Proposal Solicitation and Selection for the 21st Century

Palle Møller[a,b], Glenn Miller[a], Brett Blacker[a] and Meg Urry[a]

[a] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[b] on assignment from the Space Science Department of ESA

ABSTRACT

Writing, reviewing, and selecting the proposals which are to define the science program of any state–of–the–art observatory/space mission are all tasks which have grown in complexity, and as a consequence large amounts of time and effort are currently being invested in this process by proposers as well as reviewers. Viewed from the opposite vantage point, the currently used solicitation and selection process is a significant operational expense: Mailing paper copies of proposals and gathering reviewers for panel meetings and a "Time Allocation Committee" involves a large amount of logistical support and time by the observatory staff. Finally, the batching of proposals into yearly cycles (or six month semesters) increases the time from concept of a scientific idea to receipt of actual data which decreases the ability to respond to new scientific developments and also increases the general operational overhead of handling a large batch of observations.

In this paper we explore two experimental steps towards an optimal proposal selection process: streamlining the current process via "paperless" and "groupware" technologies, and use of a "steady state" process which accepts submission of and reviews proposals continuously. The pros and cons of each approach are examined and we demonstrate that not only are the enabling technologies available, but when resources are considered in a global manner we can identify both major improvements in the process and significant reductions in the expenditure of resources.

Keywords: Proposal Submission – Proposal Review – Phase I Process – Paperless Review – Continuous Submission

1. INTRODUCTION

Observational astronomy has, within the past two decades, evolved away from being done solely on a large number of comparatively small, often locally placed and privately operated, telescopes belonging to individual universities, towards larger facilities serving large communities. With an appreciable number of major observatory projects, both ground and space based, currently under way and several of those near completion, it is clear that the trend towards centralized operation in modern observational astronomy is continuing.

In parallel with the trend towards larger observatories, the modes of operation have also been changing towards becoming increasingly complex. In particular concepts such as "Remote observing", "Service Observing", and "Queue Scheduling" are now commonly considered as operational concepts at both ground and space based facilities. The goal of introducing such scheduling modes is to reach higher efficiency in the use of the available observing time, and to maximize the scientific value of the data from any given observation.

One downside to aiming for those goals is that the structure of the process by which we select, prepare and execute programs, becomes increasingly complex and resource intensive. Therefore, either the organization responsible for operating the observatory must have a large staff to support such a complex service–oriented structure, or a large amount of this work must be delegated to the proposers themselves and the reviewers of the proposal selection process.

While several other papers at this conference are focusing on the level of operational complexity offered to observers who have already been allocated observing time (in todays language commonly referred to as "Phase II"), we shall here concentrate on the optimization of the process which leads to the selection of proposals ("Phase I").

Our analysis of the present day HST Phase I process, as well as several possible improvements to it, is based mainly upon the experience gathered, from trial and error, at STScI during the support of HST cycles 0 through 7. We do, however, stress that this does not limit the use of our analysis to HST, or even to space missions. Many large ground based facilities are moving towards operational concepts coming ever closer to those used for space missions, and major parts of our analysis could be applied to modern ground based observatories with considerable success.

The paper is organized as follows: In Section 2 we introduce what we, for lack of a better term, shall name the "Basic Theorem for Data Quality Optimization". In truth it is not actually a theorem, but rather a guiding principle which we use to identify those process changes which are actual process improvements in a global sense. In Section 3 we analyze the currently used Phase I process. In sections 4 and 5 we introduce two experimental ideas for process improvements, and we use our guiding principle from Section 2 to compare them to the present day process, thereby identifying which improvements they would provide. In sections 6 and 7 we summarize our findings.

2. BASIC THEOREM FOR DATA QUALITY OPTIMIZATION

Operating a modern observatory at a zero support mode of simply handing out observing time and leaving the observer with the responsibility to do the rest, is clearly a non–optimal mode of operation. On the other hand, attempting to operate an observatory at a complexity level which is beyond that which it can support will increase the probability for process errors, which in turn will cause a drop in efficiency. Therefore clearly:

- **The optimal operational concept for any given observatory is where one strikes the best balance between resources available and selected operational complexity**.

While this appears to be a self–evident statement, we consider it a basic theorem for improving observing efficiency and data quality, and in what follows we shall return to this statement as the decisive element. However, the true advance in efficiency does not come from stating the self–evident, but by identifying precisely which resource to balance against which operational concept. In particular the resources considered should not be limited to those supplied by the observatory (in what follows referred to as "internal resources"), but also those supplied by the community being served by the observatory (external resources).

3. ANALYSIS OF TODAYS PROPOSAL SOLICITATION AND SELECTION PROCESS

3.1. Goal of the Phase I Process

The Goal of the Phase I process is to *select the proposals which are capable of bringing about the greatest advance of knowledge, for the smallest allocation of observing time*. The selection process is therefore one which has to take into account several elements such as:

- **Scientific goal.** What is the expected advance of knowledge?
- **Scientific feasibility.** Are the proposed observations likely to provide the expected advance of knowledge?
- **Technical feasibility.** Is it technically feasible to successfully complete the proposed program?
- **Resource balance.** Is it a reasonable investment of observing time to obtain the expected advance of knowledge? Could the same observations be done with lesser resources? Could the same conclusions be reached by a much simpler observation?
- **Proposer resources.** Do the proposing team have access to the resources needed to produce and publish the final result?

An excellent summary of the inherent problems and pitfalls in trying to devise a fair and unbiased selection process for the UK telescopes was given by Edmunds, 1996.[1] This part of the process faces exactly the same problems whether one considers a ground based or a space based observatory, and we shall therefore not dwell further on this issue but refer the reader to the detailed discussion given by Edmunds.

3.2. Resources Available, and Necessary, for the Phase I Process

As already pointed out by Edmunds, it is important at the outset to admit that the final selection cannot ever become a fully objective, unbiased process. The responsibility of the observatory is to ensure that the process becomes as objective as possible, for an acceptable expenditure of resources. To be able to perform such a selection, the observatory relies upon the help from experts in the community (in what follows named reviewers or panelists) who review and rank the submitted proposals.

For the purpose of the present analysis we shall consider the following as a resource which has to be spent, or provided, and therefore needs to be taken into account for the purpose of optimizing the balance as stated in our optimization theorem above:

1. **Documentation.** The best way to make it possible for the reviewers to identify excellent science is to ensure that proposers have easy access to the information they need to write the best possible proposal. The documentation which should be considered here is not merely instrument handbooks. Basic information (such as "How do I include a figure in my proposal") should be easy to find, or hundreds of proposers may spend a day each searching for this information, easily resulting in the loss of a full year of research time total. When preparing documentation for Phase I it is important that the organization of the documentation is such that basic information is easy to find. Technical details are often not necessary for a Phase I proposer, and might often better be left out of Phase I documentation.

2. **Proposal writing.** The total resource spent writing Phase I proposals is quite substantial. For Cycle 7 STScI received 1298 proposals. Assuming a mean investment of four working days to complete and submit an HST proposal, this amounts to 20 years of work for a full time researcher at a single HST submission deadline. Any effort the observatory can provide to make it easier to write a proposal, or even to make it easier to realize that a prospective proposal is not feasible/competitive, is clearly going to be a good investment in terms of integrated expenditure of resources.

3. **Proposal reviewing.** This also is a substantial resource. While it clearly takes less time to read and evaluate a proposal, several reviewers (at STScI three primary reviewers, nine in total) are assigned to each proposal. Typically of order 100 astronomers, spend 1-2 weeks each reading and reviewing proposals, and another 3-4 days discussing the proposals during the panel meetings. This again amounts to about 4 years total time spent. The same argument can be made here as above. Anything the observatory support staff can do to improve the quality of submitted proposals, or even to ensure that only proposals with a realistic chance of competing for time gets submitted, can drastically reduce the total time spent in this step of the process.

4. **Observatory support staff.** Members of the support staff are responsible for the availability of documentation, proposal templates and proposal handling software, are available to answer questions from proposers, they receive, sort, copy, and forward proposals to reviewers, perform review for technical feasibility, they support the panel review process, answer questions from panelists, and provide logistic support during the panel meetings.

5. **Expenses.** Printing and mailing of large amounts of documentation. Printing, copying and mailing of large amounts of proposals to reviewers. Travel and hotel costs for more than 100 panelists.

6. **Time.** While the time lapsed from the submission of a proposal to the time when data are received is not an obvious resource, it does, in an indirect sense, have a substantial impact on the final cost/benefit ratio. At present there is typically one full year between HST submission deadlines. An idea for an important observation can be conceived by a proposer at any time, and will in the mean be 6 months away from the next submission deadline. Nominally a cycle starts observing about 10–11 months after the Phase I deadline, and lasts for 12 months. Hence, in the mean an observer will receive data 22-23 months after the idea was conceived. In the nominal worst case (getting an idea just after a deadline, having data obtained in the end of the cycle), the wait–time is 3 years. The real worst case is often even longer because of scheduling constraints. The evolution of science is so rapid that in many cases an observation which would have been a major step forward three years earlier, may hardly be interesting anymore three years later. While there is no special extra cost to delaying the observations for three years, it may however have rendered the observations far less valuable, thereby effectively decreasing the value of the selection process. Also, a long lead–time for proposal execution makes it difficult for proposers to respond to important scientific developments.

3.3. Summary of Analysis

While it is stressed that the process of proposal selection is one of crucial importance for any observatory, we have identified areas in which it is extremely resource demanding. The process presently used at most observatories is that of inviting proposals at annual (or bi–annual) submission deadlines, and then collect a large number of experts who, in a very intense and somewhat stressful couple of weeks, proceed to select what they believe are the best of the submitted batch of proposals.

- This process is expensive (copying, mailing, travel and lodging)
- This process produces extremely long execution lead–time
- This process produces strong peaks of work–load, both internally and externally
- It is not evident that this process produces the best observing proposals
- It is not evident that this process produces the best review

Below we shall describe two alternative submission and selection strategies, and investigate how the introduction of those strategies would influence the expenditure of resources. The two alternative processes we here will discuss are "Paperless Review", and "Steady State Submission".

4. PAPERLESS REVIEW, PROS AND CONS

The principles of a paperless (or fully electronic) review are fairly obvious, and we shall hence address this concept first. Considering the presently available technologies, we can identify two possible modes. Proposal submission, distribution to panelists, and return of grades from panelists could all take place via either email or by posting proposals directly to the WWW. Naturally appropriate permissions must be set to ensure that strict confidentiality is never compromised.

4.1. Paperless Review: What would we gain?

The potential gain in moving away from the cumbersome handling of large amounts of paper is quite significant, both in terms of internal and external resources. In terms of the numbered items listed in Section 3.2 above they are:

1. **Documentation.** The move towards electronic documentation is already far advanced, and has substantially reduced the costs of documentation printing and mailing.
2. **Proposal writing.** Essentially all observatories have now moved to electronic templates and electronic submission. One of the outstanding problems still not completely solved is the inclusion of figures, and in general the submission of formatted (PostScript) proposals. Grey–scale figures can come out quite different on different printers, paper–size standard is different in different countries, inclusion of color figures is difficult to support (printing is slow and costly, copying is difficult). In a fully electronic review process, notably one which is WWW based, the figures never have to get printed, they can be viewed directly over the Web. Essentially this will improve the quality of submitted proposals, and make it simpler and safer to submit proposals with figures. It does require, though, that reviewers have access to high quality color monitors.
3. **Proposal reviewing.** The full review is done in the comfort of the reviewers own home/office. Reviewers have direct access to high quality figures supporting the scientific case. The text of the proposal is easily printed if the reviewer prefers a paper copy. The travel to and from the panel meeting is no longer necessary.
4. **Observatory support staff.** The support staff no longer needs to print, copy, and mail a large number of proposals. Instead proposals will be circulated by email, or posted on high security Web–sites with restricted access. Support of the review process will, as discussed in Section 4.2 below, likely increase a little due to the lack of face to face contact of the panelist. The remainder of the support tasks will remain constant.
5. **Expenses.** Will plummit.

6. **Time.** Paperless review will have no impact on the lead–time for proposal execution.

In summary: Introduction of paperless review would, at first glance, significantly lower direct expenses, and also introduce a welcome saving on other internal resources. Proposers would be given a new and powerful tool to produce better proposals in an easier way, reviewers would save travel time and have better proposals to review. As discussed in Section 3.2, even a small saving produced in those two areas will add up to a major reduction in total external resources.

4.2. Paperless Review: Potential Problems

One major argument often raised against paperless review is the lack of interpersonal contact which sharpens the arguments for and against any given proposal during the panel meetings. While it is true that this contact is lost, it is not a priori clear if this is an overall disadvantage, or in fact an advantage, for the process.

First of all, as already strongly stressed by Edmunds[1] the few excellent *must do* proposals, as well as the obviously flawed proposals, are easily picked out by all (or most) panelists. The remainder of the batch of proposals are all good, and there is no such thing as unique, objective ranking of those. The actual final ranking of those proposals will be a function of who serves on the given panel that particular cycle, what is their own scientific interest, what do they believe is the most interesting question to pose exactly at this moment. Also, as much as one would like to eliminate this, there are interpersonal relations which will also flavor the ranking.

One may argue that a hidden weakness in a given proposal which is only recognized by a single panelist will cause that proposal to be discredited during the panel discussion. However, while this is certainly the case, and we do see this happening from time to time, we also sometime see the strong personality of a single panelist dominating a panel discussion to the point where that single panelist has far too much influence on the final ranking. The support staff screens all proposals for all cases of primary conflicts of interest prior to the meetings, but the communities are too small to eliminate all secondary conflicts.

In a paperless review, after votes and comments have been received by the support staff, a preliminary ranking will be produced. Where the support staff finds that there may be reasons to suspect that a given proposal may have been erroneously ranked (high dispersion, or a comment from a single panelist pointing out a potential flaw in an otherwise highly ranked proposal), some action should be taken to invite a discussion of the given proposals (*e.g.* tele–conference or email discussion). This will not completely make up for the lack of face to face contact, but it will serve many of the same purposes – for a much lower cost in resources.

Another argument often put forward in support of the panel meetings is that the community to observatory contact is maintained in this way. At STScI it is true that the panels do actually physically meet at the institute itself, and the panelists do get to meet and interact directly with the support staff. However, this interaction happens at the time of peak–workload for the support staff, the busiest time of the cycle, and also during a stressful work–intensive time for the panelists.

Even in a paperless review it is certainly necessary that the support staff and the panelists get to know each another. It is equally important that the panelists get to discuss some proposals together so that they can learn the "calibration" of each other panel member. In a paperless review process this does not have to happen in the stressful atmosphere of the review itself, but could happen sometime mid–cycle when there would be plenty time to take out a full day for the panel to meet and discuss procedures, discovered flaws and suggestions for improvement, and also discuss a handful of typical proposals and some "trouble cases" from the previous cycle. The full value of such mid–cycle meetings will become evident after the introduction of the "Steady State Review" (Section 5) and we shall return to a more detailed discussion there. Mid–cycle meetings would, of course, mean that the resource reductions under items 3, 4 and 5 would be less than outlined above.

5. STEADY STATE SUBMISSION, PROS AND CONS

By a Steady State Submission (triple–S) scenario we mean a process without a submission deadline. Submission of proposals will be accepted at any time. The aim is to achieve a continuous process flow where proposals are submitted and reviewed at a steady rate, and are executed at a steady rate, without any process "bottle–necks".

Having an annual (or bi–annual) deadline for submission of observing proposals is a tradition at most "visitor–mode" operated ground based observatories. At observatories where observers are present during observation, and where time consuming instrument changes have to be minimized, it is necessary to have a well defined and stiff schedule well ahead of the actual observing time. Having a deadline at such observatories makes good sense.

For HST, or for any other "service–mode" operated telescope, all one really needs is to have Phase II proposals ready in time to get them onto the telescope a few weeks ahead of execution. Working backwards from that, adding time for Phase II preparation and peer review, it would not be unrealistic to aim at a continuous process where a proposal received by STScI could, on a regular basis, go to the telescope for execution some 7-8 weeks later. It should be precised that such a short lead–time could be the *goal* for a substantial fraction of proposals, it will never be possible to *guarantee* execution of all proposals in such a short time. Implementing such a system would clearly mean re–thinking the entire Phase I, as well as several changes to Phase II, but would bring about many improvements. Here we shall list those we find of most immediate value.

5.1. Steady State Submission: How would it work?

Any proposer would be free to submit a proposal at any time. Once received, the proposal undergoes a technical review, and is at the same time sent to panelists for a science review. The panel is in permanent session, but does not meet. Proposals are reviewed at home (could be a paperless review as described in Section 4). Once the proposal has obtained a final grade it enters the pool of graded proposals. When a proposal has entered the pool it can stay there forever, or one can choose to assign a maximum "pool swim–time" to each proposal after which time it is removed. When schedulers (on a continuous basis) determine that they need new proposals, they pick the topmost in the proposal pool which will fill the free slot in the schedule. Only then the PI is asked to submit a Phase II proposal. It is fast, it is efficient.

5.2. Steady State Submission: What would we gain?

- **Proposal lead–time.** The extremely long time from the birth of a scientific idea to the execution of HST observations was commented on in Section 3.2, Item 6. In particular this has as a consequence that presently observations for a given cycle will typically *not* be available before the deadline of the next, so writing proposals based on previous data means that proposers typically have to wait for two, sometimes three or four, cycles before following up on HST observations. In a triple–S process proposers can follow up on observations immediately.

- **Hardware response time.** Telescope and instrument performance can in practice be substantially different that anticipated during construction, and in addition, instruments and telescopes can change behavior with time. Preparing long cycles means that we increase the probability of having a large pool of proposals *not* well suited for the telescope in its current state, and not being in a good position to react fast to such changes as they occur (NICMOS recently provided an unfortunate reminder of this). In a continuous submission scenario the normal response time to telescope performance should be of the order 7-8 weeks (the time for proposers to respond to a changed environment), so calling in extra "Δ–cycles" and delaying accepted proposals will no longer be necessary.

- **Cycle–length, SMs, and proposal rejection rate.** The fairly frequent Servicing Missions (SMs) are a constant disruption to the cycle structure in any case. To prepare for a possible delay in a given scheduled SM, STScI usually calls in extra proposals to create a "contingency buffer". Similarly, at a ground based observatory with service observing one might wish to accept more proposals than can be executed in a given semester, to ensure a high efficiency even under unexpected weather conditions. Such a buffer of accepted proposals will automatically increase the execution lead–time if the proposals are carried over for execution in later cycles/semesters, or will cause a general loss of resources (both support staff and proposer Phase II work is lost) if they are not carried over. To get rid of the SM contingency buffer one might prepare for a single very short cycle following the SM. This could result in strong variations in the proposal rejection rate, effectively making it easier to obtain observing time in some cycles, and harder in other cycles. A highly variable rejection rate is clearly a situation one should try to avoid. In a continuous submission scenario we shall always pull the proposals as we need them, and hence we do not need any special plans to prepare for SM delays. The actual rejection rate at any given time will be a slowly evolving function of time, always reflecting the current need for HST observations. In a ground based scenario it also makes it possible to call in Phase II proposals as needed, when it becomes clear which weather conditions there is a lack of proposals for.

- **DD proposals and scientific response time.** Just as it is difficult for proposers to quickly respond to hardware changes during long cycles, it is also difficult for them to respond to important scientific developments. The only way a proposer currently may be able to shorten the extreme lead–time for HST observations is by submitting a Directors Discretionary time (DD) proposal. DD proposals are much more resource demanding, and therefore much more expensive for STScI to handle, than are normal proposals. In a continuous submission scenario proposers can react to scientific developments on a zero delay timescale, and most (probably all) of the DD proposals should hence disappear.

- **Deadline induced proposals.** When reading HST proposals it is clear that most are excellent proposals which would have been written and submitted under any submission scenario. One does, however, also get the impression that a lot of proposals seem to be last–minute–half–committed efforts written mostly simply because "there is a deadline and the next is going to be a long time from now". In a continuous submission scenario proposers will have time to think more about such proposals and either decide not to write them after all, or be able to write high quality proposals instead. It is our belief that a lot of the "bottom 20-30% proposals" which are spending panelist review resources for no good reason, are of this category. It is realistic to expect that allowing continuous submission will bring the *number* of proposals down, but the *quality* of proposals up.

- **Science induced proposals.** The ideal situation is that of only receiving proposals which are triggered by a scientific idea or an interesting new result. Such proposers are highly motivated for the right reasons, probably leading to excellent proposals. Under the current system you may have to wait 6-12 months before you can submit a follow–up proposal to your interesting result. After having submitted it you still have to wait another 12-24 months for your data. This situation is strictly non–optimal.

- **Panel resources.** The total panel workload will be approximately unchanged under a continuous submission scenario, except that the panelists no longer have to deal with the "Deadline Induced Proposals" described above. Panel work will, however, be evenly distributed over the time each panelist serves. Panelists will not need, as they do now, to take out a solid 2-3 weeks of work to read proposals and come to STScI for the meetings.

- **Phase I resources.** The Phase I workload at STScI is currently extremely uneven with an intense peak between Phase I deadline and Phase II deadline, followed by a quiet period where the Phase I team tries to recover from the stress. It would certainly be a less stressful working environment if the work could be distributed more evenly.

5.3. Steady State Submission: Potential Problems

As seen from the list above there are many advantages to be found in implementing a triple-S process. It is, however, also clear that one cannot have panel meetings in session permanently, so the triple-S process must by default mean that panelists perform their review at home, and only submit grades to the support staff. It would be natural to do this in the form of a Paperless Review as described in Section 4. Several potential problems can be identified in such a scheme, but none of them seem to be unsolvable.

- **How can reviewers rank proposals which arrive one by one?** Clearly one would need a "start–up phase" to boot the system and calibrate the reviewers. For this purpose we would need a reasonable set of typical proposals (could be from previous cycles for that matter, but they should preferable be *real* proposals in the sense that the panel is actually allocating real time). The panel would be invited for a short and non–stressful visit where the process is described in detail, and they would proceed to discuss a set of proposals spanning a wide range of quality. The panelists then keep this set of proposals as a reminder of their own calibration, and when a new proposal arrives they compare its scientific quality to their set of calibration proposals. This should ensure a consistent review by each panelist.

- **How is long term consistency achieved?** Actually this can be done much better than at present. Presently STScI has a 100% change of panel members for each cycle, hence no long term consistency in terms of panel memory. With panels serving over longer time, the ideal situation would be to substitute panel members at a slow rate, hence maintaining consistency while new members are brought in and being calibrated. The panels would meet at regular intervals (once each year, or each other year) introducing and calibrating new panelists, de-briefing the leaving panelists. The panel will discuss problematic cases identified since last meeting, and also discuss a handful of representative proposals spanning a wide range in grades to calibrate the new members. Since there is no deadline the panels can meet at random times, ensuring a non–stressful atmosphere.

- **How do we adjust fractional time given to sub–disciplines?** This could be arranged in many different ways. The bottom line in all of them must be that a Time Allocation Committee (TAC) should meet at given intervals to discuss how well the process ran since the last TAC meeting, and set priorities and rules for allocation of observing time until next TAC meeting. Those rules could be a given fraction of time to each panel, it could be a "sliding average fraction" of time requested in each panel, a combination of the two, or something entirely different.

- **Duplications?** Duplication is certainly an issue here. Several scenarios for how to treat duplications can be worked out, but they are too convolved to be included in the present discussion.

Again applying our theorem in terms of the resources listed as numbered items in Section 3.2 above, we find:

1. **Documentation.** Documentation only needs updating when there is an actual change. There will be no need to produce new documentation at each arbitrary deadline.

2. **Proposal writing.** Proposals will be *Science Induced*, not *Deadline Induced*. Expected to create fewer but better proposals.

3. **Proposal reviewing.** The review is done in the comfort of the reviewers own home/office. Proposals will arrive as a steady flow, there will be no high intensity peak. Fewer and better proposals will be easier to review.

4. **Observatory support staff.** The support staff (both Phase I and Phase II) will have a continuous work–load, no high intensity peak.

5. **Expenses.** The drop in expenses will come from the paperless review. There will be no extra saving by going to a triple–S process.

6. **Time.** Will plummit.

6. DOES ANY OF THIS APPLY TO GROUND BASED FACILITIES AT ALL?

As we pointed out at the outset, this paper is not intended to be addressing only space missions. It is true that the analysis we have presented is based mainly upon the experience we have gathered during the support of HST, and therefore is not generally applicable to all observatories. It is never the less our hope, and our belief, that as more and more observatories move towards more automated, more complex, modes of operation, that if they first turn their head towards HST to try and identify how many of the lessons we have learned might be of value to them, they will find that transition much smoother. For a modern observatory operated in service mode there are many parallels to space missions, and there is a vast experience collected at STScI which we shall be happy to share with colleagues at other observatories.

7. CONCLUSIONS

We have presented an analysis of the Phase I process currently used at STScI, and we have introduced a method to help identify which alternatives to the current process will be improvements in a global sense. Using this method, we have further analyzed two alternative processes: a "Steady State Submission" scenario, and a "Paperless Review" process. We have found that both of those processes would result in improvements to the current process, and neither of them have obvious unsolvable problems connected with them.

REFERENCES

1. M. G. Edmunds, "Telescope Time Allocation: Logic or Lottery?", *New Observing Modes for the Next Century*, eds. T. A. Boroson, J. K. Davies, and E. I. Robson, ASP Conference Series, **Vol. 87**, pp. 22-29, 1996.