More Bang for the Buck:
Lessons from a Decade of Keck Mainland Observing Operations

Gregory D. Wirth[a], Robert I. Kibrick[b], Robert W. Goodrich[a], James E. Lyke[a] and Jeff A. Mader[a]

[a]W. M. Keck Observatory, Kamuela, HI 96743 USA
[b]UCO/Lick Observatory, University of California, Santa Cruz, CA 95064 USA

ABSTRACT

Most astronomers love observing remotely because they can gather their data while avoiding the inconvenience of traveling from their home institution to the observatory; however, running a true, facility-class remote observing system can impose a significant burden on the observatory support staff. Although usage of the remote observing capability at Hawaii's W. M. Keck Observatory has risen substantially in the nine years since we began allowing astronomers to observe from California and elsewhere, Keck's system has evolved to provide better operational performance with less manpower. We describe improvements to our system in several key areas: scheduling remote observing requests; maintaining the virtual servers used by remote observers; monitoring the health and safety of equipment at remote observing stations; and abandoning the distinction between "local" and "remote" observing. These measures have substantially lessened the effort required to operate the system, provided new capabilities for training and troubleshooting by the technical staff, and enhanced the convenience and reliability of the system for observers and support staff at the remote sites.

Keywords: remote observing, VNC protocol, observation scheduling

1. INTRODUCTION

The W. M. Keck Observatory (WMKO) has offered its observers the opportunity to carry out observations remotely from sites outside of Hawaii for close to a decade. A long-term development effort involving WMKO and UCO/Lick Observatory has resulted in a hardware and software system that is reliable, robust, and convenient for the astronomers to use. The current state of this system has been described at length in the accompanying review by R. Kibrick,[1] which presents the valuable perspective of the lead developer on the project. In contrast, the present paper will put less emphasis on the technical aspects of the system, and will instead focus on what lessons have been learned from the Observatory's perspective of operating the system.

2. BACKGROUND

When the Keck Telescopes first entered science operations circa 1993, observations were carried out in the traditional way with the astronomer and the observing assistant[*] both present at the Mauna Kea summit. But nearly as soon as the telescopes had commenced science operations, observers were eager to get off the mountaintop. A recently-unearthed memo from Caltech astronomer Gerry Neugebauer dated September 22, 1994, reads in-part: "I have been thinking about the possibility of doing remote observing from Waimea." Later in the memo, he adds "All my experience indicates that the non-technical problems are as serious and as difficult as the technical ones."

Here, in a nutshell, is the crux of the issue faced by the astronomers and the observatory staff. On the one hand, the observer wants to be able to carry out his or her observing program as painlessly as possible. In 1994, this meant being able to operate the instruments from the relative comfort of Keck's Waimea headquarters, instead of gasping for air at

---

[*]In Keck parlance, the observatory staff member who operates the telescope is known as the "observing assistant."

the telescopes atop the 13,796 ft summit of Mauna Kea. At this early stage in the operation of the Observatory, the staff were probably fully focused on the main task of getting the telescopes and instruments operating optimally and producing great science; observer convenience was likely to be of secondary concern. Nevertheless, the observers' desired ultimately trumped institutional inertia: by 1996, Waimea became the primary observing location for both Keck telescopes.

Being able to observe from Waimea was a significant step toward making observing more convenient, but astronomers still faced the prospect of having to travel all of the way from California (or elsewhere) to Hawaii in order to complete their observing runs. For most people, this meant that even a single night at the telescope entailed an investment of at least 3 days on the road, with no guarantee that the weather and equipment would yield any usable data. Furthermore, those who faced travel restrictions (due to a medical condition or obligations such as teaching) faced the difficult decision of either traveling in spite of their restrictions or, alternatively, relying on others for the crucial task of gathering their observations. Once again, astronomers began to voice their desire for greater convenience by being allowed to observe with Keck remotely from somewhere on the U. S. mainland.

In 2002, the first experiments with a VNC-based[†] mode of Keck "mainland remote observing" commenced at Santa Cruz, California,[2] and shortly thereafter WMKO agreed to implement this as a new observing capability. This plan was adopted in spite of some internal resistance, for the belief among the WMKO support staff was that adding such capability could involve significant additional work for the support astronomers while also presenting the challenge of working with observers who were at a distant site instead of being in the same room.

The common theme in both of the above cases is that observers and observatory staff have not always seen eye to eye in regard to remote observing. For observers, the top priority is that remote observing maximize the convenience of observing; in contrast, for the observatory staff the top priority is to minimize the amount of additional work that remote observing will entail. Put another way, the observers think primarily about the impact of remote observing (i.e., the "bang") and the support staff concentrate on the cost (i.e., "bucks"). The challenge for the designers and operators of the remote observing system is to maximize the return on investment, and thus the remainder of this contribution will explain Keck's approach to maximizing this "bang for the buck" metric.

3. COMPONENTS OF A REMOTE OBSERVING SYSTEM

The two essential components of a functional remote observing system are fairly obvious: a method for providing the remote observer with access to instrument controls, and a system which allows the remote observer and the observing assistant (and the support astronomer, who will generally be involved at the start of the night) to communicate. In fact, one could choose to operate a remote observing system with nothing more then a remotely-displayed instrument GUI and a telephone. Such a no-frills system could prove adequate for a situation involving a very simple instrument and only infrequent need for remote usage. However, our experience suggests that a remote observing system that aspires to be convenient, robust, and reliable should include these additional components:

Handling observer requests: A method is needed for accepting requests from users who want to observe remotely and scheduling their usage of the requested remote site(s).

Managing VNC servers: The system must provide a means of launching the required VNC server processes prior to the observing run and shutting them down again when the run is completed.

Launching VNC viewers: The observers at the remote site require a way to start up the VNC viewers such that they connect to the proper VNC servers.

Monitoring remote computers: Prior to each observing run, a person or process must ensure that the computers and networking equipment to be used in remote observing are fully functional.

While any remote observing system should have these components in some form, the particular implementation of these components could be either simple or sophisticated. The design which provides the optimal bang-for-the-buck at a particular observatory will depend on how frequently the system is put to use.

[†]Virtual Network Computing[1] (VNC) provides a desktop that can be simultaneously shared among multiple sites. Multiple VNC viewers can connect to any given VNC server to provide this shared environment.

Keck's formal mainland remote observing project begin in 2002 as a relatively simple system which was adequate for limited use. As shown in Figure 1, even as recently as 2005 the system was used on only 10% of the observing nights at Keck. As such, our operation began in what we might call the "low-usage domain." In this case, the most important consideration is expediency; getting the system working quickly is paramount, even if substantial effort is required to operate the system each time it is used. In the low-usage domain, it may not be wise to invest heavily in making the operation highly efficient because the usage does not justify the effort required to do so.

With time, observers have grown to appreciate the convenience of observing with Keck remotely from the mainland; correspondingly, demand has now risen to the point that our mainland observing system is used on nearly 50% of Keck nights. With two telescopes in operation, the system is thus used almost every day on average. If we were attempting to operate with the same cumbersome techniques that we employed in the early days of the system, a substantial amount of staff time would be required to operate our system. Instead, we've gradually identified areas where investing the effort to make the system work more reliably and efficiently could reduce the effort required to keep the system operating. We provide specific examples to these enhancements in §5.
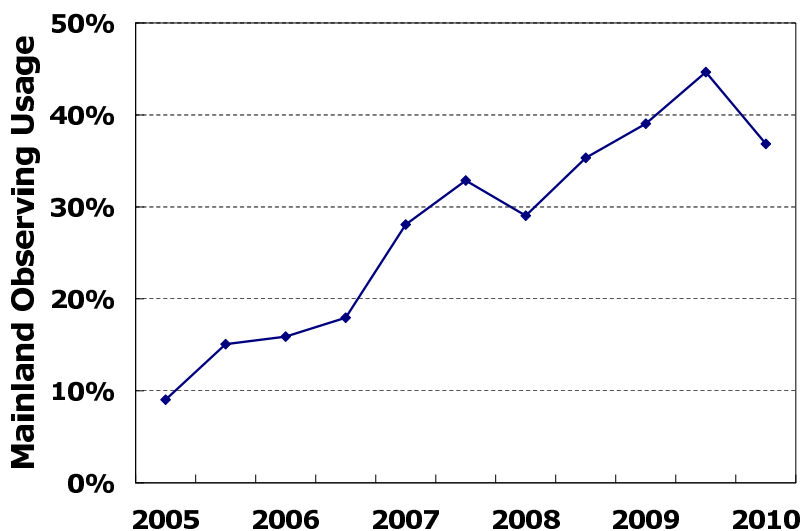

Figure 1. Fraction of Keck observing nights scheduled for mainland observing per semester since 2005. Although mainland observing was relatively infrequent 6 years ago, nearly half of scheduled nights now rely on this capability.

4. SIX LESSONS FROM KECK'S REMOTE OBSERVING EXPERIENCE

The current budgetary environment for astronomy is difficult, and we presume that most observatories are considering ways to provide services in the most cost-effective manner. Implementing remote observing can help save observers the cost of travel, but offering this service at a particular observatory only makes sense if this capability can be provided at reasonable cost. In the hope that our experiences at Keck will be relevant to other observatories, we now present six key lessons we've learned over the last decade as our remote observing system has evolved into the "high-demand" regime.

4.1 Enforce uniformity among sites

In retrospect, the single best design decision we made was to equip all of the remote (and local) observing sites as similarly as possible, because doing so has meant that we can rely on software and hardware to always work consistently. All of our mainland observing sites are equipped with Polycom video-conferencing systems which function well and reliably. The computers at each site were originally all SPARC/Solaris workstations running the FVWM window manager, but with the recent demise of Sun Microsystems we've been forced to seek alternatives and now employ Linux-based systems (configured to run FVWM for consistency) at our newest mainland observing sites.

On rare occasions (typically emergencies) we have agreed to support remote observing at sites other than our dedicated ones. In every case, the effort required to set up the VNC connections turned out to be greater than we had hoped, generally

because of configuration differences between the remote computers and the local ones. It's clear from our experience that being able to rely on a known configuration for the video and computer systems at the remote site saves a substantial amount of observatory effort.

4.2 Automate the tedious tasks as demand justifies

In the low-demand regime, you can afford to put up with cumbersome steps in the remote-observing process as long as the system is infrequently used; however, as demand increases these inefficiencies can contribute substantially to the effort required to operate the system. We can point to three examples of ways that we have been able to use process automation to reduce the amount of time we spend maintaining the system:

Handling usage requests was originally accomplished exclusively through manually-generated email exchanges. The observer would begin by emailing the Mainland Observing Coordinator at Keck, who would then email both the support astronomer assigned to the run and the Mainland Site Coordinator at the requested observing site(s). Once all of these parties had returned email conveying their approval, a final email from the Mainland Observing Coordinator to the observer would confirm approval of the request. While this approach worked well at first, the process of composing and tracking all of these emails eventually became the leading manpower drain on the system. As described further in §5.3, we have now replaced this awkward method with an automated system that tracks the status of all requests in a database and generates email automatically; this has reduced the effort for processing requests to a negligible amount.

Managing VNC servers was done by hand at first: the Mainland Observing Coordinator at Keck would regularly consult the upcoming telescope schedule to determine which observing accounts would be in use, and would issue a command to launch the required VNC servers. Now, as described in §5.1, the task of launching the VNC servers is completed by a fully-automated process which once a day reads the telescope schedule, launches new VNC servers for upcoming runs as needed, and terminates those for observing runs that are now completed. The manpower required for this aspect of the system has been reduced to zero.

Launching VNC viewers at the remote sites to connect with corresponding VNC servers at Keck was originally a tedious task which involved the remote user having to enter numerous commands in order to start the required ssh tunnels and to launch particular viewers on particular screens. The scripts[3] now deployed at each remote site will query the appropriate VNC host machine at Keck to determine which viewers are running, launch the appropriate ssh tunnels, and start the corresponding VNC viewers as part of a single command. In the near future, we plan to streamline this process further by replacing the command-line interface with a GUI that will prevent the remote observer from entering the incorrect arguments to the script (which happens only rarely but causes needless confusion).

Although the effort required to automate these steps was significant, the fact that we operate in a high-demand regime means that these modifications have more than paid for themselves through ongoing manpower savings.

4.3 Regularly check remote computers and networks

Nothing causes stress and aggravation for the support staff like learning on the day of an observing run that the computers, video-conferencing units, or VNC servers are not functioning properly. Keck uses two methods to verify that these components are always kept in good working order so that they are ready when needed:

Pre-run checkout of the remote video-conferencing units and computers is performed in the days prior to the observing run by the observers and/or the mainland site manager. By bringing up the VNC viewers and making a Polycom connection to the remote observing rooms in Waimea, the remote personnel ensure not only that the equipment and software at the remote site is fully functional, but also that the computers, VNC servers, and video-conferencing units at Keck are working well. Additionally, these sessions serve as a refresher to the observing team on how to operate the remote equipment and ensure that they know all required passwords.

Automated daily checks on the health of critical routers and workstations at the remote sites are performed automatically each day as described in §5.2.

Both of these techniques have proven effective in uncovering potential problems before they affect observing.

4.4 Treat all observing as potential remote observing

For several years, we made a distinction between observing runs which involved mainland observers and those which did not. When only local observers were involved, the observing team would launch their instrument software such that it was displayed directly to their desktop in Waimea; in contrast, teams involving remote observers would start out by bringing up their VNC viewers first and then launching their instrument software within the VNC session. Thus, this approach required us to maintain two distinct sets of startup instructions for Waimea-based observers and often proved confusing for observing teams. Eventually, we decided that it was not necessary to draw any distinction between "local" observing and "mainland-only" mode observing; instead, we have switched to a mode we call "All VNC All the Time" under which instrument software is always run within the VNC desktop environment regardless of whether any remote observers are slated to participate. This not only unifies our instrument startup/shutdown procedures, but offers important additional benefits as described next.

4.5 Be open to unanticipated uses for your remote observing system

Although the original impetus for providing remote access to the instrument desktops was to enable observing from the mainland, we've found several additional unanticipated applications for the system that further justify its development.

Observing assistants (OAs) now have the ability to "eavesdrop" on instrument operations. Since the OAs are the only staff on duty all night, they form the first line of defense in troubleshooting instrument problems during the night after the support astronomer has departed for home; thus, it behooves them to become familiar with the essential aspects of instrument operations. However, the OAs operate the telescopes from the Mauna Kea summit under most circumstances, and are generally in a different physical location from the observing team, whether that team is observing from Waimea or an off-island site; this separation has long prevented the OAs from learning more about instrument operations. With the advent of "All VNC All the Time" mode, the OA can easily bring up a copy of the observing team's instrument VNC sessions in a "view-only" mode that permits them to witness the instrument operation without interfering. VNC thus permits the observing assistants to develop into more effective troubleshooters.

Local troubleshooters now have the ability to view exactly what the observer is seeing when attempting to resolve a problem. Instead of having to drive into the office from home, the support astronomer can respond to a trouble call at night by bringing up the observers' VNC sessions from home, which generally proves adequate to identify and fix the problem. VNC thus makes the support astronomers more effective troubleshooters, as well.

Remote instrument teams can now interact directly with Keck staff on instrument tasks within VNC. In the past, the WMKO staff had virtually no hands-on time with new instruments until they were delivered to Mauna Kea. Now, staff can use VNC viewers to run the instrument remotely from the site where the instrument is being built and thus can provide useful feedback to the instrument development team during the critical late phases of construction, before final design decisions have been rendered. Additionally, VNC allows the mainland-based instrument development teams to interact with the WMKO staff in troubleshooting the inevitable instrument problems that must be resolved during the commissioning phase at the telescope.

Multi-site observing teams now routinely exploit the mainland observing system to collaborate on observing projects, overcoming the geographical constraints which otherwise keep them apart. Whereas teams split among different campuses (or even institutions) once counted on being able to get together at the telescope, they now use observing nights as their chance to rendezvous via the video-conferencing and VNC sessions.

The advent of "All VNC All the Time" mode operation at Keck has thus enabled new uses of VNC that were not previously anticipated.

4.6 Be prepared to draw the line on what is "too remote"

We pointed out in §1 that observer demands for greater convenience in observing led to the original shift from operating instruments on the Mauna Kea summit to operating from Waimea, and eventually led to operating from our dedicated mainland sites. But with today's technology there's no need to stop there. Observing could be carried out from observers' private offices instead of dedicated remote observing rooms. For even greater convenience, observing could take place

from observers' homes. For the ultimate in convenience, we could even allow observers to operate instruments from their laptop or other mobile device (e.g., smartphone).

All of these options may be possible to implement given sufficient effort by the observatory staff, and so the question for the observatory is where to draw the line. If observers are allowed to drive this decision, they will naturally select the option that is most convenient for them; thus, the observatory staff must carefully consider the implications of this choice. The right answer for a particular observatory will depend on how much effort the organization is willing to invest in making various observing modes possible.

Keck currently requires off-island observing to take place at the dedicated sites that have been established at Caltech, Swinburne, Yale, and the various University of California campuses. Our reasons are:

- Dedicated sites are guaranteed to be configured properly for running the instrument VNC sessions and communicating with the Keck video-conferencing units, and thus require no additional WMKO staff time. As noted above, any attempt to observe from a non-standard site will involve several hours of effort to set up and test the configuration.

- Dedicated sites are protected by battery backup systems to power computers and routers in case of (relatively brief) power outages, and by ISDN lines that provide an alternate communications channel to Keck if the Internet fails.

- Computers at the dedicated sites have been specifically tuned for peak performance (i.e., optimized TCP window size) over the long-latency connection to Hawaii, and are thus more responsive under VNC operation than a typical machine would be.

- Computers at dedicated sites are equipped with multiple monitors so that the observers can view all critical components of the instrument control interface at once, making it easier and more efficient to observe from these machines.

- Computers at dedicated sites are configured with enhanced security provisions to minimize the possibility of unauthorized access to Keck computers; usage of other devices could expose the Keck network to malicious intruders.

At present, we find these reasons sufficiently compelling that we plan to continue restricting off-island access to the established network of dedicated sites.

5. CASE STUDIES IN PROCESS AUTOMATION

In §4.2, we made the case for developing software to automate certain tasks which otherwise would require significant manpower to complete. In this section we present three examples of processes we've successfully automated to save time and effort.

5.1 VNC server manager

Under the Keck operations model, each observing team is assigned a particular computer account on the Keck network. This mode allows different observing teams to use the computers at the same time without interfering with each other. Similarly, we run separate sets of VNC servers for each observing team's use. Initially, the support astronomers at Keck were responsible for starting up the required VNC servers prior to the run and shutting them down when the run was completed. This arrangement worked adequately when mainland observing was used infrequently, but as usage rose this task became a burden to the staff.

To automate the process of starting and stopping VNC servers, we developed a script called `kvncmanager` which does the following:

1. reads the telescope schedule via an internal web service;

2. compiles a list of all instrument accounts which are scheduled for use within the next 5 days;

3. determines which of these accounts already have VNC servers running, and which do not;

4. launches VNC servers for those accounts which do not already have servers running;

5. checks to see whether an observing run ended the previous day, and if so, terminates any existing VNC servers for that account.

This script executes daily as a `cron` task and has virtually eliminated the need for staff to manage the instrument VNC servers.

5.2 Remote computer monitoring process

Keck's remote observing system involves over 30 different computers and routers distributed among a far-flung network of sites. A key requirement is to ensure that these devices are in working order at all times, just in case any given site is needed on short notice to serve remote observers. To monitor the status of both the local and remote computers essential to the system, we developed a script called `mainland_status` which uses the Unix `ping` utility to verify that these devices are operational and communicating with Keck. This task is scheduled to execute every morning via the Unix `cron` utility. When the `mainland_status` script is unable to successfully `ping` a remote device, the script sends email to the remote site manager as well as to the Keck Mainland Observing Coordinator, reporting the problem and prompting them to investigate the source of the failure. This has proven effective in identifying problems such as unresponsive computers, allowing us to fix the problem prior to an observing run.

5.3 Remote observing request tracking system

The process of handling remote observing requests has, until recently, been the dominant ongoing manpower expenditure in our mainland observing operation. Historically, mainland observing requests were initiated when an observer sent an email message to the Mainland Observing Coordinator at Keck. Each incoming request requires two types of checks:

- **Validation** is performed to determine that the observer's request is for the correct night, instrument, and site. This step traditionally required the Mainland Observing Coordinator to consult the telescope schedule after receiving each request in order to verify that the supplied information was correct.

- **Approval** of the request is required from (a) the manager of the requested mainland observing site and (b) the support astronomer assigned to the run. Consultation with the mainland site managers ensures that remote sites are not already scheduled for another use (e.g., remote observing at Lick Observatory[4]) on the night in question. Consent from the support astronomer is sought in order to verify that the observer is sufficiently experienced and capable of completing the observing run remotely.

Often, the information supplied in the observer's emailed request was incomplete, requiring a further exchange of email between the parties to obtain all necessary details. The requests for approval emailed to the mainland site managers and the support astronomers sometimes did not yield a timely response, requiring followup email exchanges to obtain approval. The manual method of tracking approval was not without fault and sometimes led to delays in the approval process when confirmation requests were not immediately acted upon. In sum, the manual approval method was expedient to set up but cumbersome and inefficient to operate.

To create a more robust and efficient system, we designed and developed an automated procedure for accepting and handling requests. The key feature of the new system is the Mainland Observing Request Form, a web-based form which the observer can access only after logging into his or her personal account on the WMKO public website. The form helps ensure that observers supply all of the information we need to process their requests.

Submitting the request via the web form is only the first of a number of steps that need to be completed before a request is formally approved. The newly-automated process of handling observing requests follows the same general procedure employed in the manual approval method described above. In outline form:

- The observer submits a request via the web form, invoking a CGI processing script that performs data validation, checking to see that:
  - all required fields have been completed;
  - the telescope, instrument, and principal investigator specified on the form agree with that listed on the telescope schedule;

– the requested mainland observing site is available on the specified dates.

If errors are found, the input is not accepted and the observer is prompted to revise and resubmit the form or contact WMKO staff for assistance.

- If the input is acceptable, the system creates a new database record and stores the input in the entry.

- The system sends an email to the WMKO Mainland Observing Coordinator seeking preliminary approval for the request.

- The Mainland Observing Coordinator receives the email and verifies that the request appears valid and that all required information was received. If not, he contacts the observer directly for clarification. For convenience, he can then grant preliminary approval in any of three ways: by sending a return email message; by clicking a hyperlink in the email message; or by logging into the web page for the request tracking system, locating the corresponding record, and clicking a button on a web form to grant approval.

- After preliminary approval is granted, the system consults the telescope schedule to determine which support astronomer is assigned to the night, then sends email requesting approval to the support astronomer and to the site manager at the specified mainland site(s).

- The support astronomer determines whether the request presents any conflict (e.g., the observer is not capable of operating the instrument remotely) and either grants or denies approval accordingly via one of the methods described above.

- The mainland site manager(s) consult their site calendar to verify that the remote observing room is not yet reserved for the specified night(s) and that no other conflicts (e.g., planned power outages) exist, then either grants or denies approval accordingly via one of the methods described above.

- If any of the required approvals is denied, then the system rejects the requested and sends email to the WMKO Mainland Observing Coordinator, who will then contact the observing team to discuss the reason the request was rejected.

- If all required approvals have been granted, then the system grants final approval and generates an email to that effect which is sent to the principal investigator, the person who submitted the mainland observing request, the support astronomer, and various support staff at WMKO and the remote sites.

In practice, this flow is augmented by an additional "watchdog" process that runs once daily and takes action if any of the following conditions are met:

- **Telescope schedule in conflict.** If the instrument and date on any approved or unresolved request no longer match the scheduled instrument for that date according to the WMKO telescope schedule (as can happen when the schedule is changed), then the process sends email to the WMKO Mainland Observing Coordinator alerting him to the problem.

- **Remote site overbooked.** Similarly, if the remote site has booked more mainland observing teams than it can accommodate (normally just one, but two at certain sites) then the process warns the Mainland Observing Coordinator via email so that he can resolve the conflict; for example, another site may be able to accommodate the remote observing team.

- **Approvals are unresolved.** If any of the required approvals are outstanding, the process sends email to the responsible parties prompting them to take appropriate action.

Whether it's because the process of granting approvals is now so much more convenient, or whether the automated reminders of outstanding requests are highly effective at generating a response, the typical time required to grant approval for a mainland observing request has decreased significantly since the new system was deployed: 75% of requests are now approved within 3 working days.

6. COST OF A REMOTE OBSERVING SYSTEM

To fully evaluate whether the numerous benefits of such a system will outweigh the costs, we need to know the expenses involved. There are both capital costs (to equip a site) and operational costs, and these apply to both the observatory and to the remote sites.

6.1 Capital expenditures

The main expense in setting up our mainland observing system was the effort to research, select, purchase and install the hardware, and to design, implement, and document the software that operates the system. We estimate that the total effort expended by the key developers over the 10-year lifetime of the project would amount to one person-year of time.

Because the remote observing rooms at WMKO headquarters were already equipped with the required computers and video-conferencing systems to support remote observing in Waimea, adding the capability of observing from off the island incurred little additional cost: the only capital expenses for WMKO were two workstations used as dedicated hosts for the VNC servers, plus a router capable of supporting six ISDN circuits. We estimate the cost at about $10,000, plus a minor amount of labor to set up these devices.

The cost to equip a mainland observing site depends on whether the site is designed to support just the "eavesdrop" mode only or the full "mainland-only" mode, as summarized in Table 1. A relatively simple eavesdrop-mode site can be assembled for $5,000–$15,000, depending primarily on the quality of the selected video-conferencing system. For mainland-only-mode operation, WMKO policy requires that the site be configured with additional equipment to safeguard against power or Internet failures, and these extras raise the basic cost of such a site to $15,000–$25,000. To these estimates should be added the manpower required to install, configure, and test the equipment, plus the cost of outfitting the room with required furniture.

Table 1. Estimated Cost to Equip a Mainland-Observing Site

Items	Eavesdrop Mode	Mainland-only Mode
Video-conferencing system	Y	Y
Computer	Y	Y
Monitors (4 large LCD)	Y	Y
Speakers	Y	Y
Uninterruptible Power Supply (Regular)	Y	
Spare computer		Y
ISDN router		Y
ISDN circuits		Y
Uninterruptible Power Supply (High capacity)		Y
Budget	**$5,000–$15,000**	**$15,000–$25,000**

6.2 Ongoing expenditures

At WMKO, the ongoing operational expense of the mainland observing system includes both the manpower required to coordinate the system and the rental cost of the six required ISDN circuits (about $3600/yr at present). Thanks to the numerous innovations described above that lessen the workload on the staff, we estimate that the operational burden on the Mainland Observing Coordinator amounts to about 2 hours per week. The support astronomers on staff are required to approve or deny mainland observing requests, but the effort required for this is negligible. Overall, the impact of operating the mainland observing system on the observatory itself is quite minor; i.e., < 0.1 FTE.

At the mainland locations, the site manager is responsible for approving or denying incoming mainland observing requests, tracking usage of the site to prevent conflicts, and performing minor support tasks such as instructing users in the operation of the computers and video-conferencing equipment, issuing keys and parking passes, etc. Even at the most heavily-used site (CIT), this amounts to under 10% of one staff person's time; for most sites, the effort is much lower. At eavesdrop-mode sites this is the only operational expense; mainland-only-mode sites incur the additional rental cost for three ISDN circuits at a typical cost of $1,800/yr. Thus, only a modest operational expense is involved at our mainland observing sites.

7. CONCLUSION

We have described in detail the challenges we've faced in minimizing the operating expense of the mainland observing system for WMKO and our strategies for doing so, as well as the benefits that can accrue from implementing this capability. Although the WMKO system is heavily used and required significant effort to develop, our ongoing efforts to automate the processes of handling remote usage requests and maintaining the required VNC servers have significantly increased the robustness and efficiency of operations while simultaneously reducing the required operational manpower to a very limited level. The current WMKO system shows that it is possible for a remote observing system to deliver substantial benefits at modest burden to the observatory.

ACKNOWLEDGMENTS

REFERENCES

[1] Kibrick, R. I., Wirth, G. D., Allen, S. L., Deich, W. T. S., Goodrich, R. W., Lanclos, K., and Lyke, J. E., "Operating a wide-area high-availability collaborative remote observing system for classically-scheduled observations at the W. M. Keck Observatory," in [*Telescopes from Afar*], Gajadhar, S., ed. (2011).

[2] Kibrick, R. I., Hayes, B., Allen, S. L., and Conrad, A., "Remote observing with the Keck Telescopes from multiple sites in California," in [*Advanced Global Communications Technologies for Astronomy II*], Kibrick, R. I., ed., *Proc. SPIE* **4845**, 80–93 (2002).

[3] Wirth, G. D., Kibrick, R. I., Goodrich, R. W., and Lyke, J. E., "Operating a wide-area remote observing system for the W. M. Keck Observatory," in [*Observatory Operations: Strategies, Processes, and Systems*], Brissenden, R. J. and Silva, D. R., eds., *Proc. SPIE* **7016** (2008).

[4] Kibrick, R. I., Wirth, G. D., Gates, E. L., Grigsby, B. J., Deich, W. T. S., Lanclos, K., and Allen, S. L., "A shared approach to supporting remote observing for multiple observatories," in [*Observatory Operations: Strategies, Processes, and Systems III*], Silva, D. R., Peck, A. B., and Soifer, B. T., eds., *Proc. SPIE* **7737** (2010).