Remote observing with the Keck Telescopes

Al Conrad[a], John Gathright[a] and Robert Kibrick[b]

[a]California Association for Research in Astronomy
W. M. Keck Observatory
Kamuela, Hawaii 96743 USA

[b]UCO/Lick Observatory
University of California
Santa Cruz, California 95064 USA

ABSTRACT

Remote observing has become a popular option for Keck Telescope observers. This is not surprising, given the costs and hazards of traveling to the Keck Telescope site (located atop the 4160-meter summit of Mauna Kea), coupled with the reduced mental clarity that many observers experience when working at that high altitude. Often, some observers and support staff are unable to travel to the summit due to various temporary or chronic medical conditions.

We describe the remote observing capabilities currently provided at the Keck Headquarters in Waimea (located approximately 32 kilometers from the Keck summit and at an elevation of 850 meters), as well as the subset of capabilities now available from the mainland via the Internet.

The bandwidth available between the telescope and the remote observing site determines which of several remote observing software and networking architectures is most cost-effective. We describe our operational experience with several different architectures, differentiating between those used at Keck headquarters and those used for remote observing from the mainland. Methods for optimizing bandwidth are explored, including the pipelining of image readout with data compression and transmission to the remote site. Tradeoffs between network bandwidth, security, and portability of software to remote observing sites are also explored.

Keywords: remote observing, remote operation of telescopes

1. INTRODUCTION

We separate our topic into two parts. In Section 2 we discuss remote observing from Keck Headquarter in the town of Waimea on the Big Island of Hawaii. In Section 3 we discuss remote observing from mainland sites in California.

In both cases, the primary advantage of remote observing versus working at the Keck summit is avoiding the ill effects of working at high altitude.[1] In addition, remote observing from mainland sites significantly reduces travel costs and travel time.

2. REMOTE OBSERVING FROM KECK HEADQUARTERS IN WAIMEA

2.1. Overview

The town of Waimea is located 32 kilometers northwest of the Mauna Kea Summit. The driving distance is 77 kilometers and the time to drive, including a twenty minute acclimatization stop at the Hale Pohaku mid-level facility (elevation 3,000 meters), is approximately two hours. The 30 minute climb from Hale Pohaku to the summit requires four wheel drive. Most observers who work from the summit arrive one day early and spend a non-observing night at Hale Pohaku to acclimatize before their first night of observing.

Further author information- (Send correspondence to A.C. or R.K.)
A.C.: Email: aconrad@keck.hawaii.edu; Telephone: 808-881-3812; Fax: 808-885-4464;
J.G.: Email: johng@keck.hawaii.edu; Telephone: 808-881-3087;
R.K.: Email: kibrick@ucolick.org; WWW: http://www.ucolick.org/~kibrick; Telephone: 408-459-2262; Fax: 408-426-3115

Observers who work from Waimea operate from one of the remote operations rooms located at Keck Headquarters. The work stations in these rooms communicate with data taking and telescope control computers at the summit via a dedicated link. The current link, a 1.5 Mbit/sec T1, will be upgraded to a 45 Mbit/sec DS3 in August, 1997.

Together with the advantages of avoiding the ill effects of altitude, we have found that greater accessibility to telescope management, instrument specialists, and the engineering staff, all of whom have their offices near the remote operations rooms, has helped visiting observers to better utilize the facility. Direct contact between observers and technical staff stands out as one of the primary advantages of remote observing from Waimea. In contrast, the lack of this type of direct support and contact stands out as one of the major disadvantages of remote observing from the mainland.

The personnel involved in an observation includes the principal investigator (PI) and the observing assistant (OA). The PI conducts the observation and the OA operates the telescope and assists the PI. In addition, colleagues, instrument specialists, engineers, and technicians may assist with the observation either from the control room, from their homes, or from a home institution on the mainland.

For the purposes of this discussion, we restrict ourselves to remote operations issues related to the role and location of the PI and OA only. Excluding the mainland case discussed in the next section, the four possible PI/OA configurations are given in Table 1.

Table 1. Personnel configurations for observing from Waimea or the summit. Number of observing runs for each configuration for first 4 months of 1997.

configuration	summit	HQ	Keck-1	Keck-2
1	PI, OA	–	2	35
2	PI	OA	1	0
3	OA	PI	22	24
4	–	PI,OA	23	0

For simplicity, and because it is the most interesting case from the point of view of remote observing, we further restrict our discussion to configuration number 4: both PI and OA observing remotely from Waimea. Note that in both configurations 2 and 4 there are still CARA staff members on the summit. Those staff are responsible for starting up and shutting down the telescope.

2.2. Data Protocol

Observing software divides roughly into the five categories given in Table 2.

Table 2. Link Protocol for Major Software Categories

software system	Link Protocol
instrument image display and quick look	X
instrument status and control screens	X
telescope status and control screens	X
acquisition/guider camera images	X
data reduction	ftp & NFS

As can be seen from this table, for the first three categories the work stations in Waimea serve only as glorified X-terminals. To direct the display of a summit application to a particular screen in Waimea, we use the desktop menu provided with the window manager running on the remote host. For example, this line from a Waimea user menu

```
rsh summithost summitapp -display `hostname`{$DISPLAY}
```

will evaluate to

```
rsh summithost summitapp -display hqhost:0.2
```

when selected from the third monitor on the multi-headed work station called "hqhost" in Waimea.

The primary advantage of the X protocol method is simplicity. With this approach, all software runs on summit machines just as it would for summit observing. No configuration changes or redistribution of software is needed to switch between summit versus remote observing.

Since we use a dedicated display host on the summit as well, there is no difference between the summit and remote configurations other than the speed of the link (see Fig. 1).

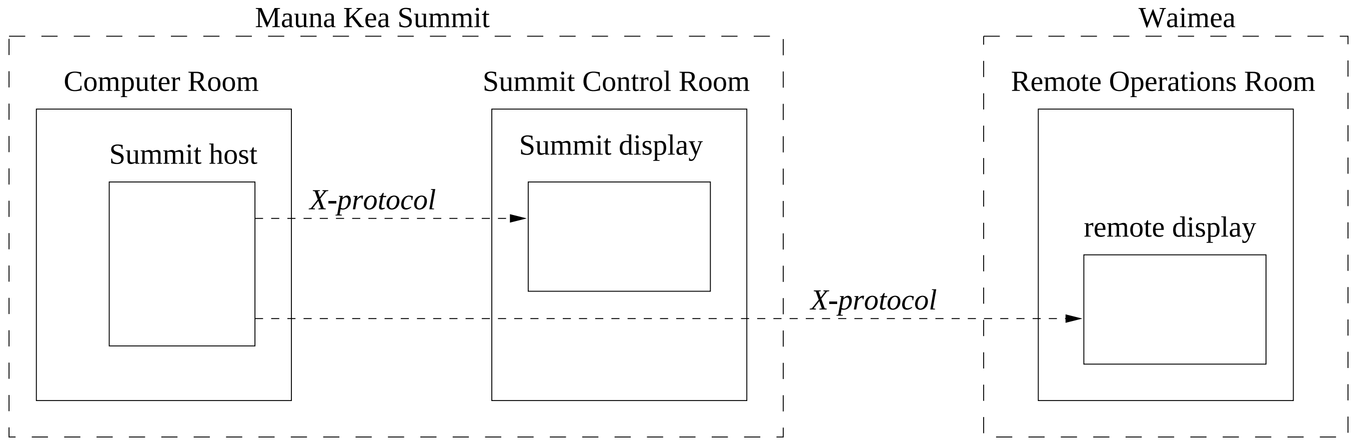

Figure 1. Summit/Waimea Configurations

The disadvantage of the X protocol method is performance. The throttling effect of the 1.5 Mbit T1 link on simple, text-only displays is negligible. For displays that require large bitmap transfers, however, the effect is noticeable. Most notably these are image readout displays, but also included in this category are applications that convert their line graphics to X-bitmaps (e.g, DataViews). When remote observing on both Keck-1 and Keck-2, currently a zoom/pan operation on the bitmap CCD image display may take as long as five seconds. Although this is marginally acceptable for viewing the result, the subsequent buffering of mouse clicks and mouse motions puts the observer out of synch with the application. This effect can result in a time loss which is significantly higher than the five second delay. In this case a distributed approach like that described in section 3.3 below would improve the performance of remote operation over the T1.

Our experiments with remote observing from Waimea with a 45 Mbit link show little performance degradation with bitmap displays run remotely. This is not surprising, given that the 45 Mbit rate exceeds the 10 Mbit rate achieved on the backbones in place at Waimea and the summit. Given the anticipated upgrade to a 45 Mbit link, we plan to continue with the X protocol approach when remote observing from Waimea.

Note that bandwidth required for a zoom/pan operation will scale with the pixel resolution of the remote display, not with the pixel resolution of the CCD. In particular, assuming that the resolution of our remote displays remains at approximately 1K-squared pixels, the 45 Mbit bandwidth that is sufficient for the 2K-squared detectors we use on the our current generation of instruments, will also be sufficient for the 8K-squared detectors due on Keck in near future.

2.3. Video Conferencing

Personnel in the summit and remote control rooms communicate with one another via a dedicated video conferencing system (PictureTel Venue 2000) running over a portion of a dedicated T1. In the cases where the PI and OA are separated (configurations 2 and 3 in Table 1), this system has proven to be critical for successful observing. Although the system is seldom used to actually view a diagram or hand drawn picture, we have discovered that the visual cues that come from body language, knowing that someone else is listening in, etc. are essential for two people collaborating on the operation of telescope and instrument.[2] This effect is born out in the literature.[3]

Before acquiring the dedicated PictureTel system, we investigated generic systems that utilize TCP/IP and work station technology.[4] Although cost and flexibility favor the latter, we have found that the superior performance and features available from dedicated systems favor the former. Surprisingly, audio quality stands out as the most important factor favoring the dedicated system. Other features include better picture quality, preset camera positions, and a larger field of view. In addition, the dedicated video system isolates the link used for video conferencing from the link used for summit to Waimea computer network traffic, thus conserving network bandwidth and preventing the loss of audio and video contact should the computer network fail.

2.4. Issues and Hurdles

2.4.1. Analyzing Usage

One of the greatest challenges we have faced with remote observing is determining which cooperating processes are consuming disproportionate amounts of the available bandwidth; or, more simply stated, "Who is hogging the net?" Although a 45 Mbit link accommodates normal traffic, we have found that a large block transfer which is network limited (typically a tape back up of an NFS mounted file system) can effectively lock out communication between two event-limited processes. To date we have not found a method of throttling these block transfers and rely on operational rules to prevent this type of transfer from interfering with operations.

Prior to the advent of switching hubs, various tools that run on an ancillary work station set up in promiscuous mode (e.g., etherman and etherfind[5]) could be used to quickly isolate an offending application. Although switching hubs have given us greater performance on the backbone, they, unfortunately, prevent us from using these tools for analyzing traffic on the dedicated link. Recently we have found some relief using vendor-supplied methods which run directly on the switch. Ultimately we will need to procure one of the special purpose packages available for analyzing a switching network (e.g., 3com Transcend).

2.4.2. Separating Operational Machines from the Office Environment

In general, computer LANs must trade off the reliability of autonomous work stations versus the efficiency and uniformity of centralized servers. Although most office environments favor centralized servers, work stations needed for observing favor the reliability that comes with autonomy. One of the challenges we have faced with remote observing is making the transition from what has historically been an exclusive office computing environment to one that is a combination of office systems and operational systems.

2.4.3. Cross Mounts

Returning to table 2, we see that images to be reduced during the night are accessed via ftp or NFS cross mounts. For accessing data partitions, NFS cross mounts present a simple uniformity, or "black box," to the observer. Unfortunately, we have found that observers need to understand the internals of the black box to avoid common pitfalls. The most common error occurs when running an operation that reads and writes a large image (8 to 16 megabytes) in such a way as to impose an inadvertent round trip on the summit link. A typical example is running IRAF "rfits" on an HQ computer with a source and destination disk on the summit. We have also experienced difficulty with NFS when a mount failure results in a long timeout, not only for the application accessing the remote disk, but for all applications running on that workstation. We have considered disabling NFS cross mounts and forcing ftp transfers, but thus far we are continuing with attempts to educate observers on how to best avoid these pitfalls.

2.5. Use of WWW

Although we have started using WWW methods such as HTML, CGI, and java to facilitate observing preparation (i.e., running simulators, specifying instrument configurations, and etc.), to date we have not experimented with these methods for real time observing. Although it is likely that this technology will someday become the appropriate method for remote observing, current difficulties writing to local disks and providing data security make it inappropriate for this purpose today.

2.6. Summary

Remote observing from Waimea is now routine and the majority of Keck observers take advantage of it. Although the inefficiencies of the X-protocol result in some delays with our current T1 connection, the simplicity of this approach is overwhelming and, at DS3 bandwidth, suffers no loss in performance when compared to summit observing. The dedicated video conferencing system, together with its high quality audio component, is essential for remote observing. The hurdles we have faced include analyzing network usage, separating operational machines from the office environment, and managing cross mounts.

3. REMOTE OBSERVING FROM THE MAINLAND

While there is a clear mandate and policy at Keck Observatory to provide remote observing from Waimea, there is currently no similar mandate or policy to directly support remote observing from the mainland. Accordingly, with the exception of occasional support of various mainland remote observing demonstration projects, there is no significant ongoing allocation of technical resources to support mainland remote observing on a routine basis. As a result, mainland remote observing is still at the proof-of-concept level, as contrasted with the routine operation level that has been achieved for remote observing from Waimea.

However, various segments of the Keck observing community have expressed significant interest in and have requested specific instances of support for a variety of mainland remote observing capabilities. In an effort to provide immediate, short-term solutions in response to these requests, software staff at Keck Observatory and UCO/Lick Observatory have informally collaborated to develop several software applications to provide a subset of remote observing capabilities from the mainland via the Internet. These applications (see Section 3.5.1) are continuing to evolve in response to changing network configurations and security concerns, but progress on their development is currently constrained by the lack of any formal commitment of resources towards their development.

3.1. Motivation

As described in the previous section, remote observing from the Keck Headquarters in Waimea affords many advantages over observing at the Keck summit. However, for observers whose home institutions are located on the mainland, remote observing at Waimea still involves many of the same costs with respect to airfare, surface transportation, lodging, meals, and travel time. These costs often exceed \$1,000 per observer for a typical observing run of 2 nights. Thus, the total travel-related costs for a mainland observing team (3 observers on average) remotely observing from Waimea can be quite high and are roughly the same as (and in some cases higher than) the costs of having that same team observe from the Keck summit.

In addition, a significant number of observing runs currently involve only a single night and as competition for observing time continues to increase, some observing runs will be assigned only half a night. In both cases, it becomes more difficult to justify not only the direct travel and lodging costs but especially the travel time for such short observing runs. Finally, many of the same temporary medical conditions which prevent observers from traveling to the Keck summit (e.g., ear or sinus infections) also prevent them from flying on a airplane to Hawaii; several of the specific mainland remote observing sessions we have supported during the past several years have been motivated by such medical concerns. For all of these reasons, remote observing from the mainland is often very desirable.

3.2. Assumptions

While the network links from the Keck summit to Waimea can be relied upon to routinely provide adequate bandwidth for remote observing, the same cannot be said for the existing links from the Keck summit to the mainland. The current Internet route to the mainland provides adequate bandwidth for many remote observing functions during nighttime hours most of the time, but its bandwidth, reliability, and availability cannot be guaranteed at any level. While complete outages are infrequent, brief interruptions or reductions in bandwidth can occur unpredictably. While conventional public dial-up telephone lines may provide somewhat more predictable performance, even these are subject to sporadic interruptions and in general do not provide adequate bandwidth even when working reliably.

Accordingly, until such time as the links to the mainland provide at least the same level of bandwidth, reliability, and availability as provided by the link between the Keck summit and Waimea, we assume that mainland remote observing involves an observing team and that at least one of the members of this team is located in Hawaii, either at

Waimea or at the Keck summit. This on-site observer assumes responsibility for the observing program in the event of interruptions of the network link to the mainland. (With the exception of engineering tests involving a NASA satellite link,[10] all of the mainland remote observing conducted to date has included at least one member of the observing team located on-site in Hawaii.) Coordination between the on-site observer and the observing team members on the mainland is usually accomplished by keeping open a telephone call (often connected to speaker-phones at each end) or by using digital audio tools to carry conversations over the Internet.

Under this definition of mainland remote observing, the need for real-time access to acquisition/guider camera images from the mainland becomes less significant, since the on-site observer can provide local confirmation of target acquisition and can monitor guider performance. Similarly, the on-site observer assumes responsibility for instrument set up using the GUIs provided on the on-site machines. Accordingly, providing mainland observers with timely access to science images for collaborative quick-look quantitative analysis becomes the primary function of this type of mainland remote observing. We will refer to this mode of operation as "eavesdropping" to distinguish it from the "full-up" remote observing currently provided at Waimea.*

3.3. Mainland remote observing software models

Remote observing from the mainland requires a reliable data link that provides adequate bandwidth between the remote site and the Keck summit at a cost that is not prohibitive. The bandwidth that is required depends on the software model that is used. Several different models are possible and have been used to date:

1. All observing software runs at Keck summit with X displays routed to remote site
2. All observing software runs at Keck summit but image data is relayed to remote site
3. Components of the observing software are run both at the Keck summit and the remote site

3.3.1. X Displays routed to remote site

The first model parallels the current mode of remote observing from Waimea, and thus shares all of its advantages and disadvantages, as described in the previous section of this paper. Since in most cases the link bandwidth between mainland sites and the Keck summit is significantly lower (and often less predictable) than that between the summit and Waimea, those disadvantages that are bandwidth-sensitive become more significant. Given such reduced bandwidth, the primary disadvantage is that many interactive operations, especially those involving interactions with bit-mapped graphics image displays (e.g., pan and zoom, color map adjustments, and other cursor-based interactions) become painfully sluggish, with the result that the remote observer's mouse motions and button presses rapidly get out of synch with the software running at the summit. Unless the remote mainland site has access to a relatively high-speed link (i.e., one approaching 1 Megabit/sec) that provides reasonably consistent bandwidth, this model of remote observing from the mainland is extremely difficult to use, and results in significantly reduced operating efficiency. Further, for those remote sites connecting to the Keck summit via the Internet, other limitations currently preclude the use of this model (see Section 3.5.3).

3.3.2. Relay of image data files to remote site

The second model overcomes many of the limitations of the first, since it allows image display software (e.g., SAOimage, ximtool, SAOtng, etc.) to be run locally at the remote site and thus provide excellent interactive performance for typical operations such as panning and zooming. However, this model requires the installation of a very small amount (less than 50 Kbytes) of software at the remote site (see section 3.5.1) and also imposes a link bandwidth requirement that scales with the typical size of images to be transferred. For the purposes of this discussion, we will focus on the two existing optical instruments at Keck, the HIRES Spectrometer[6] on Keck-1 and the LRIS Spectrometer[7] on Keck-2, since these currently involve the largest detectors and hence the largest images. Both of these instruments employ 2K by 2K CCDs that are digitized to 16-bits per pixel, thus yielding full-frame images (including prescan and overscan pixels) that occupy approximately 9 MBytes. Depending upon the inherent contrast and read noise of the images to be transmitted, lossy compression techniques[8] can be used to reduce the size of these images

*Even if network bandwidth improves sufficiently to permit full-up mainland remote observing, "eavesdropping" will still be useful as part of service observing at Keck, which may become increasingly popular in the years to come.

to between 1 to 2 MBytes while still retaining sufficient information to allow the compressed image to be adequate for quick look quantitative data analysis.[†]

A fundamental requirement of remote observing from the mainland is that this quick look quantitative analysis of science images should be available with minimal delay following completion of each exposure. Our goal has been to reduce this delay to 3 minutes or less. Using the existing HIRES Sparc-10 data taking computer, compression of the image prior to transmission requires 15 to 30 seconds elapsed time, and decompression at the receiving end requires approximately the same time, assuming a comparable machine at that end. Assuming the worst case values, this requires that a 2 MByte compressed image be transferred across the link in 120 seconds in order to meet our 3 minute goal. This imposes a minimum average link bandwidth requirement of 17 Kbytes/sec, or roughly 133 Kbits/sec.

3.3.3. Running observing software components at both sites

The third model provides additional advantages over the first two. In addition to allowing the image display software to be run locally at the remote site, it also allows the transmission of image data to be overlapped with readout of the CCD. The HIRES and LRIS CCDs both currently readout at a rate of 30 Kpixels/sec in single-amplifier readout mode and 60 Kpixels/sec in dual-amplifier mode; both systems employ at pipelined image display server so that image pixels are displayed in real-time as the image is being read out of the CCD.[9] These readout rates result in full-frame readout times of slightly less than 3 minutes and 1.5 minutes, respectively. By overlapping image transmission to the remote site with CCD readout, overall utilization of the link bandwidth is significantly increased. In addition, the remote observer begins to receive image data as soon as the CCD readout commences. If the link bandwidth to the remote site is comparable to the CCD readout bandwidth (i.e, 60 Kbytes/sec for single-amplifier mode or 120 Kbytes/sec for dual-amplifier mode), then an uncompressed image can be transferred without penalty and the image can be displayed in real-time at the remote site just the same as would be displayed if observing at Waimea or the Keck summit. If the link bandwidth is within a factor of two of the CCD readout bandwidth, an uncompressed image can still be transferred within the goal of completing the transfer within 3 minutes after readout completion. Otherwise, for lower values of link bandwidth, lossy image compression will still be required as part of the image transmission pipeline in order to complete the image transfer within the specified goal. However, such pipelined image compression is likely to be significantly less effective than the compressing a fully read out image (as is done in the second model).

The major disadvantage of the third model is that it requires distribution (and possible porting) of a significant body of software to the remote observing site. Once this software is installed, there are ongoing costs of maintaining this software and keeping it in synch with upgrades to corresponding software components at Keck.

3.4. Existing alternatives for data links to mainland sites

Currently, there are only a few available data link options that might provide adequate bandwidth between the remote site and the Keck summit at a cost that is not prohibitive:

1. Slip or PPP connection via conventional analog public dial-up telephone line
2. Internet connection
3. Experimental TCP/ATM connection via NASA ACTS Satellite[10‡]

The first option is available from nearly all mainland sites, but even with the advent of 56Kbaud modems, the bandwidth often provided is still about a factor of two less than what is needed to meet the specified goal. Since most of the data delivered to the modem is already compressed, any compression capability within the modem is unlikely to yield any significant improvement over the modem's inherent bandwidth. Further, while high-speed modems have been successfully used in other remote observing applications within the mainland,[11] our current experience with 33.6Kbaud and 28.8Kbaud modems indicate that it is often difficult to obtain a mainland-to-Hawaii dial-up connection of sufficient quality to allow modems to establish connections at these baud rates. Quite often, the line

[†] while higher compression ratios, typically of order 100, are often acceptable for visual qualitative analysis of direct images, quantitative analysis of weak spectral features limits the acceptable compression ratios to factors of 5 to 10.

[‡] http://kronos.lerc.nasa.gov/acts/acts.html

quality is such that the modems fall back to lower baud rates, such as 14.4Kbaud. In addition, we have often experienced line drop outs of sufficient duration to result in loss of the modem connection, thus requiring a re-dial to re-establish the connection.[§]

The second option could be made available to most existing Internet sites, subject to various security constraints (see Section 3.5). Our experience to date indicates that the link from the Keck summit to the UCO/Lick facilities in Santa Cruz, California usually provides adequate average bandwidth (typically ranging between 30 Kbytes/sec to 120 Kbytes/sec) during night-time (HST) observing hours to support either remote observing models two or three (see Section 3.1). On some nights, this link can intermittently provide bandwidth comparable to that obtained from Waimea, thus making the first observing model marginally possible, except for other constraints (see Section 3.5.3)

The third option is experimental, and is currently available only from the California Institute of Technology in Pasadena, California. This experiment is of limited duration and the high-speed satellite link that it employs is only available for parts of selected nights as determined by NASA.[10] Further, this experiment has demonstrated that while a high-speed satellite channel can provide the bandwidth required for any of the three mainland remote observing models, the long round-trip transmission times that result from the very long transmission paths associated with geostationary satellites prove to be a serious obstacle for the first observing model. In particular, the X protocols associated with the transmission of bit-mapped graphics (as used both by image display servers and various graphical user interfaces) appear to involve sending a large number of small packets, resulting in sluggish interactive performance for such applications, despite the inherently high bandwidth of the satellite link. Similarly, normal TCP buffering mechanisms do not work efficiently given these long round-trip delays, and SunOS network software and kernel patches are required in order to obtain adequate TCP performance. Packets sent between California and Hawaii via this ACTS satellite are subject to an unusually long round-trip transmission time (typically 0.55 seconds) because this satellite is located over the Eastern United States at 100 degrees West longitude. However, despite these difficulties, the experiments at Caltech have clearly established that even a non-optimally located geostationary communications satellite can provide a reliable high-speed link that can be effectively used for remote observing provided that transmission protocols and observing modes are carefully chosen and tailored to compensate for the long round-trip transmission times.

3.5. Current Operational Experience with Mainland Remote Observing

Our primary experience to date involves mainland remote observing via the Internet, primarily from Santa Cruz and other sites within California. Shortly after the commissioning of the HIRES Spectrometer in late 1993, we conducted initial experiments using the first remote observing model (i.e., re-directed X displays). These were eventually abandoned due to unacceptable interactive performance and occasional unexplained stalls during which displays would mysteriously fail to update (see section 3.5.3 below). Subsequent experiments in 1994 and 1995 involved running some instrument control GUIs (e.g., xhires[12]) and applications (e.g., guider eavesdrop) on machines in Santa Cruz and sending instrument command and status messages across the Internet to the Keck summit. These were ultimately abandoned due to security restrictions imposed either in Santa Cruz or at Keck.

In 1995, a filtering router was installed at the Keck summit to provide an increased level of security for the computers responsible for instrument and telescope control; this router has been programmed to reject incoming connection requests from nearly all outside sites on the Internet, while allowing summit machines to initiate outgoing connection requests to outside machines. The installation of this router complicated our plans for remote observing from the mainland, since it was no longer possible to initiate connections to the Keck summit from arbitrary remote sites. While it is possible to manually reprogram the filtering router to explicitly permit incoming connection requests from specific remote sites, operationally this is a difficult process and one that is potentially disruptive to network connectivity at the Keck summit. Since many mainland remote observing requests occur on short notice (especially requests resulting from temporary medical conditions), reprogramming of the summit router to allow temporary access from specific mainland remote observing sites is neither a practical nor a secure alternative. [¶]

[§] It is possible that these problems are confined to the particular long distance carrier we are using, and that connections originating from other mainland sites and/or involving different carriers would not be subject to these problems.

[¶] Keck is in the process of evaluating SecureID methods as an alternative to the filtering router. For remote observing model number one (X-protocol) a SecureID system can control access by user, whereas a filtering router can only control access by machine. Moreover, unlike the filtering router, it can be reconfigured on short notice. On the other hand, because they employ cooperating applications on opposite sides of the firewall, it is not clear that a SecureID system can be used for either of the other remote observing models. We are still investigating the applicability of SecureID for this type of distributed system.

About the same time, there was increasing interest in supporting mainland remote observing to sites other than Santa Cruz. Accordingly, we needed to develop a remote observing mechanism that was easily distributed to outside sites but which complied with Keck security constraints. As a result, the "shadowmon" suite was developed at Lick, and successfully used to support the automated relaying of compressed Keck images (science and guider) to a variety of California remote observing sites (i.e., Santa Cruz, Berkeley, San Diego, and Livermore) as well as to the Space Telescope Science Institute (STSci) in Baltimore, Maryland and the Center for Astrophysics in Cambridge, Massachusetts.

3.5.1. The "Shadowmon" Suite

The Shadowmon suite consists of a very small set of simple utility programs, some of which run at the Keck summit and others which run at the mainland remote observing site. A receiver program, called "server", is packaged (as a compiled executable) into a compressed Unix tar format file which also contains copies of the fcompress/hcompress image compression software tools[‖] developed and distributed by the Space Telescope Science Institute. Both SunOS and Solaris versions of these tar files have been prepared, and they are approximately 30 Kbytes long. These compressed tar files are currently distributed to mainland remote observing sites via anonymous FTP, although in the future such distribution might be done via the WWW.

At the remote observing site, the observer creates a writable directory into which the remote observing images will be received. An appropriate version of the compressed tar file containing the server software is ftped into this directory and unpacked. The remote observer then invokes the server program from that directory, at which point this program will begin listening for connection requests that originate from Keck.[**] No further action is required by the remote observer. Once corresponding software is started at Keck, any subsequent images from the specified instrument will be automatically relayed to the remote observer's machine and will be written into this directory using the same file name as was used to record that image on the computer at Keck; the server program will write only into the current directory and will not write into any other directory on the disk. As receipt of each new image commences, a message is written to the screen of the remote observer's machine, and several beeps are issued when the complete image has been received and decompressed. Barring network disruptions or congestion, the relayed image will be available for use at the remote site in under 3 minutes from the time it was written to disk at Keck. Although the relayed images have been subjected to slighty-lossy compression, for purposes of quick-look quantitative analysis, they are virtually indistinguishable from the uncompressed images. If higher levels of compression can be tolerated, transfer times can be further reduced.

Observer's can also request the OA to save an acquisition/guide camera snapshot to a Keck disk and the shadowmon script will queue it for compression and transmission to the remote site; since these images are much smaller than the science CCD images, they are typically transferred to the remote site in about 15 seconds. This allows remote observers to verify that observing targets have been correctly acquired.

On the Keck side, the shadowmon script is invoked by CARA staff (and not by the remote observer) and usually is run on the data taking computer associated with the selected instrument. Parameters to this script identify the location (i.e., Internet IP address) of the remote observing computer on which the server program is (or will be) running and specify the location of a spooling area on a disk at Keck where copies of images will be compressed. An optional parameter allows adjustment of the level of data compression to be used. As each new image for the selected instrument is recorded on disk at Keck, the shadowmon script queues it for subsequent compression and transmission to the remote site. Since the network connection to that site is initiated from the Keck side of the link, there is no conflict with the security restrictions of the filtering router at the Keck summit.

The shadowmon suite assumes that the remote observing site does not impose its own security restrictions that would reject the incoming connection requests from the shadowmon software running at Keck. In such circumstances, remote observers would need to petition their local network authorities to provide a specific exemption for shadowmon-related connection requests. Such requests would originate from a specific IP address at Keck and would be directed to a specific port on a specific host at the remote site. Thus, although local authorities would be required to "open the door", they would not need to open it very far. Even if some intruder mounted an attack against the

‖available via http://www.stsci.edu/software/hcompress.html

**Any connection requests that originate from sites other than Keck are logged and rejected.

Table 3. Tradeoff between network bandwidth and time to compress/decompress images for a typical image (9 MByte uncompressed, 2 MByte compressed)

Link Bandwidth in KBytes/sec	Link Bandwidth in Kbits/sec	Time to Transfer Uncompressed Image	Time to Transfer Compressed Image	Time to Compress and Decompress	Total Time
10	80	900s	200s	60s	260s
20	160	450s	100s	60s	160s
40	320	225s	50s	60s	110s
80	640	113s	25s	60s	85s
100	800	90s	20s	60s	80s
120	960	75s	16s	60s	76s
160	1280	57s	13s	60s	73s
193	1544	47s	10s	60s	70s

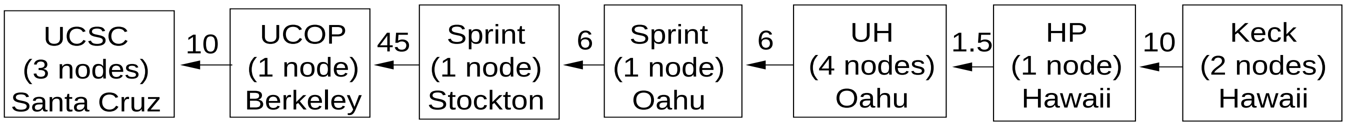

Figure 2. Current Internet route between Keck summit and Santa Cruz, California. Numbers over arrows specify speed of that segment in Mbits/sec.

remote observer's machine via this port, they should not be able to do more harm than to fill the remote observer's disk with bogus FITS files.

At the end of the observing session, a "shadowstop" script is run on the Keck side to shutdown the shadowmon script and its derivative processes. The server program is also terminated at the remote site. Since the shadowmon script is started and stopped by CARA personnel and since those personnel specify the address of the remote observing site, remote observers cannot continue to obtain Keck images beyond the end of their observing run, even if they leave the server program running on their machine.

In developing the shadowmon suite, we examined the relationship between the typical time required for compressing and decompressing an average size (9 MByte) image versus the time required to transmit compressed and non-compressed images for various link bandwidths. These results are summarized in Table 3. For link bandwidths lower than approximately 120 KBytes/sec, it is faster to transmit a compressed image, while for higher link bandwidths it is faster to bypass the compression and decompression steps and to transmit the image uncompressed. Given that the link bandwidth of the existing Internet route rarely exceeds 120 KBytes/sec, the default mode of shadowmon always compresses and decompresses the images that are sent.[††]

The results in Table 3 are based on typical compression/decompression times using a Sun Sparc-10 computer at each end. If faster machines are used so that the compression/decompression times can be decreased, then the bandwidth at which it becomes advantageous to transmit images in uncompressed-form will also decrease. Thus, under some circumstances, it may be advantageous for shadowmon to transmit images uncompressed.

3.5.2. Attributes of the Current Route Between Keck and the Mainland

Figure 2 shows the current topology of the Internet route between the Keck summit and UCO/Lick Observatory facilities in Santa Cruz. (Prior to late 1996, the Hawaii to mainland segment was 1.5 Mbit/sec T1 line between the University of Hawaii Manoa campus and the NASA Ames Research Center in Mountain View, California.) The current network configuration at Keck insures that mainland remote observing traffic does not impact the Waimea to Keck summit link.

[††]The server program is able to receive and automatically discriminate between different types of FITS files including those that are uncompressed as well as ones compressed using either the STSci fcompress utility or the Unix compress utility.

3.5.3. Unusual Internet Problems Complicate Remote Observing

In debugging the "shadowmon" software, we encountered a curious network problem that is hopefully atypical but certainly illustrative of the sorts of unexpected problems one can encounter when attempting to use the Internet for remote observing. When attempting to transfer specific image files from the Keck summit to Santa Cruz via standard ftp (or via direct TCP socket connections), file transfer would consistently stall at precisely the same spot in a given image. These stalls would usually occur during nighttime hours, while the identical image file could often be successfully transferred during daytime hours. Other image files and ASCII text files could be reliably transferred at all hours. In addition, some image files which successfully transferred uncompressed would fail to transfer when compressed using the "fcompress" utility (obtained from STSci). Thus, it appeared that some piece of network hardware somewhere along the route had difficulty transferring particular data patterns, and that this difficulty was most severe during nighttime hours in Hawaii.

In order to isolate this problem, we extracted several hundred bytes from one of the image files whose transfer had stalled, collecting the bytes from the region of the file where the stall occurred. We then generated a modified version of the Unix "ping" utility[‡‡] in which this extracted sequence of bytes was substituted in place of the standard ping test packet (that test packet is normally a linear ramp). The Unix "traceroute" utility was then used to identify all of the nodes in the route between Santa Cruz and the Keck summit. The modified ping utility was then used to individually ping each node of the route and measure the resulting packet loss from that node. While most nodes showed negligible losses, during nighttime hours the node at the mid-level facility at Hale Pohaku showed severe packet losses for these particular "poison pill" ping packets, while the immediately preceding node at UH/Manoa showed negligible losses for these same packets. At the same time, standard ping packets directed to the Hale Pohaku node showed negligible packet loss. Corresponding results were obtained when pinging the Hale Pohaku and UH/Manoa nodes from machines on the Keck summit. This implied that the data-dependent packet losses were occurring on the network segment that connects UH/Manoa and Hale Pohaku.

After contacting the relevant network support staff at the University of Hawaii (which is responsible for operating this segment), the problem was isolated to an obsolete microwave transmitter at the Hale Pohaku site. This transmitter was originally installed many years ago to provide voice telephone circuits, but, as a cost saving measure, was pressed into service as a digital communications channel despite the fact that it lacks modern hardware mechanisms for digital traffic (specifically, B8ZS signaling). During cold nights, the performance of this transmitter becomes marginal for certain specific data patterns. Since there is apparently no inexpensive hardware solution to this transmitter problem nor any immediate plan to correct it, our approach was to modify the shadowmon software so that the data is encoded prior to transmission to avoid these specific bit patterns and decoded following receipt at the remote mainland site. While the offending bit patterns can occur in any images or other binary data files, the "fcompress" image compression utility appears to increase the likelihood of such occurrences, while the Unix "compress" utility usually eliminates them. Unfortunately, the latter does not efficiently compress image data.

Unfortunately, these same offending bit patterns sometimes appear in X protocol packets, particularly those associated with the display of bit-mapped graphics regions. Thus, this transmitter problem can sometimes cause X displays that have been redirected to mainland machines to stall during display updates (see section 3.5). As a result, even though the current Internet route can often provide adequate bandwidth (at least to Santa Cruz) to support the remote display of bit-mapped X graphics originating from Keck summit machines, this transmitter problem frequently prevents such graphics from being reliably displayed. Accordingly, UH has requested that we avoid sending such packets across this route, since they tend to provoke problems with this transmitter. Fortunately, there are various plans underway to redirect the network traffic from Hale Pohaku so that this obsolete microwave transmitter can be retired, and it is hoped this process will be completed within the next year.

This network hardware problem at the Hale Pohaku site is instructive in that it illustrates the fragility of the existing Internet route between the mainland and the Keck summit. While we have some degree of control over the network hardware at the Keck and Lick ends of this route and have support staff who can be called to resolve problems when they develop, we have no similar control or support for the intervening nodes of this 12 to 14 node route. Should problems develop in those nodes during a mainland remote observing session, there is little we can

[‡‡]The ping utility uses the low-level ICMP protocol to send test data packets to a remote node which then simply echoes the test packet back to the sender. Ping counts the number of test packets successfully echoed and also measures the round-trip time for each echo. Ping derives it name for the analogous sonar operation typically performed by submarines and surface ships.

do to correct them. Further, late night hours are often the ones that network maintainers select for equipment maintenance and configuration changes, so as to minimize disruptions to the majority of their users who work during the day. While we usually receive advance notice of such scheduled maintenance on network hardware at the Lick and Keck nodes, we seldom receive warning of such activities on the intervening nodes. Thus, hardware problems or routine maintenance activities on the network hardware at intervening nodes represent an unpredictable element of risk when using the Internet for mainland remote observing, and mainland observers need to be acutely aware of that risk.

3.6. Summary

While "full-up" remote observing is not yet routinely available from the mainland, significant progress has been made toward that goal. The "eavesdropping" capability provided by the shadowmon software has been successfully and routinely used by a variety of mainland sites. It demonstrates both the significant potential as well as the many difficulties inherent in using the Internet for remote observing. The experiments conducted at Caltech using the NASA ACTS satellite have clearly established that such satellites could provide the required bandwidth, reliability, and security required for full-up mainland remote observing. Finally, the increasing availability of digital communication technologies such as ISDN services will likely open new possibilities, including mainland access to the Keck video conferencing system.

3.7. Acknowledgements

We wish to acknowledge the following individuals for their invaluable assistance in tracking down the network problems with the Hale Pohaku microwave transmitter: Jim Warner and Mark Boolootian, UCSC Communications and Technology Services; Jonathan Chock, Keck Observatory; and PuiHin Rhoads and Tjet Sun of the University of Hawaii.

REFERENCES

1. P. J. G. Forster, "Health and work at high altitude: a study at the Mauna Kea observatories," *Publications of the Astronomical Society of the Pacific* **96**, pp. 478–487, 1984.
2. "Mission to control: smile; NASA turns to videoconferencing to aid shuttle communications," *Computerworld* **29**(41), p. 28, Oct 9, 1995.
3. J. Storck and L. Sproull, "Through a glass darkly: What do people learn in videoconferences," *Human Communication Research* **22**, pp. 197–219, Dec 1995.
4. S. Casner and S. Deering, "First IETF Internet Audiocast," *Computer Communications Review* **22**(3), July 1992.
5. F. G. Fiamingo, *Unix System Administration*, University Technical Services, Ohio State University, Columbus, Ohio, 1996.
6. S. S. Vogt *et al.*, "HIRES: The High Resolution Echelle Spectrometer on the Keck Ten-Meter Telescope," in *Instrumentation in Astronomy VIII*, D. L. Crawford, ed., *Proc. SPIE* **2198**, pp. 362–375, 1994.
7. J. B. Oke *et al.*, "The Low Resolution Imaging Spectrometer for the Keck Telescope," in *Instrumentation in Astronomy VIII*, D. L. Crawford, ed., *Proc. SPIE* **2198**, pp. 178–184, 1994.
8. J. Percival and R. L. White, "Efficient Ttransfer of Images over Networks," in *Astronomical Data Analysis and Software Systems II*, R. J. Hanisch, R. J. V. Brissenden, and J. Barnes, eds., *ASP Conference Series* **52**, pp. 321–325, 1993.
9. R. I. Kibrick, R. J. Stover, and A. R. Conrad, "CCD Data Acquisition Systems at Lick and Keck Observatories," in *Astronomical Data Analysis and Software Systems II*, R. J. Hanisch, R. J. V. Brissenden, and J. Barnes, eds., *ASP Conference Series* **52**, pp. 277–288, 1993.
10. P. L. Shopbell, J. G. Cohen, and L. Bergman, "Remote observing with the Keck Telescope from California using NASA's ACTS Satellite," in *Telescope Control Systems II*, H. Lewis, ed., *Proc. SPIE* **3112**, 1997.
11. J. W. Percival, "Remote Observing from the Bottom Up: The Architecture of the WIYN Telescope Control System," in *Telescope Control Systems*, P. T. Wallace, ed., *Proc. SPIE* **2479**, pp. 33–38, 1995.
12. A. Conrad, B. Kibrick, and J. Cromer, "Two spectrograph control displays for the W.M. Keck Telescope," in *Instrumentation in Astronomy VIII*, D. L. Crawford, ed., *Proc. SPIE* **2198**, pp. 1151–1157, 1994.