

Operating a wide-area high-availability collaborative remote observing system for classically-scheduled observations at the W. M. Keck Observatory

Robert I. Kibrick^a, Gregory D. Wirth^b, Steven L. Allen^a, William T. S. Deich^a, Robert W. Goodrich^b, Kyle Lanclos^a, and James E. Lyke^b

^aUCO/Lick Observatory, University of California, Santa Cruz, CA 95064 USA

^bW. M. Keck Observatory, Kamuela, HI 96743 USA

ABSTRACT

For over a decade, the W. M. Keck Observatory’s two 10-meter telescopes have been operated remotely from its Waimea headquarters. Over the last 9 years, WMKO remote observing has expanded to allow observing teams at dedicated sites located across California to observe via the Internet either in collaboration with colleagues in Waimea or entirely from California; this capability was extended to Swinburne University in Melbourne, Australia in 2010 and to Yale University in New Haven, Connecticut in early 2011.

All Keck facility science instruments are currently supported. Observers distributed between as many as four sites can collaborate in the interactive operation of each instrument by means of shared VNC desktops and multipoint video and/or telephone conferencing. Automated routers at primary remote observing sites ensure continued connectivity during Internet outages.

Each Keck remote observing facility is similarly equipped and configured so observers have the same operating environment. This architecture provides observers the flexibility to conduct observations from the location best suited to their needs and to adapt to last-minute changes. It also enhances the ability of off-site technical staff to provide remote support.

Keywords: remote observing, observation scheduling

1. INTRODUCTION

To understand why WMKO uses classical scheduling and remote observing, it is useful to compare the alternatives. As illustrated in Table 1, observatories can select between different models to schedule telescope time and different modes to determine who (or what) conducts the observations on their telescopes. For example, under a classical scheduling model, a time-assignment committee ranks the submitted observing proposals and assigns to each highly-rated proposal blocks of telescope time; the schedule maintainer then assigns those blocks to specific dates within the telescope schedule for the upcoming semester. On those dates, the observers associated with that proposal would conduct the relevant observations, working either at the observatory or remotely from their home institutions. Alternatively, they might have a service observer at the observatory conduct those observations on their behalf.

Table 1. Scheduling models and observing modes, shown as a tabular diagram of sets, related vertically

Factor	Observing Modes			
	Classical Scheduling		Queue Scheduling	
Scheduling Models	Regular Observers		Service Observer	Automaton
Who conducts observations	Regular Observers		Service Observer	Automaton
Where observations are conducted	At home institution	“At the observatory”		Cyberspace
Where observers sleep	At home	In dorm	Either	N/A

Further author information- (Send correspondence to R.K.)

R.K.: Email: kibrick@ucolick.org; www.ucolick.org/~kibrick; Telephone: 831-459-2262; G.W.: Email: wirth@keck.hawaii.edu; www2.keck.hawaii.edu/realpublic/inst/people/wirth/; Telephone: 808-881-3866

During the last decade, a growing number of observatories have implemented a queue scheduling model. Under this model, a time-assignment committee ranks observing proposals, and each proposal receiving a sufficiently-high rank is expanded into a detailed set of observation requests that are then placed into a rank-ordered observing queue. Each night, observation requests are selected from that queue via a process that attempts to globally optimize various criteria; examples of such criteria are: matching the requirements of a specific observing request to the current sky conditions, giving priority to observing requests from higher-ranked proposals or to those requests that are time-critical, and ensuring a proper balance of observing time is maintained between each of the member institutions eligible for time on the given telescope. Typically, each observation request selected from the queue is carried out by an on-site service observer who is a member of the observatory staff rather than by the astronomers associated with the observing proposal which generated that request. However, in the case of robotic telescopes and observatories, the observation requests selected from the queue are carried out by an automaton.

Thus, in the case of queue-scheduled service observing or robotic observing, the astronomers responsible for a given observing proposal are no longer directly involved in conducting the observations defined by that proposal. As a result, these observing modes do not require those astronomers to be “at the observatory” (either physically or virtually) when those observations are carried out. Accordingly, these observing modes obviate the need for the remote observing capabilities that are the subject of this paper. In addition, these modes potentially can provide several advantages, including: more efficient overall use of telescope time, greater likelihood that observations requiring exceptional sky conditions will be carried out at the optimal time, greater likelihood that highly-ranked programs will be completed, support for programs that require only a small fraction of a night, and greater flexibility in responding to time-critical or transient events (e.g., GRBs). It is these advantages that are driving a growing number of observatories to support these scheduling models and observing modes.

Yet despite these trends, the Keck Observatory currently does not support (and has no plans to support) either queue scheduling, service observing, or robotic observing. Rather, all observing time on the Keck Telescopes is classically scheduled and all observations are carried out directly by the astronomers associated with each approved observing proposal.

There are several reasons for this. First, due to the mechanical design of the Keck Telescopes, in most cases it is not possible to change instruments on either telescope in the middle of the night, and that greatly reduces many of the benefits of queue scheduling.¹ Second, two of the most heavily used Keck instruments, LRIS on Keck I and DEIMOS on Keck II, routinely require manual reconfiguration of the instrument during the daytime so as to enable observers to select (from a fairly broad palette) the specific complement of gratings, filters, and slit masks required for a given observing program. In fact, the multi-slit masks loaded into these two instruments each day are usually specific to a particular target field that is part of a single observing program. The need for such program-specific manual reconfigurations of these instruments each day would make queue scheduling such observing programs very problematic. Third, the overwhelming majority of current Keck observers have indicated a strong preference to maintain classical scheduling of Keck Telescope time. Fourth, despite whatever inefficiencies might be attributed to classical scheduling, the Keck Telescopes remain extremely productive in terms of the papers published per telescope each year and the overall scientific impact of those papers.² Thus, it is not at all clear that adopting a queue scheduling model at WMKO would yield significant benefits nor that whatever benefits did accrue would justify the very large costs of implementing such a model.

Accordingly, the choice by WMKO to use classical scheduling and to not provide service observing effectively requires that the observers associated with a given observing program must be able to perform the observations for that program on the specific dates and times assigned to that program in the telescope schedule. This entails providing those observers with reliable and efficient real-time access to the relevant telescope and instrument control interfaces, either at the Observatory or at a remote location. The remainder of this paper describes the steps we have taken to provide that capability.

2. OBSERVING AT KECK OBSERVATORY

Between 1993 (when Keck I came online) and 1995, Keck observers conducted their observations from nearby control rooms on the Mauna Kea summit. Since 1996, most observations with the Keck telescopes have been carried out from one of the two remote operations rooms located at WMKO Headquarters in Waimea, located about 32 km from the summit; as of 2008, observing from the Keck control rooms on the summit is no longer supported. The primary motivation for this was to move observers off of the oxygen-starved summit and enable them to observe with greater safety and efficiency from a nearby facility located at a lower altitude.³ The Keck I remote operations room in Waimea became operational in 1996, and the corresponding room for Keck II (see Fig. 1) came online the following year.

The WMKO Headquarters also provides visiting scientist quarters (VSQ) where users of the remote operations rooms are provided a dark and quiet location to sleep during the daytime. Observers can prepare meals in a well-equipped kitchen in the VSQ or dine at local restaurants, several of which are within walking distance to the headquarters building.

The observer is supported by an observing assistant (OA) who operates the telescope, and a support astronomer (SA) who assists the observer in setting up and operating the instrument. During the first part of the night, an SA is present in the remote operations room and during the latter part an on-call SA can be reached by telephone at home. The OA for each Keck Telescope is usually located at the summit, but in some cases one OA (but not both) will operate one telescope from the same control room in Waimea from which the observer is running the instrument. A video-conferencing system links each remote operations room in Waimea with its corresponding telescope control room at the summit. Instrument control software runs at the summit with data written to summit disks. A periodic *rsync* copies data to disks in Waimea.



Figure 1. Remote operations room at Keck HQ in Waimea



Figure 2. Remote observing room at UCSC




Figure 3. Global overview of connectivity between mainland remote observing site, Mauna Kea, and Waimea

3. OBSERVING FROM THE OBSERVER'S HOME INSTITUTION

Efforts to enable Keck remote observing from sites in California were undertaken by researchers at the University of California, Santa Cruz (UCSC)⁴ and the California Institute of Technology (CIT).⁵ A key motivation for those efforts was to reduce the travel time and costs associated with short-duration observing runs.⁶ The remote observing facility at UCSC (see Fig. 2) became fully operational in 2001, and served as the model for similar facilities that came online at other UC campuses, at CIT, and at Swinburne University and Yale University during the following ten years (see Table 2).

Table 2. WMKO remote observing sites (UC-affiliated sites listed in bold)

Number	Name	Rooms	Resolution	ISDN backup	Location	First use
1	WMKO HQ	2	1280 × 1024	no	Waimea, Hawaii	1996
2	UCSC	1	1920 × 1200	yes	Santa Cruz, California	2001
3	UCSD	1	1280 × 1024	yes	San Diego, California	2003
4	CIT	2	1280 × 1024	yes	Pasadena, California	2004
5	LBNL	1	1920 × 1200	yes	Berkeley, California	2005
6	UCLA	1	1600 × 1200	yes	Los Angeles, California	2006
7	UCSB	1	1600 × 1200	no	Santa Barbara, California	2007
8	UCB	2	1920 × 1200	yes	Berkeley, California	2007
9	UCR	1	1600 × 1200	no	Riverside, California	2008
10	UCD	1	1600 × 1200	no	Davis, California	2008
11	UCI	1	1600 × 1200	no	Irvine, California	2008
12	Swinburne	1	1920 × 1200	no	Melbourne, Australia	2010
13	Yale	1	1920 × 1200	yes	New Haven, Connecticut	2011

The external Keck remote observing sites are primarily targeted towards observers who live within commuting distance of one of those facilities;^{6,7} no dormitories for observers are provided at these locations. They are not intended to duplicate the Waimea facility nor to operate independently of it. Rather, each is an extension of the Waimea facility. That site and those on the mainland are intended to operate in collaboration, sharing resources where practical. We rely on the existing instrument support staff in Waimea and provide video-conferencing and shared software environments so that they can most effectively support the observers at the mainland sites.

Communication between the sites is provided via research class (e.g., INET2) networks (see Fig. 3 and the Acknowledgments section). However, many mainland remote observing sites are also equipped with a series of dialed ISDN circuits (see column 5 of Table 2) that provide an alternate link to the Mauna Kea summit. In the event of a loss of Internet connectivity, remote observing traffic automatically fails over onto this ISDN backup path, and automatically falls back once Internet connectivity is restored.⁸

4. DEALING WITH INTERNET CONNECTIVITY PROBLEMS

Most of the time, the research-grade Internet paths between the Keck Telescope site and the mainland remote observing sites provide excellent bandwidth, very low rates of packet loss, and stable packet latency (see Fig. 4). Unfortunately, such excellent connectivity along these paths is not guaranteed. While some of the path segments provide redundancy, there are still many potential points of failure along these paths that can result in either a complete loss of connectivity or in seriously degraded network performance between these sites (see Fig. 5); such Internet problems, if not mitigated, can seriously disrupt mainland remote observing. Such episodes of degraded performance or lost connectivity are typically observed every few months. While some of these episodes last only a few minutes, a significant number last more than 2 hours, and in a few cases, degraded performance has persisted for more than a day (see Fig. 6).




Figure 4. Typical latency and packet loss rates along Internet path between U.C. Santa Cruz and the Keck Telescopes, as illustrated in this plot produced using the *smokeping* utility (see <http://oss.oetiker.ch/smokeping>). Most of the time, this path delivers extremely low rates of packet loss and very stable packet latency of about 60 milliseconds.




Figure 5. Increased latency and packet loss along the same path on September 2, 2009 caused by problems with an intermediate router in Hawaii. Due to the exponential-backoff algorithms that are built into the TCP/IP network protocols, a relatively small increase in packet loss rates results in a large (and often variable) increase in interactive response times, which in turn significantly decreases observing efficiency. The disruption shown here lasted over 2 hours, which is not unusual for such incidents.




Figure 6. Variation in latency and packet loss rates along same path over the past 12 months (March 2010 through February 2011), showing episodes of degraded performance every few months, with a major disruption persisting over 2 days in September 2010 and two significant disruptions in December 2010.

Given the very high value of Keck Telescope time (\$100K/night/telescope, as valued by the NASA TSIP program as implemented at WMKO), we need to ensure that the telescope does not sit idle because an observer working from a mainland site experiences Internet connectivity problems. To address this concern, WMKO established a formal policy* that defines two distinct modes of remote observing:

1. **Remote eavesdropping mode** (permitted from all external sites)

- At least one member of the observing team observes from Waimea, while the other members of the team observe from the mainland.
- Observer(s) in Waimea have primary responsibility for instrument operation, but observers on the mainland are able to operate the instrument if desired.

2. **Mainland-only mode** (permitted only from external sites with ISDN backup communications path)

- All members of the observing team observe from mainland site(s) (and at least one team member must have previous experience operating the scheduled instrument).
- Mainland observer(s) have sole responsibility for instrument operation.

To enable this second mode of operation, dialed ISDN telephone circuits and ISDN routers were installed at the Keck Telescope site and at 7 of the 12 mainland remote observing sites as shown in Table 2; these circuits provide an alternate route that can be activated during those periods when Internet connectivity is degraded or interrupted. Although this alternate route provides much lower bandwidth and somewhat higher latency than the normal Internet path, that is still preferable to a degraded Internet path that provides very unstable bandwidth and latency or no bandwidth at all.

Six of the 7 mainland sites equipped with ISDN circuits account for slightly more than 90% of the mainland remote observing activity (see Fig. 10). Each of those 7 sites has 3 ISDN lines (or 6 B-channels) and there are 6 ISDN lines at the Keck Telescope site, thus permitting simultaneous ISDN operation between Keck and two ISDN-equipped mainland sites. Each site pays its own ISDN lease charges (typically \$60/line/month) and toll charges (typically \$0.10/B-channel/minute).

In the case of a complete loss of connectivity between a mainland site and Keck, the ISDN routers at the two endpoints quickly detect the outage and automatically re-route the remote observing traffic onto the ISDN circuits in less than one minute. Unfortunately, these routers are not capable of detecting when such connectivity is badly degraded but not severed (i.e., when many of the packets being sent along the path between the sites are being dropped or severely delayed, but a small fraction of packets is still being delivered). In such cases, the ISDN routers would normally continue to route the remote observing traffic via this degraded Internet path. To remedy such situations, it is advisable for the support astronomer (SA) at Keck to intervene by issuing a command to the ISDN router at the Mauna Kea endpoint, directing it to re-route the remote observing traffic between the sites onto the appropriate ISDN circuits. Once the Internet problem has been resolved, the SA would command the router to redirect traffic back onto the Internet path. To assist the SA in making such decisions, WMKO employs a PC-based product, Orion Network Performance Monitor by SolarWinds[†], to continuously monitor connectivity, latency, and packet loss to all external remote observing workstations and routers.

5. ENSURING HIGH AVAILABILITY AND EFFICIENT OPERATION

In addition to providing an alternate route via ISDN circuits, sites that support the mainland-only mode of remote observing need to protect against other risks that could potentially interfere with remote observing, including loss of utility power or other problems that could degrade or disrupt the operation of the local workstation from which the remote observer is working. To protect against the former, such sites have all critical equipment (including the ISDN router and ISDN network termination unit) powered from a large uninterruptible power supply (UPS) sized to provide sufficient runtime to keep this equipment powered through the longest outages typical for each site. To protect against the latter, the workstations used at the remote observing sites are configured to operate standalone; they must be able to boot and operate without dependence on any outside servers (e.g., NFS, NIS) and to provide DNS fallback to the Keck DNS server in the case where no local DNS servers can be reached. Spare workstations and monitors are provided at each site in case of hardware problems.

*http://www2.keck.hawaii.edu/inst/mainland_observing/policy.html

†<http://www.solarwinds.com/products/orion/>

In addition, the UNIX workstations at the mainland remote observing sites are specially configured to provide efficient performance for remote observers. First, the TCP window size parameters in the system kernel[‡] are tuned to provide optimal network throughput across the high-bandwidth long-latency trans-Pacific networks. Second, various system housekeeping functions (e.g., network-based disk backups or disk space cleanup, both of which can seriously degrade the system's interactive performance) that UNIX systems typically run during the middle of the night (on the assumption that users are asleep) are instead scheduled to run during the daytime (HST) so as to not impact the workstation's performance during the hours when remote observing is likely to be in progress.

Keck observatory also provides daily monitoring of the network connectivity between its telescopes and the various remote observing sites; this effectively checks both the status of the network to and the status of the equipment at each site. A *mainland_status* script is invoked on a host in Waimea each morning via the Unix *cron* facility. That script uses the Unix *ping* utility to verify that all of the relevant ISDN routers and remote observing workstations at each of the mainland remote observing sites are reachable from Waimea via the network. If any are not, the script sends an email message to the Keck remote observing coordinator⁷ and to the manager at the relevant site to alert them of a potential problem. Since some remote observing hosts often go unused for many days at a time, such automated checking helps to provide timely detection of problems that might otherwise go unnoticed until a remote observing session is just about to commence.

6. ALTERNATIVE APPROACHES

An alternative approach that has been proposed is to make Keck remote observing available directly from the workstations that most observers have in their offices; this approach has been implemented successfully by other observatories, such as the NASA Infrared Telescope Facility (IRTF)[§] on Mauna Kea. While this approach may work well for operating simple instruments (e.g., photometers) whose user interfaces and displays don't require significant amounts of screen real estate, this model can be problematic when operating more complex instruments that require more screen real estate than is typically found on a user's workstation, such as the existing instruments in operation at Keck. While video-conferencing can be provided via the user's workstation using inexpensive webcams, these typically don't provide the same capabilities (e.g., remote control of camera pan/zoom, high quality video and audio echo cancellation, multi-conferencing capability, etc.) as provided by the dedicated video-conferencing systems used in our existing Keck remote observing rooms.

In addition, while the remotely-observe-from-anywhere model might provide greater convenience for some observers, it would require an unpredictable level of support from WMKO to ensure that the remote observer's workstation is properly configured. First, there is no way to ensure that an individual user's workstation is running a current version of its operating system and/or web browser (or a even web browser that is compatible with the remote observing applications) and that critical security patches have been applied. Second, an observer may have configured his or her workstation to have non-standard keyboard or mouse button mappings[¶] or unusual settings of browser and/or window manager properties, and these may cause subtle operational problems that are difficult for the support astronomers at Keck to diagnose, since they have no way to reproduce the remote observer's local operating environment. Third, most observers' private workstations have neither the level of UPS backup power that is typically provided to the equipment in our Keck external remote observing facilities, nor access to a backup, ISDN-based network path as provided at many of those facilities (see column 5 of Table 2). Fourth, such workstations typically have not been configured to prevent system housekeeping functions from running during the nighttime hours nor have their TCP window size parameters been optimized for the high-bandwidth long-latency trans-Pacific links. Fifth, most such workstations aren't capable of running standalone and are dependent on server machines that could go down. Finally, if remote observers can connect from arbitrary sites around the world, it is much more difficult to provide the level of routine connectivity monitoring (see Section 5) as we currently provide to our existing Keck remote observing facilities.

Accordingly, WMKO has no plans to support the mainland-only mode of remote observing from the workstations that most observers have in their offices. However, in exceptional cases, WMKO may support the eavesdrop mode from such workstations.

[‡]<http://www.psc.edu/networking/projects/tcptune/>

[§]http://irtfweb.ifa.hawaii.edu/observing/remote_obs/

[¶]Such non-standard key mappings can sometimes cause vulnerable software to go awry in unpredictable ways that are difficult to troubleshoot, resulting in a significant loss of observing time. An unfortunate example of such a loss that was ultimately found to have been triggered by unexpected key mappings is described at the start of an article in the August 27, 2006 edition of Time Magazine, *How the Stars Were Born*, available at <http://www.time.com/time/magazine/article/0,9171,1376229,00.html>

7. DESIGN GOALS AND IMPLEMENTATION STRATEGY

Key design goals for the Keck mainland remote observing software are that it should enable observing from mainland sites that is as reliable, secure, and efficient as observing from Waimea. In addition, it should be simple to operate and provide equivalent capabilities using the existing Keck instrument software interfaces. The software should present observers with the same “look and feel” regardless of which Keck remote observing site they are using, and it should enable multi-site shared control of the instrument.

Because there are no staff at any of the sites who are assigned full time either to develop or maintain the Keck remote observing software, we needed to minimize the volume of software that needs to be maintained at the remote sites. Accordingly, we kept the overall system relatively simple by minimizing external dependencies and maximizing the use of off-the-shelf components. Because we were working on a limited budget, we used open source and free software wherever possible. To avoid unintended interactions between the remote observing software and existing legacy instrument control software and to avoid adding additional loads onto legacy instrument control hosts, we needed to ensure a clean separation between those systems.

Several off-the-shelf software components provide the foundation atop which our thin layer of application software sits: (1) The UNIX operating system enables our system to run on different hardware architectures such as Sun/SPARC or Intel/X86. (2) *Ssh* provides secure connections between the sites. (3) The *FVWM* window manager is supported under both Solaris and Linux and is used both inside and outside of the VNC desktops. (4) VNC, or virtual network computing, provides shared virtual desktops and is the key ingredient to providing an effective and efficient collaborative framework.

8. IMPLEMENTATION DETAILS

8.1 VNC shared virtual desktops

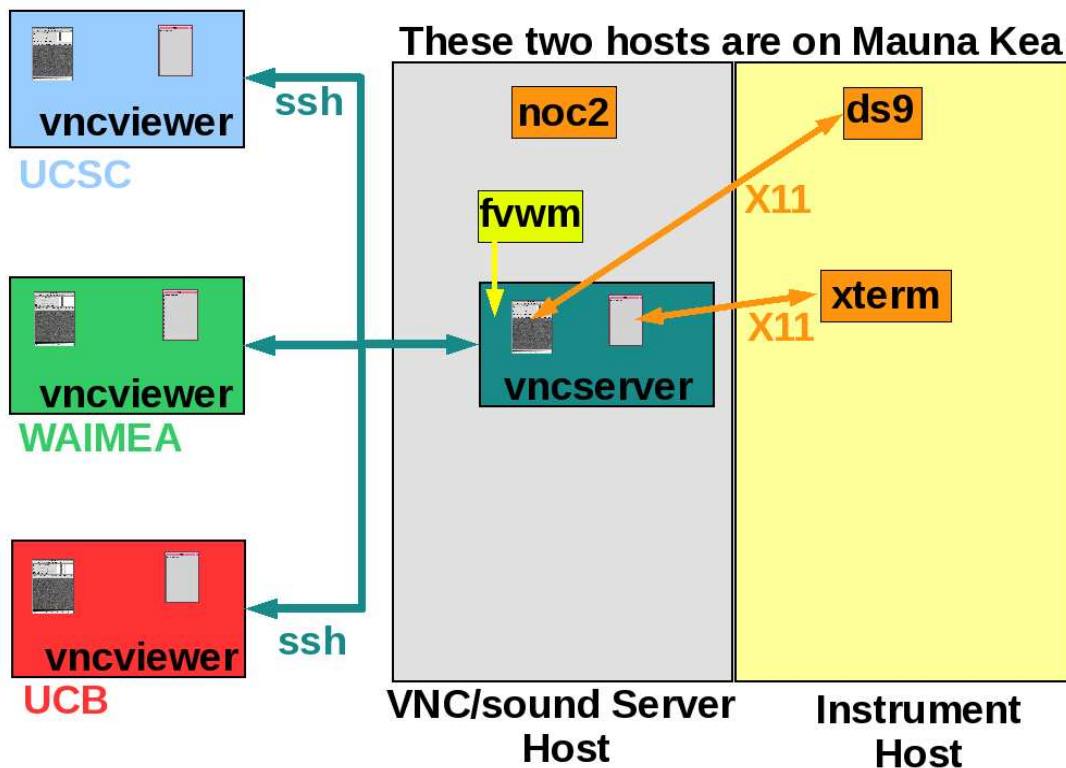


Figure 7. Overall architecture of shared VNC desktops. VNC servers run on a dedicated host at the summit and all connections from remote observing sites are via *ssh* tunnels to that host. *X*-based GUIs and image displays that run on an instrument control host redirect their displays to a VNC server running on that server host. VNC viewers running at the remote sites all see and interact with the identical shared VNC desktop. *FVWM* is the window manager for each VNC desktop. *Noc2* monitors and logs all connections to all VNC servers.

VNC operates on a client/server model (see Fig. 7). A VNC server (which in our case runs on a dedicated VNC/sound server host at the Keck Telescope site) acts as a virtual *X* server that is not directly connected to any physical display hardware. *X*-based client applications (e.g., GUIs, *xterms*, and image displays such as *ds9*) running on instrument control hosts connected to the Keck summit LAN can direct their displays to such a VNC server in the same way that they would to any other *X* server and without incurring any added overhead for doing so.

VNC viewer clients running on workstations in Waimea or at the mainland remote observing sites connect (via *ssh* tunnels) to the desired VNC server running on the dedicated server host at the Keck Telescope site. That VNC server uses the VNC protocol to relay to each of the connected VNC viewer clients at the remote sites those display streams that were directed to it by the *X*-based client applications running on the instrument control hosts. Those VNC viewer clients in turn direct their displays to the local *X* server which drives the screen of one of the LCD monitors in front of which the remote observer is seated. Thus, the VNC server host carries this multiplexing/relay load rather than the instrument control hosts.

In this way, whatever graphical output is generated by the *X*-based client applications (running on the Keck instrument control hosts) is rendered identically on the screens of the corresponding LCD monitors at each of the Keck remote observing sites that are connected to that VNC server, all without increasing the load on those control hosts. In similar fashion, each VNC viewer uses VNC protocol to send to that VNC server any keystrokes, mouse motions, or mouse clicks that it receives when it has window focus. Thus, each site has a co-equal control path (similar to the dual controls in an airplane with a pilot and co-pilot) back to the client applications running on the instrument control hosts on Mauna Kea.

VNC provides several key capabilities and advantages. Its shared virtual desktops ensure that all sites participating in a remote observing run have an identical and simultaneous view of the state of the instrument and telescope. In addition, all participating sites have the option of either interacting with the instrument GUIs and image displays or simply watching the actions taken by their colleagues at other sites, since all sites can see each other's keystrokes and mouse motions in real-time. (The videoconferencing connections between sites provide a simple and efficient means for the observers at each site to feel as if they were all working in the same room and to verbally arbitrate who does what and when.)

An additional advantage of VNC is that compared to *X* protocol, the VNC protocol delivers markedly superior interactive performance over long-latency links, such as the trans-Pacific links that connect the mainland remote observing sites to Keck.⁹ A further advantage of VNC is that session context is preserved at server end (i.e., within the VNC servers that run at the Mauna Kea summit), so that: (1) the remote sites can connect or disconnect from the session in any sequence without impacting each other, and (2) a dropped connection from a remote site does not kill the overall session.

8.2 Mapping VNC virtual desktops to workspaces on physical screens

As seen in Figs. 1 and 2, the workstations at each of the Keck remote observing facilities are equipped with multiple physical screens, each of which is treated as a separately-addressed logical screen. The *X* sessions hosted by these workstations are configured so that each physical screen is served by a separate local *X* server. For example, on a workstation with 3 physical screens, the *X* server for the leftmost screen is accessed as display :0.0, the middle screen is :0.1, and the rightmost screen is :0.2. A given application can be launched on a specific screen by first setting the *X* DISPLAY environment variable to address the corresponding *X* server; for example, to launch an application on the middle screen, the DISPLAY environment variable would first be set to :0.1. Because each screen has its own *X* server, it also has its own separate set of color maps. Similarly, if a pop-up window or pulldown menu grabs exclusive window focus, that focus grab only impacts a single screen, enabling the other screens to continue their independent updates.

Associated with each physical screen is a distinct pair of named VNC virtual desktops, one of which is a tan 8-bit VNC desktop and the other a blue 24-bit VNC desktop (the different background colors visually highlight which is which). A separate VNC viewer client is used to display each of these desktops and each viewer is displayed within a separate virtual workspace that is managed by the *FVWM* running within context of the local *X* server for that screen. Providing both 8- and 24-bit VNC virtual desktops enables the use of a wide range of *X*-based client applications. The 8-bit desktops enable the use of legacy *X* clients that require an *X* server supporting 8-bit *X* visuals; they also provide superior interactive response when adjusting the image contrast and bias for image display clients such as *ds9*. The 24-bit VNC virtual desktops are used to support applications (e.g., browsers) that require a much deeper color map than is provided by the 8-bit desktops.

Each VNC virtual desktop has a distinct name that reflects its function and which controls the workspace and physical screen position on which it normally appears. Typically, the 8-bit VNC desktops associated with the instrument control interfaces are named *controlN* and the 24-bit VNC desktops associated with image analysis operations are named *analysisN*, where N is 0, 1, or 2 depending on whether the desktop is to appear on screen :0.0, :0.1, or :0.2.

The configuration file for the *FVWM* window manager defines the number of virtual workspaces that *FVWM* will create and manage on each screen. Typically, we define 4 virtual workspaces per screen; two of the workspaces are used for displaying the two shared VNC desktops (i.e., the 8-bit and 24-bit desktops) for each screen and the other two are available for use as local desktops. The configuration file also provides a method for mapping VNC virtual desktop names to a specific *FVWM* workspace. Accordingly, when the VNC viewer client that will display a given named VNC desktop is started, *FVWM* will launch that viewer within the corresponding workspace defined in its configuration file. For example, the *controlN* desktops are started in *FVWM* workspace (1,0) on each screen, while the *analysisN* desktops are started in workspace (1,1); workspaces (0,0) and (0,1) are available as local desktops. *FVWM* also provides a workspace selection GUI that resides in the corner of each workspace and allows the user to switch between workspaces with a single click.

Together, these addressing methods enable our remote observing software to ensure that a given named VNC virtual desktop will be displayed within the same virtual workspace and physical screen position at all of the remote observing sites that are participating in a given observing run, thus ensuring that the observers at all of those sites have a consistent arrangement of shared VNC virtual desktops.

The use of multiple distinct VNC desktops improves overall interactive response, because the VNC viewer that renders each desktop connects to its corresponding VNC server via a separate *ssh* tunnel, enabling multiple TCP windows across the network to be in transit in parallel; it also enables observing team members located at different sites to simultaneously work on separate tasks in separate shared virtual VNC desktops, yet still have visibility into each other's actions.

8.3 Sharing audio alerts

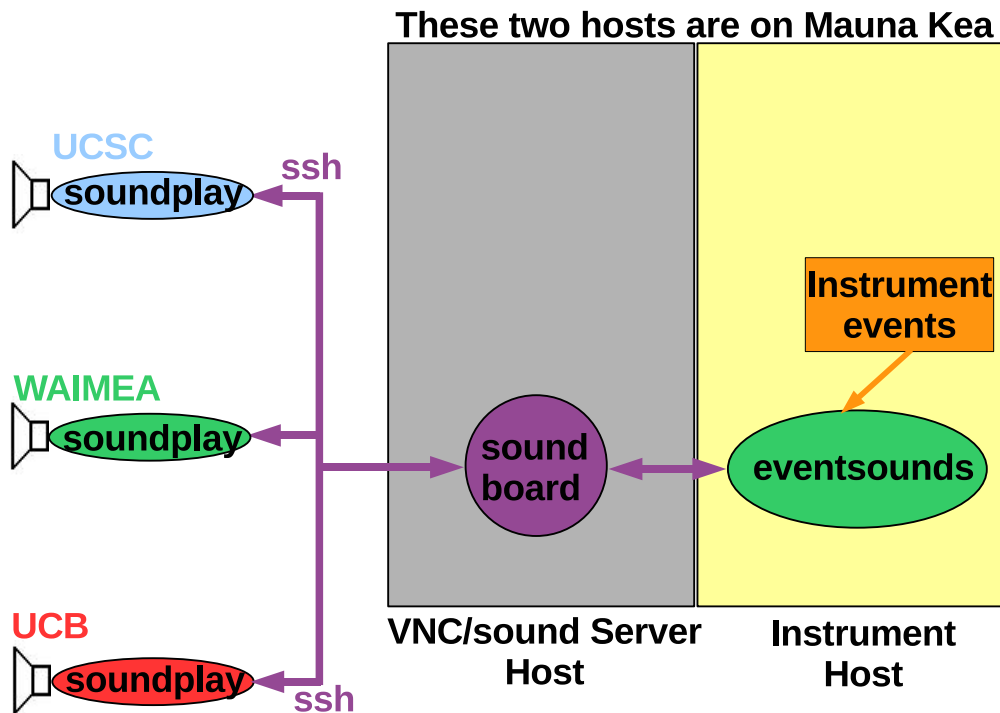


Figure 8. Overall architecture of for distributing audio alert sounds. *Soundboard* sound server runs on the same dedicated host as the VNC servers and all connections to the *soundboard* from remote observing sites are via *ssh* tunnels to that host. *Eventsounds* GUI running on an instrument host maps instrument events (e.g., *exposure completed*) into messages that request *soundboard* to transmit a sound file request to the connected clients. *Soundplay* clients at the remote sites play the sound files relayed by *soundboard*.

While VNC enables virtual desktops to be shared between the remote observing sites, it does not provide a means of sharing audio alerts. Since many of the Keck instruments enable audio alert sounds to be associated with various instrument

events, we implemented a set of *Tcl* programs (see Fig. 8) that operate similar to VNC; these programs enable such alert sounds to be shared between sites. To simplify installation of these programs, they are packaged as standalone *Tcl* starksits and thus do not require a full *Tcl* installation on the remote observing workstation.

A *soundboard* server functions like a VNC server and is used to multiplex sounds to the remote sites; it runs on the same summit host as the VNC servers. *Soundplay* clients function like VNC viewers and are used to play the sounds multiplexed by the soundboard; these clients run on the workstations at the remote sites and connect to the *soundboard* via *ssh* tunnels. An *eventsounds* GUI that runs on the instrument control hosts listens for instrument events and maps these into messages that it sends to the *soundboard* to request that it broadcast a particular audio alert sound. Note that *eventsounds* is the only piece of remote observing software that we install on the instrument control hosts.

8.4 Software needed at Keck

The off-the-shelf software needed at Keck to support this remote observing architecture includes Unix (currently Solaris), *X*, *ssh*, *super* (similar to *sudo*), *cs*, *Perl*, *expect*, *tklogger*, *FVWM*, and *RealVNC*. Layered on top of that are several custom software suites that manage the operation of remote observing sessions and monitor/schedule those sessions:

The *kvnc* suite is a set of *Perl* scripts used to manage the VNC servers that run on the VNC/sound server host at the Keck summit and the VNC viewers that run on the remote observing workstations in Waimea. These scripts consult a master *kvnc* configuration file for each instrument which specifies the number and types of VNC desktops (as well the screen on which each should appear) that are needed by each instrument; these scripts do all of the bookkeeping and addressing needed to ensure that the correct types of VNC desktops appear in the correct positions and run within the correct instrument account context. This suite also includes commands that enable support astronomers to check the status of running VNC processes.

The *noc2* suite is a set of *cs* scripts used to monitor, log, and alert remote observers and support astronomers regarding which remote observing sites are connected to which shared VNC desktops. Since all remote sites establish connections to the VNC/sound server host, it provides a centralized point for monitoring such connections. Whenever a new connection is established (or an existing connection is terminated) to a VNC server running on that central host: 1) a time-stamped log entry is recorded in a global log file and, 2) a notification message is generated in a window that pop-ups on the shared VNC desktop served by that server so that all users of that desktop will be alerted; that message lists the identity of the site that just connected to (or disconnected from) that VNC desktop and also lists the connection status of any other sites that have had connections to that virtual desktop since it was created. This suite also includes commands that enable support astronomers to view a list of all of the active VNC connections.

The *soundboard*, *soundplay*, and *eventsounds* suite of *Tcl* scripts manage the distribution of audio alert sounds to the remote observing sites, as described above. Each is packaged as a *Tcl* starksit and can thus be installed as a standalone program that does not require an underlying *Tcl* installation. This is especially important for the *eventsounds* GUI, which installs on each of the instrument control hosts, since several of those hosts are quite old and run older Solaris versions.

Additional software at Keck helps to automate the scheduling of remote observing requests and the startup and shut-down of the relevant VNC servers based on the telescope and remote observing schedules.¹⁰ As noted in Section 5, daily *cron* jobs monitor the status of the remote observing sites. Work continues on improving the software tools for dynamic monitoring of Internet connectivity between Keck and the remote observing sites and to provide support astronomers with more convenient tools to switch remote observing traffic onto the ISDN backup path during periods of degraded Internet connectivity.

8.5 Software needed at the mainland remote observing sites

The off-the-shelf software needed at the mainland remote observing sites includes Unix (Solaris or Linux), *X*, *ssh*, *cs*, *expect*, *FVWM*, and *RealVNC*. Layered on top of that is a software suite that manages the operation of the remote observing sessions:

The *kvncinst/tel* suite is a set of *cs* scripts used to manage the *soundplayer* and the VNC viewers that run on the workstations at the mainland remote observing sites. These scripts do all of the bookkeeping and addressing needed to ensure that the correct types of VNC desktops are correctly rendered and appear in the correct positions on the workspaces and screens at the remote site; they also maintain a log of relevant session events. The remote observer simply invokes these scripts with the name of the instrument and account for which the remote observing session should be established. The user is then prompted to supply several distinct passwords in order to complete the connection.

The *soundplay* Tcl script receives the audio alert sounds broadcast by the *soundboard* that runs at Keck and plays those sounds on the speaker attached to the local workstation. It is packaged as a *Tcl* starker and can thus be installed as a standalone program that does not require an underlying *Tcl* installation.

9. OTHER CONSIDERATIONS

9.1 Standardization between sites

A key factor in the reliability and maintainability of the Keck remote observing facilities has been the use of relatively standardized hardware and software at the various sites; the Keck remote observing facilities in Waimea served as the basic model which was then replicated at the various external sites. That model included the use of dedicated H.323-based video-conferencing systems manufactured by Polycom and multi-headed Sun SPARC-architecture workstations running Solaris and employing frame buffers and *X* servers capable of providing simultaneous support for both 8-bit and 24-bit *X* visuals, with separate *X* servers assigned to each monitor.⁹

Unfortunately, with the demise of the Sun SPARC workstation line and the purchase of Sun Microsystems by Oracle Inc., this component of the model is no longer available. Accordingly, the Keck remote observing software that runs at the external sites has been ported to run under generic Intel-based systems running Linux and utilizing frame buffers and *X* servers that lack the capability of simultaneously supporting 8-bit and 24-bit *X* visuals; this necessary capability was instead provided by new features present in Version 4 of *RealVNC*^{||}. As new Keck remote observing sites come online, they will employ this Linux-based version of the software; the facilities at Swinburne University and Yale University both employ this version.

The first few remote observing sites used monitors with 1280×1024 (horizontal \times vertical) pixel resolution, so the shared VNC desktops used for remote observing were sized accordingly. As newer sites came online, they typically had higher resolution monitors, either 1600×1200 or 1920×1200 pixels (see column 4 of Table 2), and many of the original sites have upgraded their monitors to these higher resolutions. Since 1280×1024 is currently the least common denominator between the various sites, that is still that default size of VNC desktops currently used for Keck remote observing. While this ensures that each VNC desktop can be fully displayed (without scroll bars or re-sampling, thus ensuring that all connected sites see the same VNC desktop rendered in the same way) at any of the Keck remote observing sites, it also means that screen real estate is not used efficiently at many of them. However, until WMKO upgrades the resolution of the monitors in its two remote operations rooms in Waimea and in its two control rooms on the Mauna Kea summit, this small-sized VNC desktop will likely remain the default for Keck remote observing.

9.2 Observing teams distributed across multiple sites

A growing number of Keck observations are now conducted by observing teams whose members are distributed between two or more external sites. While the existing VNC-based software easily supports this capability, some care needs to be taken in establishing the H.323-based video-conferencing connections between the various external sites in order to minimize audio and video latencies. While the video-conferencing systems at some of the external sites provide support for only point-to-point connections, many of the sites have systems that include an embedded multi-point controller capable of connecting four sites together in a multi-site video-conference. In general, connections should be made so that only a single trans-Pacific video-conference connection is established.

For example, in the case of a Keck observing team that is distributed between UCSD, UCLA, and CIT and where embedded multi-point controllers are available at UCLA and at WMKO, an optimal video-conferencing connection strategy would be for the UCSD and CIT sites to connect to the multi-point controller at UCLA, which would then in turn establish a connection to the multi-point controller at the WMKO headquarters in Waimea; the Keck observing assistant at the Mauna Kea summit would also establish a connection to the multi-point controller in Waimea. In this manner, any unnecessary routing of video-conferencing traffic across the Pacific is avoided and audio and video latencies are minimized.

The *noc2* software suite described in section 8.4 above alerts users of a VNC desktop whenever another site either connects to or disconnects from that desktop, noting the connection and disconnection times. This enables observers in such teams to quickly determine which of their colleagues is currently connected to a given VNC desktop, as well as to determine when a colleague not currently connected was last connected.

^{||}<http://www.realvnc.com>

10. USAGE METRICS AND STATISTICS

In reviewing the use of the external remote observing facilities over the last 5.5 years (see Fig. 9**), there has been steady growth in use of the mainland-only mode of operation. Before 2007, such usage accounted for less than 10% of mainland remote observing activity. By the end of 2009, such usage had grown to about 50% overall, as is now the dominant remote observing mode at both UCB and UCSD (see Fig. 10).

The systems that we have put in place at the mainland remote observing sites to enhance reliability and availability have so far proven quite effective. Over the last 9.5 years of operation, zero nights of telescope time were lost due to the use of mainland remote observing. However, some minutes of time have been lost during failover to ISDN circuits and observing efficiency is somewhat reduced during those periods when traffic is re-routed onto those circuits.




Figure 9. Usage of the WMKO mainland observing system for each (6-month) semester between February 2006 and January 2011. Bottom part of histogram (dark blue) represents “eavesdropping” mode usage and top part (light blue) represents “mainland-only” mode operation. Since early 2007, mainland-only mode has become increasingly popular. Semester to semester fluctuations reflect a variety of factors, such as telescope time assigned to observers affiliated with institutions lacking their own remote observing facility or nights allocated to interferometry, segment phasing, or telescope engineering.

11. APPLICABILITY TO OTHER OBSERVATORIES

Observatories like Keck that allocate most of their observing time to classically-scheduled observations typically need to provide some level of observer support facilities, such as control rooms, daytime sleeping facilities, and either dining facilities or food preparation areas. When planning such facilities, several important factors need to be considered, as illustrated in Table 1. With the advent of remote observing conducted from external sites, such planning becomes more challenging; as the fraction of external remote observations increases, the demand for support facilities at the observatory decreases while demand at external sites increases. In planning for the next generation of such observatories, it is useful to examine the impact of providing external remote observing for Keck Observatory.

**The number of “nights” displayed along the Y-axis is computed as the sum of *telescope-night-fractions* summed across both Keck Telescopes and across the number of nights in the given semester. The *telescope-night-fraction* value for a given night on a given telescope represents the fraction of that night (e.g., 0, 0.5, or 1.0) on which the observing program for that telescope included participation from observers working from at least one of the twelve external remote observing sites.




Figure 10. Usage of the mainland observing system by site since February 2006.




Figure 11. Usage of the mainland observing system by instrument since February 2006.

The solutions we have developed to support remote observing with the Keck telescopes are in principle readily adaptable to other observatories that use similar scheduling models and observing modes. In fact, many of those solutions were adopted for the remote observing program at Lick Observatory.¹¹ Given the acceptance that this mode of observing has achieved at both Keck and Lick and assuming that fossil fuel prices continue to increase the cost of trans-oceanic travel, it is almost certain for any new telescopes for which a significant portion of the observing time is classically-scheduled, observers will want to operate those telescopes from remote sites starting at first light (and prior to that, engineers will debug instruments remotely during the commissioning phase). If the new generation of such observatories is to make best use of their remote observing capability, the design work required to support this mode should be considered early in the project.

12. CONCLUSION

The effort to establish Keck “mainland-observing” facilities was an experiment that began over a decade ago. At that time, it was unclear whether such a mode of observing would prove reliable or whether it would be accepted by the Keck observing community. With a modest investment in hardware and software over the intervening years, that experiment has now been expanded successfully to include twelve external sites, and has spawned the corresponding implementation of remote observing at Lick; the resulting cumulative savings in travel costs and travel time have been substantial. The usage metrics presented here indicate that this mode of observing has earned significant acceptance within the Keck observing community. A number of factors contributed to the success of this experiment, but the most important were the decisions to avoid unnecessary complexity, and to construct the system using well-engineered and mature hardware and software.

It is also clear that remote observing from the observer’s home institution is not always the best solution for some observers and circumstances, and that observing “at the observatory” is an important option to maintain. In planning next-generation observatories such as the TMT, it is important to provide observers with the flexibility of choosing either observing mode, and to allocate sufficient resources during design and implementation to ensure that an optimum mix of observer support facilities is provided at the observatory as well as at the external sites.

ACKNOWLEDGMENTS

The network that makes all this possible has been developed over the past decade through partnerships involving universities and networking organizations in Hawaii, Australia, Washington and California with support from the National Science Foundation International Research Network Connections (IRNC) program and others. Keck Observatory is connected via the Mauna Kea Observatories Communication Network at the summit and other University of Hawaii connections on the Big Island of Hawaii. High speed connectivity from Hawaii takes place over submarine fiber provided by the Australian Academic and Research Network (AARNet) through the SX TransPORT project on the Southern Cross Cable Network, which also provides global connectivity to Australian universities and research facilities. The Corporation for Educational Networking Initiatives in California (CENIC) provides connectivity from the submarine fiber landing station in the U.S. to

CENIC's California Research and Education Network (CalREN), which serves California's universities, colleges and K12 schools. Connectivity is further extended to the global research and education networking community through the Pacific Wave distributed international peering exchange in Los Angeles, Sunnyvale, and Seattle.

We also thank the following WMKO staff who contribute to the ongoing operation of the remote observing facilities in Waimea and in California: Randy Campbell, Jon Chock, Liz Chock, Al Conrad, Scott Dahm, Grant Hill, Marc Kassis, Jim Lyke, Jeff Mader, Julia Simmons, and Hien Tran. We also acknowledge the efforts of the following individuals for their efforts in helping to set up and maintain the various external remote observing sites: Judy Cohen, Patrick Shopbell, Richard Walters, and Andy Boden (CIT); Deb Agarwal, Cedric Hui, Saul Perlmutter, Masaaki Yamato, and Tony Spadafora (LBNL); George Hau, Glenn Kacprzak, Jeff Cooke, and Michael Murphy (Swinburne University); Andrew Howard, Bill Boyd, Kelley McDonald, and Geoff Marcy (UCB); Chris Fassnacht and Daniel Wang (UCD); Gary Channan, Aaron Barth, Tammy Smecker-Hane, and Carrollan Simmons (UCI); Matthew Barczys, Chris Johnson, Leon Kopelevich, Ian McClean, Laurie Liles, and Jason Weiss (UCLA); Bill Strossman, Gillian Wilson, and Gabriela Canalizo (UCR); Peter James, David Kirkman, John O'Meara, and David Tytler (UCSD); Glenn Schiferl, Tommaso Treu, and Crystal Martin (UCSB); Marc Boolootian, Kelsey Damas, and Jim Warner (UCSC); Victoria Gardner, Robert Chang, Craig Henry, Dean Barufi, Marla Geha, Pieter van Dokkum, and Debra Fischer (Yale).

The W. M. Keck Observatory is operated as a scientific partnership among the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation.

REFERENCES

- [1] Jorgensen, I., Rodgers, B., and Crabtree, D. R., "Gemini observatory: Five years of multi-instrument queue operations," in [*Observatory Operations: Strategies, Processes, and Systems III*], Silva, D. R., Peck, A. B., and Soifer, T., eds., *Proc. SPIE* **7737**, 77370J–1 (2010).
- [2] Crabtree, D. R., "Scientific productivity of large telescopes," in [*Observatory Operations: Strategies, Processes, and Systems II*], Brissenden, R. J. and Silva, D. R., eds., *Proc. SPIE* **7016**, 70161A (2008).
- [3] Kibrick, R., Conrad, A., and Perala, A., "Through the Far Looking Glass: Collaborative Remote observing With the W. M. Keck Observatory," *ACM Interactions* **5**(3), 32–39 (May-June 1998).
- [4] Conrad, A. R., Kibrick, R., and Gathright, J., "Remote Observing with the Keck Telescopes," in [*Telescope Control Systems II*], Lewis, H., ed., *Proc. SPIE* **3112**, 99–110 (1997).
- [5] Cohen, J., Shopbell, P., and Bergman, L., "Remote Observing with the Keck Telescope Using the ACTS Satellite," in [*Final Report for NASA Grant BK-509-20-42-00-00*], 1–22 (March 1, 1998).
- [6] Kibrick, R. I., Allen, S. L., and Conrad, A. R., "Remote observing with the Keck Telescopes from the U.S mainland," in [*Advanced Global Communications Technologies for Astronomy*], Kibrick, R. and Wallander, A., eds., *Proc. SPIE* **4011**, 104–116 (2000).
- [7] 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**, 70160Y (2008).
- [8] 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).
- [9] Kibrick, R., Allen, S. L., Conrad, A., Stickel, T., and Wirth, G., "Optimizing the use of X and VNC protocols for support of remote observing," in [*Advanced Software, Control, and Communication Systems for Astronomy*], Lewis, H. and Raffi, G., eds., *Proc. SPIE* **5496**, 323–334 (2004).
- [10] Wirth, G. D., Kibrick, R. I., Lyke, J. E., Mader, J. A., and Goodrich, R. W., "More bang for the buck: lessons learned from remote observing at the W. M. Keck Observatory," in [*Telescopes from Afar*], in *this proceedings* (2011).
- [11] 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, T., eds., *Proc. SPIE* **7737**, 773712 (2010).