

Remote observing with the Keck Telescopes from multiple sites in California

Robert Kibrick^a, Brian Hayes^a, Steve Allen^a, and Al Conrad^b

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

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

ABSTRACT

Remote observing is now the dominant mode of operation for both Keck telescopes and their associated instruments. Over 90% of all Keck observations are carried out remotely from the Keck Headquarters in Waimea, Hawaii. The majority of Keck observers, however, are affiliated with research institutions located on the U.S. mainland, primarily in California. To observe with the Keck telescopes, most of these astronomers currently travel several thousand kilometers in order to sit in a Keck remote control room located tens of kilometers from the telescopes. Given recent improvements in network infrastructure and facilities, many of these observations can now be conducted directly from California.

This report describes the operation of a Keck telescope remote observing facility located at the UCO/Lick Observatory headquarters on the Santa Cruz campus of the University of California (UCSC). This facility currently enables remote operation and engineering of Keck optical instruments via Internet-2. The facility was initially located in temporary quarters and became operational on a trial basis in September 2001. In June 2002, the facility moved to permanent quarters and became fully operational in July 2002.

We examine in detail issues of Internet network bandwidth and reliability, and describe the design, routing implementation, and operation of an automated fall-back network utilizing dialed ISDN telephone circuits and routers. This report also briefly describes the status of efforts to establish Keck remote observing facilities at other California sites, and how the fall-back network design could be expanded to support multiple sites.

Keywords: remote observing, remote operation, remote instrumentation, video conferencing, Internet-2, ISDN

1. INTRODUCTION

Over the past several years, the Internet-2 initiative has spurred significant improvements in network infrastructure throughout the U.S. A multi-gigabit-per-second research backbone now spans much of the continent, connecting hundreds of major universities and research facilities. While the Internet-2 segment that connects the mainland to Hawaii does not yet offer gigabit bandwidth, that fiber link has been recently upgraded to an OC-3 circuit, providing 155 megabits per second (Mbps) between the mainland and the island of Oahu, and between Oahu and the city of Hilo (on the island of Hawaii), where the headquarters of most of the Mauna Kea observatories are located. A DS-3 circuit from Hilo to Mauna Kea currently provides 45 Mbps between those sites, and another DS-3 circuit connects from Mauna Kea to Waimea, where the headquarters for the Keck Observatory and Canada-France-Hawaii Telescope are located. Similar initiatives within California have provided high-speed connectivity between the various institutions that employ the major share of Keck telescope observers. Together, these high speed networks provide the potential for many Keck telescope observing programs to be conducted remotely from California if appropriate remote observing facilities are established.

For the past several years, we have worked to establish such a facility at UCSC. Details of the motivation and planning for that facility and the software architecture on which it is based were discussed in two earlier reports.^{1,2} Key aspects of those reports are excerpted in this introductory section to provide background for what follows.

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

R.K.: Email: kibrick@ucolick.org; <http://www.ucolick.org/~kibrick>; Telephone: 831-459-2262; Fax: 831-459-2298;

S.A.: Email: sla@ucolick.org; Telephone: 831-459-3046; Fax 831-459-2298;

A.C.: Email: aconrad@keck.hawaii.edu; Telephone: 808-881-3812; Fax: 808-885-4464

1.1. Telescope Time Assignment and Scheduling

Keck telescope time is scheduled in terms of a classical observing model in which astronomers are assigned a specific set of dates for each observing run. Unlike some recent telescopes,³ the Keck telescopes were not designed to support queue-scheduled observing. While Keck observations have occasionally been carried out by local service observers on behalf of remote astronomers, service observing is usually not requested by Keck observers nor generally supported by the Observatory.

Nearly all Keck science observations are performed directly by the astronomers (and their observing teams) who have been allocated telescope time. Time is assigned twice a year, so each Keck telescope schedule covers a six-month semester. Each semester is divided into a series of observing runs, each of which begins on a specific date. A given run is assigned to a particular observing team and some teams may be allocated more than one observing run in a given semester.

Currently, the duration of an observing run varies from half a night (i.e., two different observing teams split the time for that night) to five nights. Approximately half of all science observing runs on the Keck telescopes are for only half a night or one night.²

1.2. Run duration versus observing location

Observing time on the Keck telescopes is currently allocated to astronomers from the University of California (UC), California Institute of Technology (CIT), National Aeronautics and Spaces Administration (NASA), and the University of Hawaii (UH). Collectively, these institutions (and the various grant funding agencies that support their respective researchers) expend several hundred thousand dollars each year on direct travel costs (e.g., airfare, ground transportation, lodging) incurred by astronomers commuting to the remote operations facility in Waimea. In addition to these direct costs are the indirect costs of researchers' time consumed in travel. Such travel time is particularly inefficient for observing runs that last only a single night or one-half of a night because more time is consumed in travel between the mainland and Waimea than in observation. For such short duration runs, operation from the mainland is more efficient in terms of travel time and costs.

Remote operation from the mainland is also particularly attractive in several other situations. Even for longer duration runs, travel to Waimea can prove difficult in cases where health issues (e.g., sinus or ear infections) make airplane travel problematic, or when family considerations (e.g., impending birth of a child) or academic obligations make it difficult to be away. In addition, since most Keck instruments were designed and built by institutions in California, mainland remote operation allows engineers and technical support staff at those sites to run instrument diagnostics, conduct performance analysis, or perform system calibrations remotely.

Remote operation from the mainland offers particular advantages for short duration runs involving routine observing programs, while remote operation from Waimea offers advantages for longer duration runs.

1.3. Remote Observing from Keck Headquarters in Waimea

Most observations with the Keck telescopes are now carried out from Waimea. The town of Waimea is located 32 kilometers northwest of the Mauna Kea summit. Observers who work from Waimea operate from one of the two remote operations rooms located at Keck Headquarters (see Fig. 1). The work stations in these rooms communicate with data taking and telescope control computers at the summit via a dedicated, 45 Mbit/sec DS3 link.

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, a 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 is usually located at the summit, but in some cases, will operate the telescope from the same control room in Waimea from which the observer is running the instrument.

A dedicated, point-to-point video-conferencing system (PictureTel Venue 2000) links each remote operations room in Waimea with its corresponding telescope control room at the summit. In the typical case, where the observer and the OA are separated, this system has proven to be critical for successful observing. We have found that the visual cues that come from body language and from knowing who else is in the room are essential

for a group of people collaborating on the operation of the telescope and instrument. Audio quality has also proved to be essential. The audio quality of this system is superior to that delivered by a standard telephone connection.

The Waimea facility also provides a very well appointed visiting scientists quarters (VSQ) located a short walk from the remote operations rooms. The VSQ provides ten separate suites, each affording a quiet and dark location where observers can sleep during the day. The VSQ also includes a common kitchen and dining area, a library, a laundry room, and a recreation room. A shopping center with a supermarket and several restaurants is within walking distance from the Keck HQ. Such amenities are of particular importance for observers engaged in longer duration observing runs.

A further attraction of the Waimea facility is the availability of the full complement of Keck support staff during the daytime. 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 utilize the facility more efficiently. Direct contact between observers and technical staff stands out as one of the primary advantages of remote observing from Waimea.

1.4. The Keck Remote Observing Facility at U.C. Santa Cruz

1.4.1. Assumptions

The remote observing facility at UCSC is primarily targeted towards observers who are scheduled for short duration observing runs and who live within commuting distance to the facility. It is not intended to duplicate the Waimea facility nor to operate independently of it. Rather, the Santa Cruz facility is an extension of the facility in Waimea. The two are intended to operate in collaboration, sharing resources where practical.

Accordingly, there are no plans to provide facilities equivalent to the Keck VSQ, since it is clearly not practical nor cost-effective to replicate such facilities at each mainland remote observing site. Remote observers using the Santa Cruz facility are expected to make their own eating and daytime sleeping arrangements. Most choose to sleep at home during the day.

Similarly, we have no plans nor budget to provide on-site support astronomers. Rather, we rely on the existing instrument support staff at Waimea and provide video-conferencing and shared software environments so that they can most effectively support the observers at the mainland site. First-time users of an instrument are required to observe from Waimea so that the Keck support astronomers are not required to train novice users remotely.

Given these assumptions under which the Santa Cruz facility operates, there are several reasons why observers with longer duration runs may prefer to observe from Waimea. First, it is easier to justify the travel costs and travel time to Waimea when these can be amortized over multiple nights of observing, as opposed to short-duration runs, where the travel time exceeds the observing time. Second, the homes of many mainland observers are neither dark nor quiet during the daytime. Thus, for runs that last longer than one night, observer efficiency on subsequent nights may be significantly higher as a result of the superior daytime sleeping arrangements provided at the Keck VSQ. Third, Waimea provides a certain level of isolation from on-campus sources of interruptions and distractions that can sometimes interfere with observation planning or data analysis activities. While such interruptions can possibly be deferred for a half-night or one-night run, they become harder to avoid when a run lasts several nights.

In some cases, an observing run may be conducted by an observing team in which some of its members observe from Waimea while others elect to observe from the mainland. In particular, the mainland remote observing facility can provide opportunities for students to participate in observing runs despite limited availability of travel funds. In the case of observing teams having 4 or more active participants, Keck Observatory support staff foresee advantages to distributing such teams between sites rather than trying to fit the entire team into the limited space of the remote operations room in Waimea. If an observing team is composed of observers from widely separated mainland sites (e.g., Santa Cruz and San Diego), then simultaneous remote operation for multiple mainland sites may prove desirable (See Sect. 5).

An observer operating a Keck instrument remotely from the Santa Cruz facility would conduct observations much the same as if operating remotely from Waimea. In both cases, the observer uses a secure protocol to remotely log in to the given instrument's observing computer, which is located on the Mauna Kea summit. In both cases, the observer would run the identical software on that same observing computer, and the results would be displayed via X11 on the observer's workstation at the remote site. In addition, in both cases, the observer would use a similar video-conferencing system to interact remotely with the observing assistant at the summit who is operating the telescope. The major operational difference is that the remote observer on the mainland interacts with the support astronomer in Waimea via video-conferencing, while a remote observer in Waimea has direct face-to-face contact with the SA. A more subtle difference is that the remote observer on the mainland is separated by several time zones from the OA and SA, and thus breaks for meals during daytime setup activities are often not in sync.

A key objective of the mainland facility is to provide observers with the flexibility to choose the observing site that is most practical and productive for them. That choice depends on the needs of the members of their observing team and the specific circumstances and duration of a given observing run. That objective can best be achieved if the mainland sites are operated in close collaboration and cooperation with the Waimea facility.

1.4.2. Equipment and services provided

The remote observing facility in Santa Cruz (see Fig. 2 provides a subset of the equipment and services provided in Waimea. These include:

1. A similar complement of remote observing workstations that run the same software as those in Waimea;
2. External disk, tape, and writable CDROM drives for image storage and transport;
3. A video-conferencing system that connects to Waimea and Mauna Kea via IP-based H.323 protocol;
4. A standard telephone and speaker-phone for local and long distance calls;
5. ISDN telephone lines and routers to provide a backup data path to Mauna Kea (see Sect. 3);
6. An uninterruptible power supply (UPS) to protect against all but the most lengthy power outages;
7. Support equipment including printer, copy machine, and FAX;
8. Amenities including a microwave oven, refrigerator, and bottled water dispenser with hot and cold taps.

These items are divided between two nearby rooms, with the computing and video-conferencing equipment in one room and the support equipment and amenities in the other. These two rooms are located adjacent to the offices of the UCO/Lick Observatory Scientific Programming Group. This group developed the software for several of the Keck optical instruments*. During the day, its staff can provide expert software assistance to observers using those instruments and assist in resolving any initial configuration or instrument start-up problems that arise during the afternoon.

Because of the 2 to 3 hour time difference between California and Hawaii (California observes Daylight Savings Time while Hawaii does not), the remote observers' end-of-night activities overlap with the start of the work day at the remote site. This enables the software staff to hear first hand from the observer about any software problems that occurred during the night or about any suggestions for software improvements. Such direct contact between the observers and the software staff has proven beneficial to both and results in improved software for current and future instruments.

*The Keck High Resolution Echelle Spectrometer (HIRES), the Echellette Spectrograph and Imager (ESI), and the Deep Imaging Multi-Object Spectrograph (DEIMOS)

2. NETWORK RELIABILITY CONCERNS

The successful operation of a mainland remote observing facility is critically dependent on the reliability and bandwidth of the network connections between Waimea, Mauna Kea, and the mainland facility. Astronomers are unlikely to embrace this mode of operation unless they can be confident that potential network performance problems will not significantly reduce observing time or efficiency.

2.1. Experience with the Waimea to summit link

Remote operation from Waimea was initially viewed with considerable skepticism by a number of observers. This skepticism was due in part to concerns about the performance of the network between Waimea and Mauna Kea. The record, however, indicates that the reliability of the summit to Waimea link has been extremely high, and it has earned the confidence of most astronomers.

The summit-to-Waimea path consists of a single segment that spans only about 40 kilometers. It involves only a single provider (GTE/Hawaiian Telephone), and Keck directly controls the active network gear at each end. The entire link is private and therefore secure. If a connectivity problem develops, it is relatively straightforward to isolate the source and assign responsibility for repair. If the failure is at the summit, there are already staff on site who can quickly swap in spare hardware, and if it is in Waimea, there are on-call staff who can respond.

2.2. Concerns regarding the wide area link

The Mauna Kea-to-UCSC link is much longer and complex. It consists of several dozen segments that span three thousand kilometers, and few segments are backed up by redundant network paths. It involves six different network organizations, including UCO/Lick Observatory, UCSC Communications and Technology Services, CalREN-2, Abilene, UH, and Keck Observatory.

If a network connectivity problem develops during observing, troubleshooting requires isolating the source, assigning responsibility, and getting the responsible party to correct the fault. While a variety of network troubleshooting tools (e.g., ping, traceroute, mtr) can often rapidly isolate the portion of the network that is at a fault, contacting the responsible party can be difficult, especially if the problem is not at one of the endpoints (i.e., Santa Cruz or Mauna Kea). Even if the observer or support staff is able to contact those responsible for maintaining the malfunctioning network segment, working with them to resolve the problem can often be a lengthy and difficult process, especially when such faults occur in the middle of the night (which is when optical and infrared astronomers using ground based telescopes typically observe).

Under our current operating model for mainland remote observing, we are not budgeted to provide any staff for troubleshooting wide-area-network problems that may arise during observing, and even if we were, in most cases, all that such staff could do would be to report the problem rather than to solve it. Keck observers operating remotely from the mainland certainly should not be expected to troubleshoot or resolve such problems themselves. Further, Keck Observatory policy specifies that such wide-area-network troubleshooting activity is not within the scope of responsibility for either the observing assistant or support astronomer. (However, solution of any network problems within the Keck observatory summit network remains the responsibility of Keck support staff, just as it would be were the astronomer observing remotely from Waimea.)

While interruptions of the network between Mauna Kea and the mainland have been relatively infrequent, occasional interruptions do occur. Most last only a few minutes, but some have persisted for several hours. On one occasion the interruption lasted most of a day, when an air conditioning unit in a network equipment room in Hilo failed. Hardly any of these interruptions have been located at either the Santa Cruz or Mauna Kea end points, and thus most could not have been resolved directly either by Lick or Keck Observatory support staff.

Rather, most of the network interruptions we have encountered thus far have occurred either at one of the Abilene gigapops or within the University of Hawaii network domain. About half of these interruptions turned out to be due to scheduled network maintenance activities[†] which we only found out about after the fact; the remaining outages have been the result of various hardware failures or accidental misconfiguration of routers.

[†]Many network maintenance activities are scheduled to occur in the middle of the night, since most network users who are not astronomers are likely to be asleep at that hour.

While we have recently been added to the list of sites that receive advance notice of scheduled maintenance activities on the Abilene backbone network, we still receive little or no notification of impending maintenance activities that occur within other network domains along the path between Santa Cruz and Mauna Kea. When we have advanced notification of a scheduled outage, observers can be better prepared for dealing with the cutover to the backup path. In some cases, it may also be advantageous for them to manually force a cutover to the backup path in advance of the scheduled outage, so as to avoid having an automatic cutover occur during a time critical portion of the observing program.

In the absence of a backup data path, the consequences of such network interruptions for observers working from a mainland remote observing site can be severe (especially if the interruption lasts more than a few tens of seconds) and can result in significant loss of observing time. During such interruptions, mainland observers are unable to operate the instrument or receive any telemetry from it or the telescope. (They also lose their video-conferencing connection to the summit and to Waimea, which although inconvenient, is not fatal, since a telephone call can be used instead.) If the outage persists, network connections time out, display windows close, and sessions are disconnected; this results in additional lost observing time, since it can take several minutes to reestablish these sessions and to bring up and configure the various display windows once network connectivity is finally restored. In the meantime, the remote observers typically have little or no information regarding the source of the interruption or how long it will take to repair, and this can cause extreme frustration and even panic.

The consequences of such network interruptions are perceived to be so severe as to require a backup data path. It is currently the policy of Keck Observatory that in the absence of such a backup path, remote observing from the mainland is not permitted unless one member of the observing team is in Waimea and able to carry on the observing program in the event of a network interruption. Further, an observing program in which the entire observing team intends to operate from a mainland site will not be permitted until the backup data path from that site has been demonstrated to work.

3. THE BACKUP DATA PATH

To protect against such network interruptions, we have installed a backup data path between the remote observing room in Santa Cruz and the Keck telescopes using dialed ISDN telephone lines.⁴ Three ISDN BRI lines terminate directly in the remote observing room, while three others terminate in the Keck I telescope dome. (By providing a backup ISDN pathway directly between these two locations, we bypass outages that might occur anywhere along the Internet-2 path, including even those within our own building or the UCSC campus.) Network routers (Cisco 2621) containing both conventional fast Ethernet and quad ISDN BRI interfaces have been installed at both sites[‡] (See Fig. 3). These routers provide both inverse multiplexing and dial-on-demand capabilities for the ISDN lines. Static routes have been established on the facility routers at both sites so that all packets flowing between Keck and the Santa Cruz facility are routed through these two ISDN routers.

3.1. Overview of the dialed ISDN backup solution

Many large institutions have used dedicated point-to-point leased lines to provide network connectivity between the various sites of their enterprise. Often, these dedicated leased lines are backed up by dialed ISDN circuits. Accordingly, many ISDN routers already provide built-in methods for dynamically re-routing network traffic from a dedicated point-to-point line to a set of dialed ISDN backup circuits. The mechanism is quite simple. If the router detects a loss of link status from the interface to the dedicated line, traffic is re-routed to the interface for the dialed ISDN circuits. When the router's interface for the dedicated line detects that link status for that line has been restored, the dialed circuits are shutdown and traffic re-routed to the dedicated circuit.

A similar method cannot be used to detect most losses of connectivity across an Internet path. Since in most cases the interruption will not be due to a problem with the local Ethernet segment, the router's Ethernet interface will continue to report valid link status. Thus, we also cannot rely on the router's Ethernet interface to directly detect any interruption of network connectivity that occurs beyond the local segment.

[‡]Both the routers and the ISDN ST interfaces at both sites are on battery backup power.

Unfortunately, at the time we were designing our backup path, we did not find any ISDN routers which provided a built-in method for using dialed ISDN lines as a backup path for anything other than a dedicated point-to-point circuit. To implement a method that would enable dialed ISDN lines to serve as a backup path to a long-distance Internet connection (which is neither dedicated nor point-to-point) we had to combine a number of different methods that the router does provide: general route encapsulation (GRE) tunneling, bandwidth-on-demand, floating static routes, and open shortest path first (OSPF) routing.

3.1.1. GRE Tunneling

A tunnel is simply an encapsulation method that hides data and also simulates a point-to-point link. It can be used to make indirectly-connected physically-distant sites appear logically as if they are directly-connected neighbors. There are different types of tunnels, but because our main concern is encapsulation, a generic route encapsulation (GRE) tunnel was selected for simplicity.⁵ Unfortunately, backup functionality is not built into a GRE tunnel, which means the tunnel cannot detect a underlying loss of Internet network connectivity on its own.

3.1.2. Bandwidth-on-Demand

Bandwidth-on-demand is an automatic method by which successive ISDN calls are placed to establish as many ISDN B-channel connections as needed to satisfy the currently demanded bandwidth. As more bandwidth is demanded, more calls are placed, and the respective B-channels are bonded together to form a logical channel whose aggregate bandwidth is the sum of the bandwidths of the bonded B-channels. As the demand falls off, calls are automatically terminated and the corresponding bandwidth is reduced.

3.1.3. Floating static routes

A floating static route to a destination is simply a static route whose administrative distance exceeds that of any dynamic route to that same destination. It is considered floating because it only becomes active when there is no other route to that destination.

3.1.4. OSPF Dynamic Routing

The OSPF routing protocol is one of several such protocols typically used to maintain routing tables based on the current connectivity status between a given router and its neighboring routers within the local OSPF domain. Each of these routers running OSPF will transmit OSPF “HELLO” packets to its neighboring routers every 10 seconds. These packets serve two functions. They are used to exchange routing information between neighboring routers, and they provide a mechanism for determining which of the neighboring routers is currently reachable.

If a neighboring router fails to respond to repeated HELLO packets, then OSPF concludes that that router is unreachable, and it removes from its local routing table any dynamic routes that depend on that unreachable router. Once that neighboring router resumes responding to HELLO packets, OSPF concludes that that router is once again reachable, and any dynamic routes involving that router are restored.

Unfortunately, like GRE tunnels, OSPF routing, by itself, has no ability to directly detect a loss of network connectivity between two distant sites that are connected by a indirect path which crosses multiple OSPF domains. However, while neither GRE tunnels nor OSPF routing can directly detect such a loss individually, that deficiency can be overcome if both are used together, i.e., if we run a dynamic routing protocol through a tunnel.

3.1.5. How it works: fail-over and fall-back

Our implementation of the backup data path uses these four mechanisms in combination. First, a GRE tunnel across the Internet is defined between the ISDN router in Santa Cruz and its companion router on Mauna Kea. Second, on each of these routers, a floating static route is defined which will route to the ISDN dialer interface any traffic that is addressed to the other site; whenever traffic is routed to that dialer interface, bandwidth-on-demand will automatically cause ISDN calls to be placed. Third, we establish a private OSPF domain that runs across the tunnel so that our ISDN routers appear to each other as sole OSPF “neighbors”.

Under normal operation, the OSPF running on each of our two routers will establish dynamic routes that will route all traffic to the other site via the tunnel. (Those dynamic routes take precedence over the floating static route.) Each router will attempt to exchange OSPF HELLO packets with its “neighbor” (i.e., the other site) via that tunnel. As long as the Internet path between the two routers remains intact, those HELLO packets will continue to flow across the tunnel, and the OSPF on each router will conclude that its neighbor is reachable. Accordingly, each router will maintain the dynamic “tunnel” route to its neighbor.

If the underlying Internet network connection between the two sites is interrupted for more than 30 consecutive seconds (i.e., if each router fails to see a reply from its neighbor to three consecutive HELLO packets)[§], then OSPF will conclude that the neighbor is currently unreachable, and it will remove from its routing table the dynamic tunnel route to that neighbor. Once OSPF has removed that dynamic route, the floating static route to that neighbor becomes active, and all regular traffic to that neighbor is routed to the ISDN dialer interface.

As packets start flowing to that interface, the Santa Cruz router’s bandwidth-on-demand mechanism initiates a sequence of dialed ISDN calls to the respective B-channels of the end router on Mauna Kea. These calls complete within a matter of seconds. Once these calls complete, packet flow between the two routers (and hence the two sites) is restored, albeit at a lower bandwidth than is available over the regular network path. This entire process occurs automatically in less than a minute and without any manual intervention required by staff at either site.

Even as all regular packets are being routed to the ISDN dialer interface, OSPF will continue to send HELLO packets to its neighbor via the tunnel. But as long as the Internet connection remains down, those HELLO packets will never arrive at the other end. However, once Internet connectivity is restored, those HELLO packets will resume flowing across the tunnel in both directions. The OSPF on each router will then conclude that its neighbor is once again reachable via that path, and will re-instate the dynamic tunnel route to that neighbor[¶]. That dynamic route takes precedence over the floating static route, so regular packet traffic will once again be routed to the tunnel rather than to the ISDN dialer interface. Since packets are no longer flowing to that dialer interface, the router’s bandwidth-on-demand mechanism will hang up any ISDN lines that are still connected.

Alternative dynamic routing protocols could also be used to perform this function. OSPF is a link-state protocol that is in contrast to a distance-vector protocol such as the Enhanced Interior Gateway Routing Protocol (EIGRP). Either OSPF or EIGRP could have been chosen; although they run on completely different protocols, their results would be similar. In this case, convergence time is the most important factor, and both OSPF and EIGRP converge faster than any other routing protocol. Other important issues include bandwidth consumption and scalability. Both OSPF and EIGRP are very similar in those respects as well.⁶ We chose OSPF over EIGRP because the former was already running on the UCSC campus network and thus we had a local base of expertise that we could tap. In addition, EIGRP is proprietary to Cisco routers while OSPF is not.

3.2. Operational Results

Fail-over from the primary network path to the backup ISDN path is mildly disruptive because no packets flow between the two sites for either the 30 seconds it takes OSPF to decide that the network is down nor during the additional few seconds it takes for the ISDN calls to complete. In addition, once packet flow resumes over the ISDN path, a 30 to 40 second backlog of accumulated packets suddenly is delivered to the remote observer’s displays. This is mildly amusing for the next few seconds as various displays of dynamically changing data (e.g., clocks and telescope position displays) update with dizzying speed. Once the back-logged packets are processed, display updates resume at their normally expected rates.

[§]Although OSPF can be configured to fail over in less than 30 seconds (i.e, fewer than 3 dropped HELLO packets), we decided on that threshold to avoid triggering use of the ISDN lines for momentary network disruptions or occasional bursts of network congestion that might result in the loss of a single HELLO packet.

[¶]The routers have been tuned to provide appropriate hysteresis so as to avoid excessive route switching in cases where the Internet-2 network path is unstable.

In contrast, fall-back from the backup ISDN path to the primary network path is totally transparent and painless. When fall-back occurs, the only indication one receives is that interactive response time improves significantly. In both cases, the transfer between paths is totally automatic and requires no manual intervention at either site. The ISDN backup path has performed flawlessly, kicking into action in every case where it was needed.

While the backup ISDN pathway provides only a few percent of the bandwidth available via the Internet-2 network path, it allows the remote observer to continue observing (albeit with significantly reduced efficiency) and to maintain context within their observing program so that they can resume working at full efficiency once the regular network path is restored. While some interactive operations (e.g., panning and zooming various image displays) can become annoyingly sluggish while the ISDN path is active, useful observing can still be carried out. Our experience has been that most network interruptions are relatively brief (e.g., comparable to the time required for a router to be rebooted, a cable to be moved to a different interface, or a hot-swappable card to be replaced), so that these periods of reduced efficiency generally do not last very long.

Even more important is the fact that the backup ISDN path has prevented the severe disruptions (e.g., session disconnection and shutdown of display windows) that observers typically experience from a complete loss of network connectivity. The time required to recover from those disruptions can often exceed the duration of the network interruption itself, thus resulting in an even greater loss of observing time.

3.3. Operational Costs

We have installed three ISDN lines at each site, yielding an aggregate bandwidth of 384 Kbps (two 64 Kbps B-channels per ISDN line times three ISDN lines). Should it prove necessary, additional lines could be added. (At some point, it will become more cost effective to switch to an ISDN PRI interface rather than to install more ISDN BRI lines.) Most of the time, these ISDN lines are idle; they accrue long distance toll charges only during those times when Internet-2 connectivity is interrupted *and* connectivity between the two sites is required. ^{||}

Under the current State of California long distance contract, it costs \$0.07 per minute to operate each ISDN B-channel for all state-to-state calls and at all hours of the day. If all six B-channels are in use, this comes to \$0.42 per minute or \$25.20 per hour of operation. The bandwidth-on-demand feature of the ISDN routers attempts to minimize those toll charges by maintaining only as many connections as are needed to sustain the demanded bandwidth. However, most of the time that the primary network is down, the demanded bandwidth is sufficiently high that all six B-channels will be in use. While these ISDN toll charges are variable depending on usage, they have yet to exceed \$50 in any month, since the lengthiest network outages we have observed either occurred prior to the installation of the ISDN lines or occurred during times when remote observing was not in progress. The total fixed lease charges for the three ISDN lines at both sites cost about \$5,000/year, or about \$70 per line per site per month.

The value of Keck telescope observing time is conservatively estimated at \$1 per second per telescope, or \$3600 per hour. This is based on the annual cost of operating the telescopes plus the cost of amortizing the construction cost of each telescope over its expected operating lifetime. From this perspective, these ISDN lines more than pay for themselves if they are able to prevent the loss of as little as two hours of observing time per year. They have already done so for the current calendar year.

3.4. Operational Issues

Routing protocols such as OSPF and EIGRP are typically used not only to optimize the routing of packets but as a mechanism for various Internet service providers (ISPs) to exchange routing information within a given routing domain. However, in our application, the existence of the private ISDN backup path between Santa Cruz and Mauna Kea is not something that we want advertised to other routers on the Internet. Otherwise, for example, were that route to leak out and become a part of more global routing tables, then in the event of a

^{||}Whenever the remote observing facility is not in use (and hence the backup path is not needed), we disable the ISDN dialer interface on the Santa Cruz router to avoid placing any unnecessary calls and incurring needless ISDN charges. At the start of any mainland remote observing run, we re-enable that dialer interface and verify the operation of the backup data path by temporarily severing the connection between our ISDN router and the UCSC campus network.

network interruption on the circuit between Oahu and the mainland, those routers might try to re-route some of their traffic over the very limited bandwidth of our ISDN circuits.

Accordingly, when we first discussed with our local campus network engineer our plans for an ISDN backup path using OSPF, he was extremely skeptical of the idea and concerned that if our OSPF routes were to leak out to the campus and beyond, that they might create routing problems for the campus and regional networks, and that we might find our ISDN circuits flooded with extraneous traffic, thereby increasing their operating cost and greatly reducing their utility for remote observing. As a result, in implementing our router configurations, considerable effort was expended to ensure that the OSPF routing in our two ISDN routers is defined as its own independent OSPF routing domain with no connection whatsoever to any other external OSPF routing domain. Our two ISDN routers advertise their respective OSPF routes only to each and to no other routers.

The ISDN backup data path also raises various network security issues. For obvious reasons, many of those details cannot be safely discussed in such a widely available publication. We simply note that the routers themselves enforce access lists that are extremely restrictive and that firewall technology is also employed to provide additional protections. In addition, the ISDN routers themselves are equipped to utilize the caller I.D. feature provided by the telephone lines and will only accept ISDN connections from known sites. Calls are only initiated from the mainland to Mauna Kea. Under normal operation, the ISDN routers in Santa Cruz do not accept any incoming ISDN calls. In addition to the security implemented within the router, other software layers enforce both host-based and user-based authentication.

4. RECENT OPERATIONAL EXPERIENCE

Our remote observing facility at UCSC operated on a trial basis in rather cramped temporary quarters from September 2001 through June 2002. At the end of June, it moved to larger, permanent quarters which are more comparable in size to the remote observing rooms in Waimea. The electrical power in the new quarters is provided by a massive constant online UPS, which should enable the facility to continue operating through all but the most lengthy of power outages.

Since September 2001, the facility has been used to conduct remote observations with two Keck optical imaging spectrographs: the Echellette Spectrograph and Imager (ESI), and the Low Resolution Imaging Spectrograph (LRIS). During each of the remote observing runs, the backup data path was activated at least once (and often several times), typically for periods of between 5 and 20 minutes each. During one of these runs, the backup path activated 4 times during a single two hour period; we later found out that that period coincided with a window of scheduled maintenance on an Abilene router.

4.1. Remote engineering and commissioning support

In addition to its somewhat intermittent use for remote observing, the facility has been quite heavily used for remote engineering and support of existing Keck optical instruments such as HIRES and ESI. But its heaviest usage to date has been in support of commissioning activities for the Deep Imaging Multi-Object Spectrograph (DEIMOS), which was successfully commissioned on the Keck II telescope during June and July 2002. During those months, the facility was in use nearly every day and on many of the nights. It enabled technical support staff in Santa Cruz to conduct an extensive array of reliability tests and calibration sequences without the dislocation and expense of traveling to Mauna Kea or Waimea.⁷

4.2. Unexpected use of the facility following September 11, 2001

The tragic events of September 11, 2001 resulted in many unanticipated and painful disruptions to the lives and plans of millions of people. One such disruption was the complete grounding of all U.S. civilian air traffic for a period of several days, including flights between the mainland and Hawaii. A team of astronomers from the California Institute for Technology (CIT, or Caltech) was scheduled to fly to Hawaii on September 13 to begin a 5-night run with LRIS on Keck-I. Suddenly, they had no practical way to get to either Waimea or Mauna Kea to conduct their observing run.

On the afternoon of September 12, the Caltech team asked if they could attempt their observations from the UCSC remote observing facility. We had never planned for the facility to be used by non-local observers

nor for such a lengthy observing run. Neither had we operated LRIS from UCSC before, but we agreed that there was little to be lost by trying. The Caltech team left Pasadena early on the morning of September 13 and drove the roughly 600 kilometers to Santa Cruz (since air traffic was still grounded), arriving late in the afternoon. That evening they were successfully operating LRIS from UCSC well before sunset in Hawaii; all five nights of the run were successfully carried out from UCSC. According to the nightly observing efficiency statistics maintained by Keck Observatory (these measure the total time during the night that the instrument's shutter was open and collecting photons versus the number of hours of nighttime that the telescope dome was open), the efficiency achieved during those five nights of LRIS remote observing from Santa Cruz equaled or exceeded the efficiency of the preceding LRIS observing runs conducted from Waimea under comparable weather conditions. The performance of the facility under these somewhat trying circumstances greatly exceeded everyone's expectations. It also provided a compelling demonstration of the effectiveness of our X-display-based approach to remote operation.²

The most difficult problem that arose during this observing run were the sleeping accommodations for the Caltech team, as we were unable on such short notice to obtain lodging for all five nights at a single location. Since motel check-in and check-out times were in conflict with the hours that the observers would either be sleeping or observing, on several nights charges were incurred at two different motels for the same night. Motels are also not designed for daytime sleepers, since that is when housekeeping activities normally occur, and these are typically very noisy. If a remote observing facility intends to support on a routine basis multi-night observing runs for non-local observers, then it should provide adequate daytime sleeping accommodations similar to the VSQ facility at Waimea, since motels cannot realistically meet this need.

4.3. Practical Lessons Learned: LCD screens .vs. CRT monitors

As noted in Sect.1.3, video-conferencing is a valuable component of remote observing. Accordingly, it is important to take note of its implementation details. Most video-conferencing systems utilize video compression techniques that transmit a complete image frame at intervals much less frequent than the video frame rate. In between such full frames, the compression scheme transmits only those image features that have changed since the last full frame was transmitted.

In IP-based video-conferencing systems, the instantaneous bandwidth required to transmit the compressed video stream thus depends on the amount of motion in the scene being transmitted. The greater the amount of motion, the more bandwidth is required. If the available bandwidth is capped, then increased motion reduces the amount of bandwidth available for overall resolution of the image. Most users of such video-conferencing systems have observed this phenomenon: if their conference partners gesture with rapid hand motions, the overall resolution of the transmitted image will momentarily be degraded, while if they sit still, the image becomes much sharper.

Because of this effect, cathode ray tube (CRT) monitors should be avoided when specifying the hardware components for the observer's workstations in a remote observing room (as well as the monitors for the workstations for the observing assistant at Mauna Kea and those for the support astronomer in Waimea). Due to the beating between the vertical retrace frequencies of the video conferencing camera (typically 24 to 30 Hz) and that of CRT monitor (60 to 90 Hz), coupled with the very sharp peak in CRT phosphor intensity, if the camera is pointed towards a CRT monitor, the resulting camera image exhibits a very rapid and annoying "flicker". Worse yet, the video-conferencing system perceives this flicker as relevant motion, and wastes precious bandwidth transmitting it to the other end-point of the conference. As a result, the quality of the transmitted scene is unnecessarily degraded.

This problem can be reduced by arranging the layout of the remote observing room so that the video-conferencing camera points at the observer but avoids pointing at any CRT screens. An even better solution is to replace the CRT screens with LCD screens, since those do not generate flicker in the camera image. An additional advantage of LCD screens is that it is sometimes extremely useful to be able to point the video-conferencing camera directly at the OA's or SA's screen, especially in those cases where the display generated by a particular application is not amenable to distribution via a shared virtual desktop. Further advantages of LCD screens include higher reliability, lower power consumption (and thus a longer run time on UPS power during a power outage), lower heat dissipation, and smaller footprint.

5. EXTENDING REMOTE OBSERVING TO OTHER CALIFORNIA SITES

The success of the mainland remote observing facility at Santa Cruz has motivated other institutions to begin work on establishing similar facilities at their sites. A remote observing facility at Caltech is now nearly complete (including installation of ISDN lines), and will become fully operational once the ISDN router configurations there and at Mauna Kea have been updated with the appropriate tunnel, OSPF, and ISDN line parameters. A similar facility is currently being assembled at the San Diego campus of the University of California (UCSD), and other UC campuses are also considering similar plans. All of these sites connect to the CalREN-2 high-speed research network and thus all share similar Internet-2 network bandwidth to Mauna Kea.

While the backup data path model should scale relatively easily to additional sites, it raises various operational and administrative questions with regard to the shared use of the ISDN lines and router at the Mauna Kea summit. So long as only a single mainland site is conducting remote observations on any given night, there is no difficulty, since that site can have exclusive use of the ISDN lines and router on Mauna Kea. Currently, the Keck Observatory policy is that they will only allow mainland remote observing from a single site on any night, so the problem has at least been temporarily resolved on an administrative level. Longer term, as additional remote observing facilities become operational, conflicts between sites for use of the shared Mauna Kea resources are likely to arise, and we will soon need to consider adding more ISDN lines and one or more ISDN routers to the Mauna Kea summit. While there is little technical difficulty involved in implementing such additional infrastructure, the economic and administrative aspects will likely prove more challenging.

6. REMAINING CHALLENGES

Most Internet-based communication relies on the transmission control protocol/internet protocol (TCP/IP). Because the initial design for TCP/IP was optimized for use either on short, high-speed networks or slow long-distance networks, long-distance high-speed networks such as our Internet-2 path to Hawaii present significant challenges with respect to achieving optimal use of the available bandwidth. Careful tuning of TCP/IP parameters (e.g., the TCP window size) and enabling of more advanced features (e.g., selective acknowledgment) is needed to achieve optimal performance. Significant research efforts, such as the Web100 Initiative,⁸ are currently underway to develop more automated schemes for performing such TCP/IP tuning. In our case, the task of optimal TCP/IP tuning is even more difficult, because the tuning parameters that would be optimal for operation over the Internet-2 path are significantly different than those that are optimal when running over the ISDN backup path. This remains a significant challenge and one towards which we intend to direct further effort.

During this year our remote observing facility at UCSC has successfully made the transition from an experimental trial to an operational facility, and other sites will be making similar transitions in the coming months. In some respects, many of the technology aspects of these facilities have moved forward more rapidly than have the administrative and policy aspects. The most significant remaining challenge is to develop a coherent administrative and policy structure that will enable these relatively autonomous facilities to function effectively and cooperatively with each other and with the central facility in Waimea.

7. CONCLUSION

While operation from Waimea will continue to provide advantages for longer duration observing runs, remote operation from the mainland should provide significant reductions in travel costs and travel time for short-duration runs. The extent to which it reduces overall costs remains to be seen. In the case of telescopes and instruments in Chile that were remotely operated from a European Southern Observatory (ESO) facility in Germany during the 1990s, the costs of a dedicated satellite communications channel plus the costs of providing on-site instrument specialists at the remote facility effectively cancelled out any savings in travel costs.⁹ By avoiding the high costs of a dedicated communications channel and leveraging the existing investment in Waimea facilities and staff, the mainland sites should yield a net reduction in overall costs. Over the long term, oil prices and travel costs are likely to continue to climb, while advances in optical communications technologies should cause communications costs to fall. Those trends should make mainland remote observing an increasingly cost-effective option.

ACKNOWLEDGMENTS

We wish to acknowledge the combined efforts of the University of Hawaii, the Gemini Telescope consortium, the National Science Foundation, the Department of Defense, the University Corporation for Advanced Internet Development, and the Corporation for Education Network Initiatives in California toward the establishment of a reliable, high-bandwidth network pathway between the Mauna Kea summit and the U.S. mainland. We also deeply appreciate the efforts of UCSC campus network engineer Jim Warner, Keck Observatory system manager Jon Chock, and Eric Berling and Louis Kemp of Cisco systems for their advice in designing an ISDN backup network. We also thank Keck Observatory support astronomers Greg Wirth and Bob Goodrich for their assistance in remotely operating the LRIS, ESI, and DEIMOS instruments from Santa Cruz, and software engineers Liz Chock and Julia Simmons for their helpful assistance in resolving software issues. We would especially like to thank David Sprayberry for his efforts to help define a rational policy basis for mainland remote observing. We also acknowledge the efforts of the Lick Observatory Network Information and Computing Systems group, including Jeff Mallory, Kyle Lanclos, and John Rickard, for their assistance in configuring the remote observing facility at UCSC.

Finally the authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had with the indigenous Hawaiian community. We are most fortunate to have had the opportunity to work there and conduct observations from this very special mountain.

REFERENCES

1. A. R. Conrad, R. Kibrick, and J. Gathright, "Remote Observing with the Keck Telescopes," in *Telescope Control Systems II*, H. Lewis, ed., *Proc. SPIE* **3112**, pp. 99–110, 1997.
2. R. I. Kibrick, S. L. Allen, and A. R. Conrad, "Remote observing with the Keck Telescopes from the U.S mainland," in *Advanced Global Communications Technologies for Astronomy*, R. Kibrick and A. Wallander, eds., *Proc. SPIE* **4011**, pp. 104–116, 2000.
3. T. Boroson, J. Davies, and I. Robson, eds., *New Observing Modes for the Next Century*, vol. CS-87, Astronomical Society of the Pacific, San Francisco, 1996.
4. B. Hayes, "Achieving Dependable Network Connectivity between the Santa Cruz Remote Observing Facility and the Keck Telescopes in Mauna Kea, Hawaii," *UCO/Lick Observatory internal report*, 2001.
5. Cisco Systems, "Configuring logical interfaces," <http://www.cisco.com/univercd/cc/td/doc/product/software>, 2001.
6. S. Ballew, *Managing IP Networks*, O'Reilly, 1997.
7. S. Faber *et al.*, "DEIMOS spectrograph for the Keck II telescope: integration and testing," in *Instrument Design and Performance for Optical/Infrared Ground-Based Telescopes*, M. Iye, ed., *Proc. SPIE* **4841**, 2002.
8. Web100 Project, "Web100 Concept Paper," <http://www.web100.org/docs/>, 2002.
9. A. A. Zijlstra, A. Wallander, L. Kaper, and J. A. Rodriguez, "Remote Observing at the ESO NTT and CAT Telescopes," *Publ. Astro. Soc. of the Pacific* **109**, pp. 1256–1263, November 1997.




Figure 1. Keck remote operations room in Waimea



Figure 2. Santa Cruz remote observing room

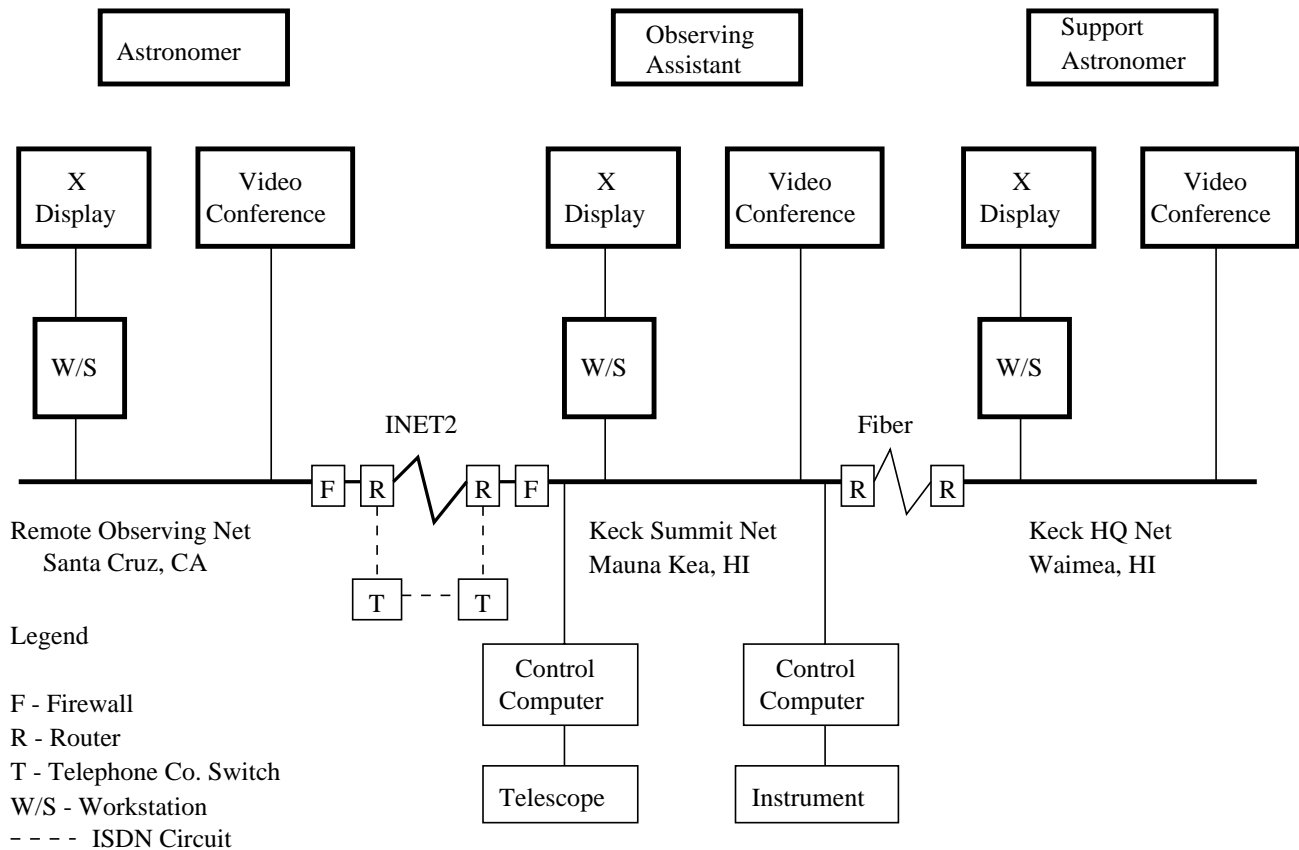


Figure 3. Global overview of connectivity between Santa Cruz, Mauna Kea, and Waimea