. This ensures that the number of scientifically productive HST orbits each year is maximized and that most observations are done using orbits having as long a visibility period as reasonably possible. Equally important, beyond merely maximizing the available observing time, the LRP process enables the implementation of multiple scientifically necessary scheduling requirements. These include synoptic or sequenced observations, observations of time-critical events (e.g. the transit of an exoplanet), observations required to place targets at particular orientations relative to the HST field of view, and coordination with other observatories (both ground and space). The scheduling process also must accommodate various limitations of the science instruments including avoidance of the SAA, persistence effects in the WFC3/IR detectors, and occasional periods of non-availability of certain instruments (e.g. during CCD detector anneals).

The optimization process proceeds in a layered fashion by identifying weeks of the year suitable for certain observations and then optimizing each week (during the SMS generation process – see Section 2.5). This creates a natural hierarchy of priorities for each visit with priority tending to go to those visits who's scheduling opportunities are either rare or which will not recur for many months or years.

A key aspect of this activity is that each week's scheduler has a significant oversubscription of visits (to achieve greater schedule efficiency), clearly knows from the LRP which visits are "must go", has good communications to know which visits failed to schedule from the prior week, and good visibility into the consequences of not scheduling a particular visit in that week.

In recent years, a significant (and increasing) fraction of the LRP consists of highly constrained visits. This is due in large part to evolution of the scientific priorities of the HST user community including the large number of exoplanet transit observations.

A consequence of the nature of the current HST LRP is that Rapid observations frequently displace a planned relatively high priority program. Depending upon the constraints and duration of the Rapid visit, the displaced program may be scheduled in an SMS in the following weeks or may be delayed significantly longer. This often becomes an iterative decision process with multiple actors as discussed in Section 2.1.

The LRP process is a continuous process during the year as newly approved, displaced, or repeat observations are placed into the LRP.

Conclusion #6: ToO's that are relatively short and are less disruptive (more lead time and less specific timing constraints) impose a smaller science cost on the LRP. Conversely, a multi-orbit disruptive ToO with demanding timing constraints in effect uses far more orbits than its face value and often creates significant disruptions to the LRP since its placement in the SMS offers fewer opportunities to accommodate other programs.

Conclusion #7: Although performing the LRP process over the entire cycle depends upon the SPIKE software to accomplish a global optimization of the usage of HST, there is significant manual intervention on time scales of a few weeks to balance competing priorities while simultaneously optimizing the schedule.

2.5 SMS Generation Work Flow

Once a proposal (or more precisely a visit within a proposal) is ready for scheduling, it moves into the SMS building process. Normal HST flow completes the construction of an SMS on a Thursday for transmission to GSFC. The execution of this seven-day SMS usually starts on Sunday evening (0 hours UT).

Figure 1: SMS Build Timeline

SMS Build Timeline:

T-17 (Friday): Generic TDRSS contacts request
T-14-10 (Mon-Fri): NCC resolves TDRS schedule conflicts w/ other missions
T-13-12 (Tues-Wed): LRP updated based upon results of prior weeks SMS
T-12 (Wed): LRP provides pool of visits for next SMS schedule
T-11-10 (Thr-Fri): SMS schedule built (nominally completes)
T-7 (Mon): TDRSS contacts set (may result in SMS mods); PASS processing starts
T-6-4 (Tue-Thr): PASS processing continues; final PASS run(s) completed
T-4 (Thr): SMS sent to GSFC
T-1 (Sunday): uplink first command loads to HST

This illustrates a key point that the SMS cannot be built in isolation since what is on the prior week's SMS impact its contents (linked visits, visits done a week early for efficiency, visits delayed that could not be scheduled, etc.). This is generally not final until Thursday of the prior week thus much of the "packing" of the SMS occurs on Friday and Monday.

Ground communications with HST both up and down are performed via TDRSS contacts. Two major constraints exist. First, since HST's ephemeris cannot be precisely known many weeks in advance, final TDRSS contacts are identified one week prior to an SMS start. Second, since TDRSS is a shared resource with many customers, scheduling opportunities are not under HST's control and must be negotiated with TDRSS command NCC at White Sands, New Mexico.

Normally the SMS is built within the expected TDRSS constraints available at the start of the SMS build process (typically 11 days prior to SMS start). Late changes to an SMS due to a Rapid observation may necessitate renegotiation of TDRSS contacts. While this is usually successful, it adds work and some degree of uncertainty. The TDRSS schedule is the origin of the Sunday evening (EST; midnight UT) SMS start boundary.

The SMS is organized into eight-hour command loads. It is not unusual for ground operations to be 24 hours ahead of the timeline in sending these command loads up to the spacecraft (although unexecuted command loads may be overwritten). Currently, any changes to the executing program need to happen at a command load boundary.

Conclusion #8: HST's operational architecture is fundamentally tied to TDRSS operations making major redesign of the scheduling and command load process impractical.

2.6 Disruptive versus non-Disruptive Observations

Due to the workflow in constructing the SMS described above, Rapid observations activated less than or equal to 21 days prior to their desired execution are considered "disruptive" by policy. Although as shown in Section 3.1, this appears to be a rather conservative definition as it allows the 11 days for SMS generation, execution anytime within the 7 day SMS, and 3 days for contingency. The degree of disruption rises as the time window shortens. This applies both to the interval prior to the start of the SMS and to the need to schedule the observation earlier within the seven day SMS. In practice, Rapid observations activated more than ~11 days prior to the start of the SMS on which they will execute are generally "non-disruptive" of the SMS building process itself (although they certainly impact the LRP).

Activation between 4 and 11 days (relative to the start of the week's SMS) is disruptive of the SMS build process at STScI. This results in additional work at STScI (e.g. re-building a nearly finished SMS, re-planning TDRSS contacts) and may result in additional work at GSFC (e.g. due to late delivery of the SMS weekend staffing may be required).

If the activation is less than ~4 days, then it is likely that the currently executing SMS will require an intercept SMS. In this case, a replacement SMS covering the remainder of the week is built by STScI and delivered to GSFC. The switchover must happen at an 8-hour command load boundary. In some cases, this might require off-hours work at both STScI and GSFC to accomplish.
Conclusion #9: Disruptive observations (either via DD request or ToO activation) represent significant additional work.

2.7 Fundamental Limits to ToO Response Time

Within the current HST ground system architecture, the fastest possible response to a decision to conduct a Rapid observation is as follows:

a) Construct a suitable APT Phase 2 submission (starting point)
b) PC and CS reviews of APT submission; find Guide Stars (2 hours)
c) Build intercept SMS and delivery to GSFC (10-12 hours)
d) Review at GSFC (4 hours)
e) Uplink to HST (4 hours)

Given the 8-hour quantization of command loads, the implied 24 hours is rather optimistic and 36 hours is probably a reasonable lower limit for Rapid response starting from an approved Phase 2 submission. Shortening this would require major changes to current legacy software systems used in proposal generation and validation.

Another unchangeable limitation arises from the use of TDRSS for HST communications. If HST is placed into a different attitude than anticipated by the TDRSS contacts schedule, then changes to the TDRSS contact schedule will need to be renegotiated with TDRSS command at White Sands. Given the need for fairly frequent downlink opportunities to avoid overflowing the on-board data storage, this creates a timing limitation in building and uplinking an intercept SMS as this is a manual process between TDRSS command and STScI.

Finally, HST operations are only staffed for 5 day/8 hour shifts and Rapid observation execution at the shortest timescales requires overtime work and optimized alignment with staffed shifts (or requires GSFC staff to come in off-hours).

A timeline of a successful rapid ToO is provided as Appendix A.

Conclusion#10: ToO response time cannot be shortened to less than 24 hours at the extreme and more realistically 36-48 hours is the shortest turn around possible for HST without a major redesign of most of the ground system.

3.0 History of ToO and DD observations

3.1 Past Target of Opportunity observations

A history of TAC approved ToO's since Cycle 20 is available on the HST Metrics page (http://www.stsci.edu/hst/metrics/TOO/Data/20180108too.txt/preWrap). Table 1 summarizes the history of these observations by Cycle. While the Cycle25 numbers are included, in Table 1, they are not used in the subsequent analysis as Cycle 25 is ongoing.

Table 1: Cycle 20-25 ToO Statistics

Cycle	Proposals	Disrupt Proposals	Disrupt Visits	Disrupt Visits Done	Non-Dis Visits	Non-Dis Visits Done	Prop w/ No data
25	21	13	16	1	27	4	17
24	18	10	19	9	36	30	4
23	24	12	13	7	63	52	8
22	18	9	12	6	44	37	8
21	15	7	10	4	19	15	6
20	16	6	6	4	44	37	4
Cycles 20-24	91	44	60	30	206	171	30

Table 1 shows the number of ToO proposals (not including DD) approved for each Cycle. The numbers of Disruptive and Non-Disruptive visits approved (under the 21 day rule), and the number of visits observed to date (essentially all proposals consist solely of either Disruptive or Non-Disruptive, not both). Also included are the numbers of proposals that have not obtained data. [Note that in Cycle 20 the policy threshold for a disruptive proposal was 14 days which biases this statistics slightly.]

Notice that of the 91 approved proposals in Cycles 20-24, 30 did not result in any executed visits (33%). Since some ToO proposals are granted long-term status to account for low probability of occurrence events, it is possible that the Cycle 24 statistics will change slightly but this basic trend will persist. This should not be regarded as a negative aspect of the ToO proposal pool since there are clear examples of "contingency" proposals (e.g. cometary impacts on Jupiter) whose presence in the pool would clearly facilitate the timely acquisition of data that prior experience has proven valuable.

<u>Conclusion #11: The number of Disruptive Proposals and Disruptive Visits approved by the TAC has increased in recent Cycles.</u>

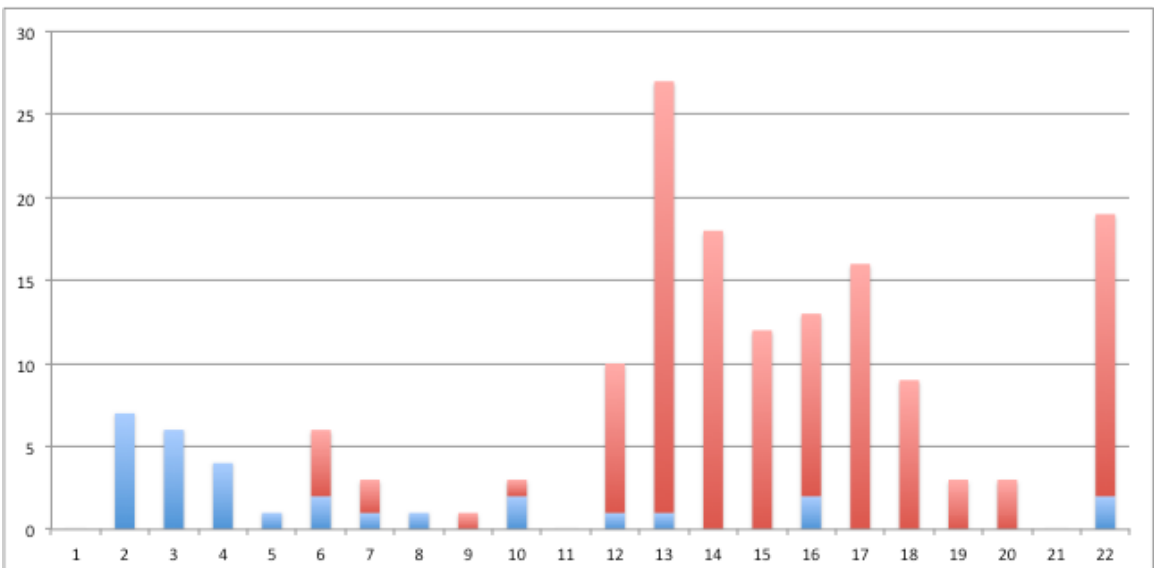

Figure 2: Days (X-axis) between PI Activation and Visit Execution. Blue shows numbers of Disruptive visits and Red shows numbers of visits approved as Non-Disruptive.

Figure 2 shows the number of days between the PI's activation request and the observation for executed ToO visits in Cycles 20-24. The 22 Day bin includes all cases >21 days. Blue shows visits classified as Disruptive and Red visits approved as Non-Disruptive.

Two caveats apply to this figure: (1) we exclude 39 visits from proposals 13677 and 14808 where multiple targets in the same fields were activated simultaneously (significantly reducing the workload and degree of disruption) – these cluster between 12 and 16 days, and (2) where a sequence of visits resulted from an activation request, only the timing of the first visit is shown. In both situations, Figure 2 remains reasonably reflective of the impact on the HST schedule and work.

Clearly the HST scheduling process is able to support fast response (2-3 days) and also accomplishes a considerable majority of the non-disruptive ToO observations within 18 days.

Conclusion #12: Disruptive ToO are frequently observed within 2-4 days following an activation request.

The overall success in scheduling ToO observations is reflected in the fact that of 162 requests for activation in Cycles 20-24, the first visit was scheduled within 21 days 143 times (88%). If proposals 13677 and 14808 were included, the rate increases to 91%.

Conclusion #13: Although policy defines Non-Disruptive ToO activations requiring >21 days, most are actually observed within~18 days.

One other interesting behavior is the frequency of events by day of the week. Figure 3 shows the numbers of ToO activations (Blue = Disruptive and Red = Non-Disruptive) as a function of the day of the week. Figure 4 shows the same for their execution dates. Note that execution times are in UT while activation is EST/EDT.

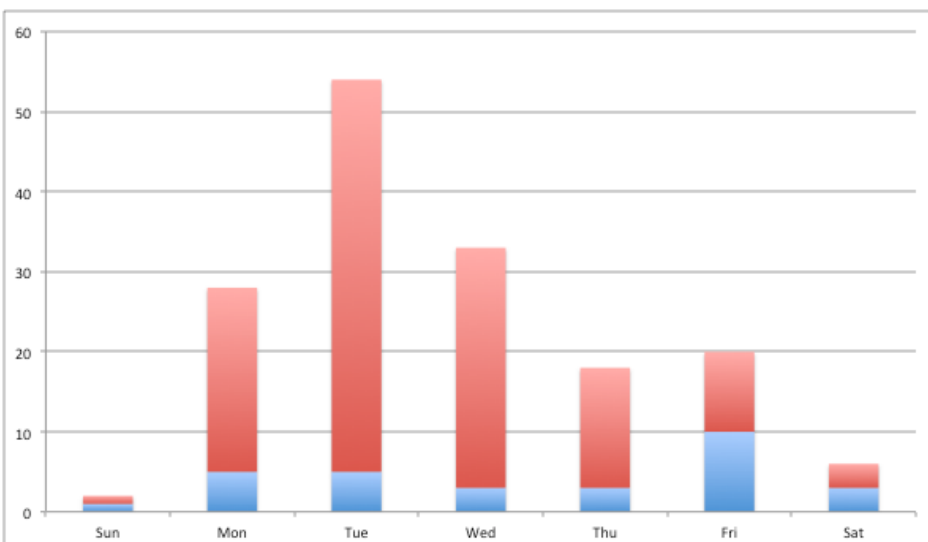

Figure 3: ToO Activation Day for Cycles 20-24 (Blue = Disruptive, Red = Non-Disruptive)

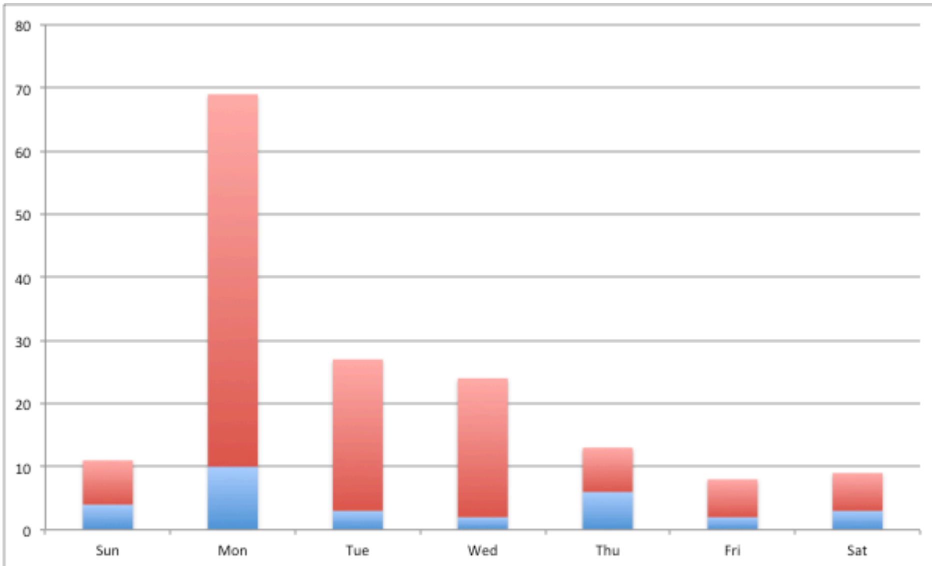

Figure 4: ToO Execution Day for Cycles 20-24 (Blue = Disruptive, Red = Non-Disruptive)

Not surprisingly, there are fewer activations on the weekends. However, there is a strong tendency for ToO observations to be executed on Monday (i.e. early in the SMS). A significant factor in explaining this trend is that re-building and re-delivering an SMS is usually preferable to intercepting the executing SMS. Thus if an activation has a observer provided "window" that permits the observation to occur on Monday rather than Sunday, the less disruptive and expensive solution will be adopted. Furthermore, PIs and PCs know that having a non-disruptive visit ready on a Wednesday can be scheduled in 12 days but on Thursday it would be scheduled in 2.5 weeks.

Conclusion #14: The STScI scheduling process tends to place ToO observations early in an SMS. This reflects the ability and dedication of the schedulers.

3.2 Past Director's Discretionary observations

A history of past DD observations is available at:
http://www.stsci.edu/hst/metrics/Discretionary/Data/20180108dd2.txt/preWrap
(n.b. While there do appear to be a few minor inconsistences in these records, they do not change the overall statistics or conclusions).

We find 121 DD proposals consisting of 411 executed visits that were submitted and approved between January 1, 2010 and January 5, 2018. These are summarized in Table 2. These statistics reflect the time to execute the first observation in a proposal. Many proposals contain multiple visits (often with specified cadence) but for the purpose of this analysis only the first visit executed is considered "rapid." This somewhat understates the degree of impact on subsequent weeks but as these have less (but certainly non-zero) impact they are ignored in this study.

Table 2: Past usage of DD proposals

Year	# Proposals	Accept to Exec <=11 Days	Accept to Exec <=21 Days
2010	9	3	5
2011	14	4	9
2012	14	0	2
2013	15	3	10
2014	12	1	3
2015	10	0	5
2016	21	3	15
2017	26	3	12
Total	121	17 (14%)	61 (50%)

As expected, there is a broad range of times between submission and execution as not all of these observations were time critical (see Figure 5).

Conclusion #15: The number of DD proposals and the number of DD proposals that meet the policy definition of "Disruptive" (i.e. <= 21 days) has increased in recent years although the numbers of significantly disruptive (T<=11 days) has remains constant and is relatively infrequent.

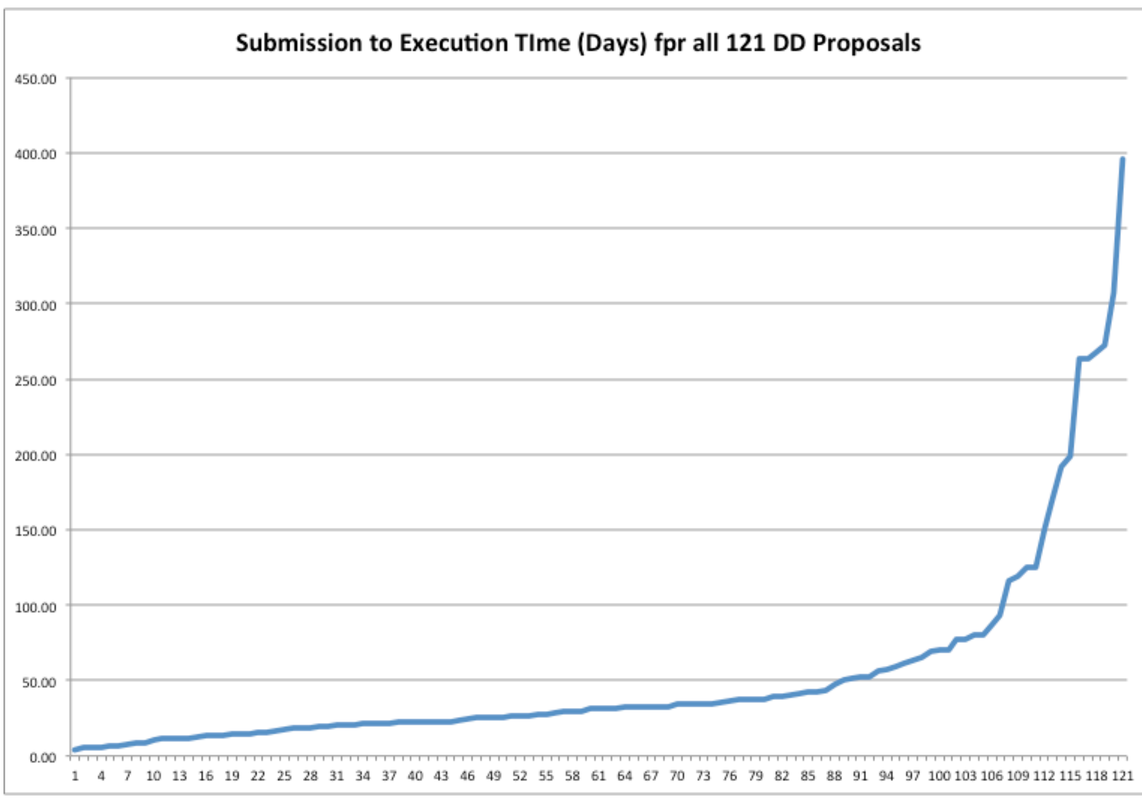

Figure 5: Submission to Execution Time for 121 DD Proposals 2010-2017. Here the 121 proposals (X-axis) are sorted by execution time in Days (shown on the Y-axis).

As seen in Figure 6, the time from submission to acceptance of a DD proposal has a median time of 6 days and is considerably faster for time critical proposals. From the total of 121 approved DD proposals, only 4 took longer than 30 days prior to acceptance.

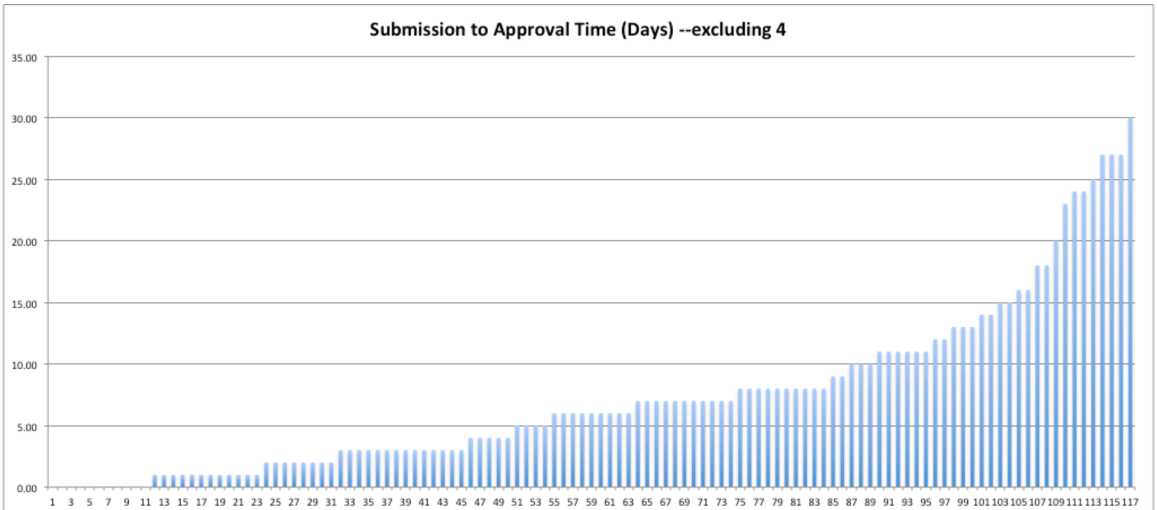

Figure 6: Submission to Approval Time for 117 DD Proposals. Here 117 proposals are sorted by time in Days between submission and approval (shown on Y axis).

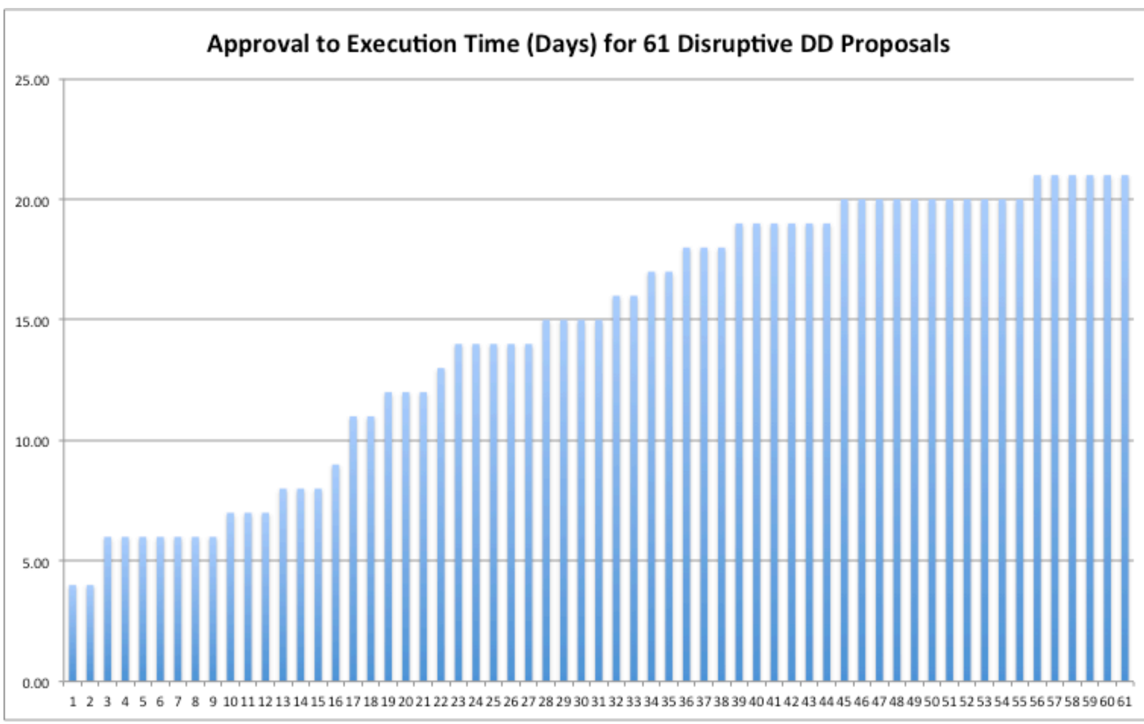

Figure 7: Approval to Execution TIme (Days) for the 61 Disruptive DD Proposals. Time in Days in Y-axis.

Using the 21-day policy for disruptive proposals post-facto, 61 of the DD proposals executed their first visit within 21 days of approval. Their distribution is shown in Figure 7. The median timescale is about 15 days with only 12 <= 7 days.

Conclusion #16: The DD approval process is generally faster than the implementation time (approval to execution) and thus does not constitute the rate determining step.

Table 3 shows the Submission and Execution Day of the Week for the 121 DD proposals. As expected, PIs rarely submit on weekends and there is a small (1.5 sigma) tendency to receive more DD proposals on Fridays. Unlike ToO proposals, there is not a strong bias towards execution on one day of the week.

Table 3: Day of the Week for submission and execution of DD proposals

	Sun	Mon	Tue	Wed	Thu	Fri	Sat
Submit	6	20	23	23	19	27	3
Execute	7	7	30	24	20	21	12

3.3 Comparison of ToO and DD rapid proposals

Three general observations regarding our actual experience with both TAC approved ToOs and short turn-around DD programs (both disruptive and non-disruptive) are worth making.

First, our staff is very dedicated to supporting these proposals and typically exerts considerable "above and beyond" efforts in assuring their success. This clearly provides pride of accomplishment, a strong sense of mission, and keeps them directly engaged with the user community.

Second, the increase in scheduling constraints with the growth of exoplanet research (and some planetary mission support observations) in recent years has made our short and long term schedules more fragile. Thus the effort required to insert ToO observations (both in schedule complexity and the implications of delaying displaced observations) has been increasing.

Third, the numbers of ToO and DD proposals which disrupt both the LRP and the SMS generation workflow has been increasing in recent years.

Our past experience (Cycles 20-24 for ToO) and CY 2010-2017 for DD) provide a look at the relative usage of these policy options for obtaining rapid turnaround observations with HST. We find:

Conclusion #17: ToO observations are more effective in providing very short (T<5 day) observations than DD observations.

Note that DD observations are not classified as disruptive vs non-disruptive but are scheduled in accordance with the scientific requirements as exist when they are accepted (e.g. they are often of the "as soon as possible" or "within the next N days").

Conclusion #18: Truly rapid observations with schedule impact (T<11 days) occur 2.1 times per year from DD and 6.4 times per year from ToO observations. This implies an average frequency of one major disruption every 6 weeks.

Conclusion #19: We are executing rapid observations (i.e. within 21 days whether or not intended as disruptive) 7.65 times per year from DD and 28.6 (conservatively) times per year from ToO activations. This implies an average frequency of one response every 10 days.

Conclusion #20: Both the overall frequency and the frequency of shorter response time observations are increasing in recent years.

Conclusion #21: The majority of the DD proposals which meet the 21 Day Disruptive criteria execute with comparable delays to the Non-Disruptive ToO proposals.

Three caveats apply to these numbers. First, a few of the very fastest turnaround activations were observations of the same target but this was a minor effect. Second, the follow-up of supernovae in search fields resulted in a single activation request for multiple targets –these are counted as only a single activation for this study. Third, some activations result in multiple visits in a short period of time (e.g. days or weeks) that produce a disproportional impact on the schedule.

4. Opportunities for Enhancing Rapid Observations

Rapid observations have proven to be scientifically successful for HST, and range from solar system observations (e.g. cometary impacts on Jupiter) to SN observations for cosmological studies. As thus there is continuing (and probably increasing) demand for such observations, we consider three categories of enhancements to the present policies and capabilities: response speed, number of accepted Rapid proposals, and complexity of Rapid proposals.

4.1 Timescale for Response

The total number of Rapid observations that disrupt the SMS development process is approximately one every two weeks. The degree of disruption varies depending upon the timescale for executing the observation, the phasing within the scheduling flow, the readiness of the proposal for execution, and the other programs on that SMS. To be considered "disruptive" requires an activation timescale <21 days. In reality, if a proposal is ready on Monday for execution the following week (i.e. the PC implementation work was completed the prior week), its degree of disruption is mainly on the LRP and not the workloads of the implementation staff.

However, requests received on Friday create the potential for off hours (evenings and weekend) work for STScI PCs and CSs if the necessary work cannot be completed during normal working hours.

Conclusion #22: Proposals asking for more than 7-8 days from activation and that are ready for scheduling on a Monday are significantly less disruptive than shorter timescales. This reflects the historical success of the scheduling team in placing ToO observations on the first or second day of a SMS.

Conversely, activation requests received after the SMS is generated and sent to GSFC (i.e. Thursday) and that expect observations prior to the subsequent SMS (e.g. about 11-12 days) are highly disruptive as the SMS will need to be re-generated and re-delivered or an intercept SMS will be required. Here the consequences are enhanced workloads, considerable susceptibility to problems in proposal implementations, and a higher likelihood of a larger impact on SMS efficiency. As discussed in Sections 2.5 and 2.6, the interactions between HST and TDRSS plus the quantization of HST command loads pose hard limits on the shortest possible activation at about 36-48 hours.

Conclusion #23: Implementation costs increase significantly for activations less than 11 (7-8 days if provided on Monday) days although existing processes can handle such proposals successfully.

We have considered the option of planning in advance for a ToO and then executing very rapidly. In order to work within the constraints of the HST to TDRSS communications systems, the position of the target would need to be known to approximately 1 degree in advance. Overcoming the current eight-hour command block boundaries might be possible with significant investment. Under the best circumstances, one might imagine a ToO for a target within a 1 degree region with a latency time of ~12 hours using a pre-planned block of HST time (i.e. doing nothing else with that time).

Conclusion #24: Hard limits preclude the HST plus TDRSS system from executing anthing other than a ultra-constrained type of ToO in less than 24 hours and 36-48 hours is a practical limit. This ultra-constrained ToO would need to reside within a 1-degree region on the sky specified in advance.

4.2 Frequency of Rapid observations

The current limitations on the frequency of Rapid observations are not hard limits but rather represent a choice involving staffing levels, scheduling priorities, and efficiency in the use of HST orbits. The current frequency of disruptive ToO and DD rapid observations is sustainable and could be increased if scientifically desired. A number of the recommendations in Section 5 are designed to decrease the staffing costs of accommodating Rapid observations. If the process is not streamlined, then the cost of increasing the frequency of rapid observations would be a need to increase STScI staffing in the impacted teams.

Finding the right balance between highly linked programs on the LRP and disruptive Rapid observation is a scientific choice that may be revisited each cycle.

If the frequency of astronomical events requesting Rapid observations becomes much larger (e.g. in the LSST era), it may be worth considering grouping such observations. Assigning certain weeks in the LRP in advance for Rapids would potentially decrease the impact on regularly scheduled programs and would permit optimization of the staffing required to create Rapid observation intensive SMS's.

Conclusion #25: The present frequency of ToO observations is sustainable and could be significantly increased if desired (although process streamlining and/or increased staffing would be required).

4.3 Complexity of ToO observations

The implementation effort, probability of success, timescale, and impact on other planned observations are dependent upon the complexity of the desired observation. Basically short, simple, unconstrained observations are easier.

Key factors include:

1) Science Instruments: The instruments with high voltage detectors (ACS/SBC, COS FUV/NUV, and STIS MAMA) require that each observation be cleared to avoid detector overlight. This results in delay, requires availability of key staff members, and may require revision of the proposed observations.
2) Timing constraints and linkages: Tight timing windows both for the initial and any subsequent observations may increase the impact of the rapid observation on other planned observations or delay its implementation if modifications are required (e.g. to avoid the SAA).
3) Length of the Rapid observation: Visits with fewer orbits are both easier to schedule and have less impact on the remainder of the SMS. Observations that do not use the entire orbit (even by only 5-10 minutes) provide considerably more flexibility for scheduling and the inclusion of necessary spacecraft activities.

Conclusion #26: Simple observations using CCD or IR detectors that do not fully use the visibility period are more likely to be rapidly implemented and have less impact on other observations and staff resources.

5. Recommendations to Streamline Rapid Observations

The following recommendations aim to reduce the workloads on STScI staff, reduce the impacts on the schedule efficiency, and maintain both TAC and DD opportunities for Rapid observations. Some of these recommendations impose some tradeoffs or additional work for the PIs.

Recommendation #1: Time critical (e.g. <3 weeks) DD proposals should be delivered as Phase 2 proposals. Rationale: Reduces iterations with PI and permits assessment at approval stage of magnitude of impact upon other programs. Reflects the observation that historically ToO observations schedule more rapidly than DD proposals (Conclusions #17).

Recommendation #2: An STScI Scientist (e.g. CS) should be delegated authority to iterate the proposal with the PC without recourse to PI for disruptive ToOs and DDs. Rationale: If we make the commitment to disrupt the schedule and workflows, priority should be shifted to getting a reasonable executable proposal onto the SMS rather than perfection. Moving from our current posture that ToO orbits are "owned" by the PI and that that PI can iterate seeking perfection will reduce implementation cost and timescale.

Recommendation #3: Disruptive ToOs and DDs should be implemented using SCHD=100 plus providing a few minutes of unused time at the end of each orbit. Rationale: Excessive orbit packing is the easiest way to induce re-work or to force the ToO/DD visit to unnecessarily displace other high priority orbits. Additional time in the visibility period may provide the scheduler opportunities to "salvage" an opportunity.

Recommendation #4: Acceptability to the PI of single Guide Star acquisition and guiding if necessary and appropriate shall be established at the time of proposal submission. If single guide star observations were not so approved, failure to find a suitable guide star pair will stop all work on that visit. Rationale: Since GS selection must be performed late in the implementation process, this decision should be made earlier and fixed.

6. Summary of Conclusions Developed in Sections 2-4

1. ToO observations always require Phase 2 (re-)submission at the time of activation.
2. Rapid observations frequently require resolution of competing priorities late in the process of generating the observing schedule.
3. Timely Contact Scientist review of Rapid observations is required for all proposals.
4. Successful and timely CS reviews, especially for proposals using the high-voltage detectors, require that the PI provide sufficient information or respond quickly to requests for clarification or additional information.
5. Guide star selection is fundamentally a risk activity during the activation of a Rapid observation that may either delay the workflow or result in the observation being un-schedulable within the desired timeframe.
6. ToO's that are relatively short and are less disruptive (more lead time and less specific timing constraints) impose a smaller science cost on the LRP.

Conversely, a multi-orbit disruptive ToO with demanding timing constraints in effect uses far more orbits than its face value and often creates significant disruptions to the LRP since its placement in the SMS offers fewer opportunities to accommodate other programs.

7. Although performing the LRP process over the entire cycle depends upon the SPIKE software to accomplish a global optimization of the usage of HST, there is significant manual intervention on time scales of a few weeks to balance competing priorities while simultaneously optimizing the schedule.
8. HST's operational architecture is fundamentally tied to TDRSS operations making major redesign of the scheduling and command load process impractical.
9. Disruptive observations (either via DD request or ToO activation) represent significant additional work.
10. ToO response time cannot be shortened to less than 24 hours at the extreme and more realistically 36-48 hours is the shortest turn around possible for HST without a major redesign of most of the ground system.
11. The number of Disruptive Proposals and Disruptive Visits approved by the TAC has increased in recent Cycles.
12. Disruptive ToO are frequently observed within 2-4 days following an activation request.
13. Although policy defines Non-Disruptive ToO activations requiring >21 days, most are actually observed within~18 days.
14. The STScI scheduling process tends to place ToO observations early in an SMS.
15. The number of DD proposals and the number of DD proposals that meet the policy definition of "Disruptive" (i.e. <= 21 days) has increased in recent years although the numbers of significantly disruptive (T<=11 days) has remains constant and is relatively infrequent.
16. The DD approval process is generally faster than the implementation time (approval to execution) and thus does not constitute the rate determining step.
17. ToO observations are more effective in providing very short (T<5 day) observations than DD observations.
18. Truly rapid observations with schedule impact (T<11 days) occur 2.1 times per year from DD and 6.4 times per year from ToO observations. This implies an average frequency of one major disruption every 6 weeks.
19. We are executing rapid observations (i.e. within 21 days whether or not intended as disruptive) 7.65 times per year from DD and 28.6 (conservatively) times per year from ToO activations. This implies an average frequency of one response every 10 days.
20. Both the overall frequency and the frequency of shorter response time observations are increasing in recent years.
21. The majority of the DD proposals which meet the 21 Day Disruptive criteria execute with comparable delays to the Non-Disruptive ToO proposals.
22. Proposals asking for more than 7-8 days from activation and that are ready for scheduling on a Monday are significantly less disruptive than shorter timescales. This reflects the historical success of the scheduling team in placing ToO observations on the first or second day of a SMS.

23. Implementation costs increase significantly for activations less than 11 (7-8 days if provided on Monday) days although existing processes can handle such proposals successfully.
24. Hard limits preclude the HST plus TDRSS system from executing anthing other than a ultra-constrained type of ToO in less than 24 hours and 36-48 hours is a practical limit. This ultra-constrained ToO would need to reside within a 1-degree region on the sky specified in advance.
25. The present frequency of ToO observations is sustainable and could be significantly increased if desired (although process streamlining and/or increased staffing would be required).
26. Simple observations using CCD or IR detectors that do not fully use the visibility period are more likely to be rapidly implemented and has less impact on other observations and staff resources.

Acknowledgements

David Adler, Merle Reinhart, Denise Taylor, and Bill Workman provided essential inputs and comments on this white paper. Tom Brown, Helmut Jenkner, Neill Reid, and Rachel Osten also provided useful comments on an earlier draft.

Acronym List

APT – Astronomers Proposal Tool
CS – Contact Scientist
FGS – Fine Guidance Sensor
GS – Guide Star
GSFC – NASA's Goddard Space Flight Center
HSTMO – HST Mission Office
PC – Program Coordinator
PI – Principal Investigator
SMO – Science Mission Office
SMS – Science Mission Specification
TAC – Time Allocation Committee
TDRSS – Tracking and Data Relay Satellite System
ToO – Target of Opportunity

Appendix A: Example of a Disruptive ToO Timeline

In this case, the PI asked for it to be onboard within 5 days. For HST, the timer starts with the activation request.

6/26/17 9:08 EDT - Activation Request. 14 orbits, first visit 3 orbits to be scheduled within 5 days. Other 3 visits 3 orbits each evenly spaced over the following 2 weeks. (Note, this was a nearby supernova event that Swift caught.)

PC, CS and PI began immediate email communications since it was a normal Monday during working hours.

6/26/17 9:20 EDT - PI submitted proposal revision of the activated visits.

6/26/17 13:00 EDT - SMSB calendar rework started.

6/26/17 14:45 EDT - PI submitted a revision to the activated visits based upon the CS review.

6/26/17 15:08 EDT - CS signs off on the visits.

6/26/17 15:25 EDT - PC completes work and set visits flight ready.

6/26/17 15:45 EDT - SMSB end-of-shift. Will continue the next day since there is time.

6/27/17 08:00 EDT - SMSB continue from previous day

6/27/17 13:00 EDT - Calendar work completed.

6/27/17 13:10 EDT - SMS Generation completed.

6/27/17 13:43 EDT - PASS work started.

6/27/17 18:00 EDT - Products delivered to FOT for uplink and to OPUS.

6/27/17 19:15 EDT - Product Load Approval granted. Note, FOT has to staff outside of nominal day-shift hours to support this.

6/28/17 17:56 EDT - First modified load transferred to the payload computer

6/28/17 18:02 EDT - First modified load transferred to the spacecraft computer

6/28/17 21:53 EDT - Modified loads begin executing

6/29/17 05:20 EDT - ToO begins executing

Note, if this was done in one continuous shot, it would have taken about 16h 45m for the STScI part of the process, so two nominal 8 hour shifts. However, everyone knew that we had 5 days to get to this observation in the HST timeline, so it probably could have been compressed by a few more hours (likely to 12-14 hours which is about the best we've ever done with a disruptive ToO which I think in that case the goal was to get it executing within 48 hours of the activation).