

Context-Linked Virtual Assistants for Distributed Teams: An Astrophysics Case Study

Sarah S. Poon

Physics Division
Space Sciences Lab
Berkeley, CA 94720 USA
sarahpoon@gmail.com

Rollin C. Thomas

Physics Division
Lawrence Berkeley
National Lab
Berkeley, CA 94720 USA
RCThomas@lbl.gov

Cecilia R. Aragon

Computational Research Division
Lawrence Berkeley National Lab
Berkeley, CA 94720 USA
CRAragon@lbl.gov

Brian Lee

Lawrence Berkeley
National Lab
Gatan Inc.
Pleasanton, CA 94588 USA
BLee@gatan.com

ABSTRACT

There is a growing need for distributed teams to analyze complex and dynamic data streams and make critical decisions under time pressure. Via a case study, we discuss potential guidelines for the design of software tools to facilitate such collaborative decision-making. We introduce the term *context-linked* to characterize systems where both task and context information are included in a shared space. We describe a novel, lightweight, context-linked event notification/virtual assistant system developed to aid a cross-cultural, geographically distributed team of astrophysicists to remotely maneuver a custom-built instrument under challenging operational conditions, where critical decisions must be made in as little as 45 seconds. The system has been in use since 2005 by a major international astrophysics collaboration. We describe the design and implementation of the event notification system and then present a case study, based on event log analysis and user interviews, of its effectiveness in substantially improving user performance during time-critical science tasks. Finally, we discuss the implications of context linking for supporting common ground in distributed teams.

Author Keywords

Computer-mediated communication, context-aware computing, notification systems, social awareness, collaboratories, cross-cultural collaboration, astrophysics.

ACM Classification Keywords

H5.3. [Information interfaces and presentation]: Group and organization interfaces—Computer-supported cooperative work.

INTRODUCTION

In a growing number of operational domains, geographically separated collaborations must collectively make critical decisions under time pressure while evaluating complex, dynamic data streams. First responders, pilots and air traffic controllers, and hospital staff are some common examples. Certain scientific collaborations also fall into this category.

Astrophysicists cooperatively operating large telescopes need to monitor complex and ever-changing data while maneuvering equipment under time pressure within tight operational constraints.

There is an obvious need for guidelines and examples of collaborative applications to facilitate effective decision-making in such situations. One obstacle to the development of such applications is the difficulty of gaining user acceptance of novel, custom-developed collaborative tools, especially in time-critical domains. For this reason, it makes sense to build lightweight tools that integrate with the existing collaborative environment. Further, such tools must be efficient and effective, providing streamlined, timely access to critical information. Event notification must be appropriately but not overly intrusive.

Studies from the aviation domain have demonstrated that *scene-linked* symbology for head-up displays (synthetic context information projected on a pilot's cockpit windshield, that appears to be placed within and move with the external physical environment) minimizes inappropriate distractions and provides significant performance advantages [2, 13]. We propose an analogy for computer displays, and define the term *context-linked* to describe a system where both task information (immediately pertinent to the job at hand) and context information (background processes and events in the environment) are directly included in a shared communication space. In our case, we chose to integrate an event notification system, which provided both task and context information, directly into the tool providing the primary means of team communication.

Our lightweight context-linked event notification/virtual assistant system (referred to here as “Bert”) was integrated with an Internet instant messaging (or “chat”) client that had already been accepted by the collaboration. Bert was developed for the Nearby Supernova Factory [1], the largest data volume supernova search currently in operation, processing over 50GB of data per night.

There has been little research on software tools for cross-cultural scientific collaborations making critical decisions under time pressure. The major contributions of this paper are 1) the case study of a novel, context-linked notification system for a time-critical domain, and 2) the potential

implications of context linking for the design of collaborative tools for distributed “control room” situations in general.

Operational challenges

An astronomer’s task of observing the heavens with a large, multi-million dollar telescope has some surprising parallels with that of operating a jet aircraft. The astronomer must monitor a large and complex set of operational data while maneuvering the telescope and any attached equipment within particular constraints; although there are some safeguards, violation of these constraints can lead to damage of the telescope or its key components. The telescope is frequently located in an inhospitable environment, such as a remote mountaintop, where hypoxia (lack of oxygen) is a constant threat. Adverse weather conditions such as extremely high winds, rain, sleet, or snow, are relatively common occurrences which necessitate closing the telescope dome immediately to prevent damage to delicate optics. Even the rise of the sun is a danger, as equipment can be damaged by sunlight focused by the telescope. Equipment failures such as the dome becoming stuck also occur. Further adding to the operational challenges, observing must necessarily be done at night, during the lowest ebb of human alertness and cognitive ability.

Due to the expense of telescope time, nights are often fully scheduled with an array of astronomical objects, which must be observed in a particular order in a tightly scheduled timetable for maximum scientific benefit. Further, the phase of the moon, upper air turbulence, fog, and changing cloud conditions throughout the night can cause unpredictable variations in the schedule, as certain objects may be no longer visible at their appointed time slots.

Cross-cultural collaboration

The Nearby Supernova Factory astrophysics collaboration has about 30 members; about half of the scientists are located in the U.S. and the other half in France. On any given night, the telescope is typically operated by a geographically separated group of two to six people. The scientists are in different time zones from each other (France, California, the U.S. East Coast) and from the telescope itself (Hawaii). Correct decisions must be made quickly and collaboratively, although some of the team members have never met each other and come from differing cultures with dissimilar assumptions, and some are not native English speakers.

Team members must be able to analyze and evaluate a large amount of data and rapidly make cognitively demanding calculations, sometimes in as little as 45 seconds, while at the same time being fully aware of changing weather conditions, the approach of daylight and other safety issues. They must focus on individually demanding and precise tasks while maintaining an overall understanding of a large amount of dynamic data affecting the telescope’s operation and safety. Further, astronomers may be operating telescopes they are unfamiliar with, in an interface that is

not in their native language, or remotely, in a different time zone or under different weather conditions, so that normal human diurnal rhythms or other environmental clues are working against their intuition.

Event notification/virtual assistant system: Bert

Due to the difficulty of the overall science task, multiple software tools were developed to ease the process and increase scientific output. Bert was developed to assist the astronomers to maintain awareness of the approach of sunrise and perform time-critical science tasks. Bert is a lightweight event notification/virtual assistant system that operates within the chat communication system used during telescope operation. Bert has two main functions: announcing relevant events, such as completed exposures or the number of minutes until sunrise, and responding to user queries for information. It has been in operational use since 2005.

The remainder of this paper is organized as follows: first, we describe the science background for which Bert was developed, then we discuss related work. The next section describes the overall architecture, design, and implementation of the Bert notification system. We then present a case study and the results of an informal user evaluation of the production system by astrophysicists using the software as they gather and analyze science data. The final sections present discussion and conclusions and suggest directions for further work.

SCIENCE BACKGROUND

One of the grand challenges in astrophysics today is the effort to comprehend the mysterious “dark energy,” which accounts for three-quarters of the matter/energy budget of the universe. The existence of dark energy may well require the development of new theories of physics and cosmology. Dark energy acts to accelerate the expansion of the universe (as opposed to gravity, which acts to decelerate the expansion). Our current understanding of dark energy comes primarily from the study of supernovae [16, 18].

The Nearby Supernova Factory (SNfactory) [1] is an international astrophysics experiment designed to discover and measure Type Ia supernovae in greater number and detail than has ever been done before. These supernovae are stellar explosions that have a consistent maximum brightness, allowing them to be used as “standard candles” to measure distances to other galaxies and to trace the rate of expansion of the universe and how dark energy affects the structure of the cosmos. The SNfactory receives 50-80 GB of image data per night, which must be processed within 12-24 hours to find potential supernova candidates immediately and obtain maximum scientific benefit from the study of these rare and short-lived stellar events.

Promising supernova candidates are sent for confirmation and spectrophotometric follow-up by SNIFS (the SuperNova Integral Field Spectrograph) [12] on the University of Hawaii 2.2m telescope on the summit of Mauna Kea (Figure 1). Candidates are spectrally imaged through a 15 x 15

microlens array on SNIFS; the spectral data are processed at the summit, and then saved to a remote database in France.



Figure 1. University of Hawaii 2.2m telescope on the summit of Mauna Kea, Hawaii, U.S.

In order to maximize the efficiency of each night of telescope operation, the scientists have parallelized operations such as telescope slew with data readout, or calibration lamp warm-up with image acquisition. This has the effect of substantially increasing the number of targets to observe in a night, requiring maximized efficiency from the person conducting the observations, commonly known as the “shifter” (since they were assigned to handle a particular “shift,” or period of time the telescope is in operation). On a typical night, the SNIFS schedule contains approximately 15 supernova spectra, 7 or 8 standard star spectra (“standard stars” are well-studied stars of known brightness and spectra which can be used for instrument calibration), and numerous other associated calibration observations. During each exposure, parallel operations must be conducted while maintaining awareness of the exact moment when the exposure will end, so as to evaluate previously unobserved targets and make an immediate followup decision. Timing is critical because the schedule is so highly optimized.

It is this portion of the supernova followup pipeline, involving full nights of observation remotely operating the SNIFS spectrograph on the UH 2.2m telescope, that the event notification system was designed to address. Since the system was put into production, the average number of supernova targets successfully observed each night has increased, and the number of human errors causing missed targets or uncompleted schedules has decreased.

RELATED WORK

Computer support for distributed scientific collaborations has been applied to various scientific fields. The concept of the collaboratory emerged in the late 1980's, defined by William Wulf in 1989 as a "center without walls, in which the nation's researchers can perform their research without regard to physical location." [21]

The Upper Atmospheric Research Collaboratory (UARC) uses chat to augment the analysis of shared data resources. The UARC interface is a data analysis application with an

embedded chat tool that provides real-time images of ionospheric phenomena and allows for synchronous discussion of the displays. Part of the motivation for this interface stemmed from the fact that the scientists felt unable to respond rapidly to interesting phenomena. [7, 15]

Another study of collaboratories has shown that chat can be used as a grounding and orienting tool for distributed research groups during real time experiments. The chat logs taken during an earthquake engineering experiment illustrated how chat was heavily used to keep the chat participants informed of the status at the physical site of the experiment. This study found that remote participants "needed only a basic sense of this in order to feel adequately grounded as participants in the experiment" and this low information requirement allowed chat to function as an adequate grounding tool. The same study also showed that trouble shooting question and answering was another common usage of the chat during experiments. [3]

In non-collaboratory related work, researchers at the University of Sussex studied the use of a system integrating tickertape CVS notifications and chat software by a software development team. They found that these notifications prompted communication in the chat, including coordination/negotiation of work, and focused code and design discussions. [8]

A number of studies have shown that persistent chat can offer many benefits to a collaboration, including creating a record of collaboration history, helping newcomers join the group and orient themselves more quickly, and facilitating asynchronous communication [6, 9, 11, 17]. Chat has been found to be a useful tool for collaborations generally, despite its limitations, and there have been a number of approaches to augment chat, often by modifying the chat interface in a variety of ways to increase readability or persistence [6, 9-11, 17, 19]. However, none of these augmentation methods included a virtual assistant. Additionally, many of these interfaces were heavyweight research prototypes and/or were not used in a production system.

BERT ARCHITECTURE AND IMPLEMENTATION

Bert was originally a side project, initiated as a response to the problem that scientists operating the telescope, especially those located far from Hawaii, seemed to be unaware of how much time they had left in the schedule. In particular, the scientists felt that without powerful environmental clues (e.g., the sky becoming brighter in the east during morning twilight) remote operators on the other side of the world working in daylight lacked a sense of urgency in completing a schedule.

A number of solutions were considered – for example, the scientists considered having the background color of the VNC (Virtual Network Computing remote desktop) display change as sunrise approached, but in the end opted to capitalize on several advantages provided by incorporating Bert into the existing instant messaging tool:

1. Utilizing the shared communication space. The chat client had already been accepted by the collaboration, and scientists were accustomed to the level of interruption produced by chat messages. Since chat was the primary means of communication during telescope operation, it was presumed that scientists would pay attention to chat messages, yet would not find them annoying or distracting. The benefits of incorporating context and task information into the shared communication space were not fully understood at the time of design, but shortly became very clear. Bert only had to announce events in one place and everyone was able to see them. Not only that, every collaborator knew that the rest of the group had also been notified. This provided a shared context that was extremely valuable when making decisions under time pressure.

2. Logging. Previous studies have demonstrated the utility of chat logs in a number of domains [6, 11]. Instant message chat rooms can be logged with a date-stamp and the identification of the participants. System logs collected in the course of normal operations capture instrument and software

status, but instant messaging captures some essence of the shifter's state of mind. This was particularly important in designing the observing system because it revealed problematic or erroneous concepts that could be countered by improving the shifting interface. The date-stamp feature is extremely important for in-depth analysis of unusual incidents (equipment/software issues) because it indicates when shifters noticed excursions from normal procedures. Also, the reactions of shifters to such excursions can be correlated with the date-stamped system logs.

Logs are also useful for training, since shifters are able to review conversations of others on shift, increasing and improving the shifter operations "culture," and disseminating "best practices" for a successful shift.

3. Speech synthesis. Inexpensive or free speech synthesis plug-ins are available for many chat clients (notably *iChatter* for Mac OS X iChat, or the *festival* system for Linux clients). Some allow specific synthetic voices to be assigned to particular chat participants, adding a dimension of differentiation between "speakers." Shifting is an

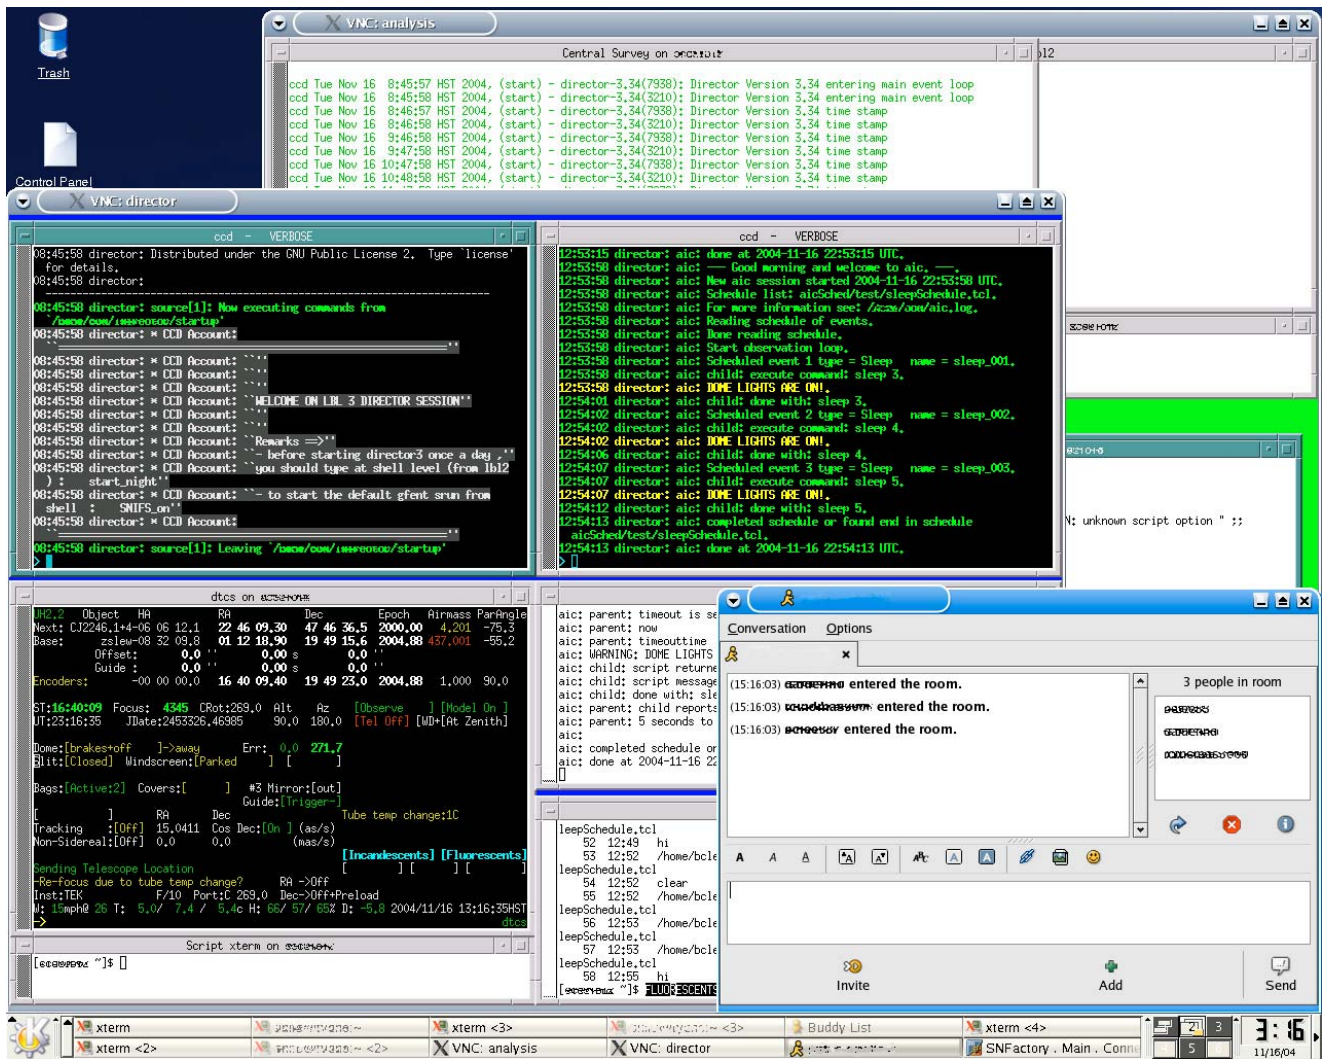


Figure 2. VNC telescope control window with chat

intensely visual task (Figure 2). The shifter must control windows in a VNC client, enter text on weather and instrument conditions into a data-taking tool, as well as examine acquisition images and spectral plots. With attention divided among such tasks, an additional visual stream for alerts is less than effective. It has long been known that spoken communication greatly increases the speed of completion of group tasks [4].

Since iChatter allows different voices to be assigned to each chat member, including a “robotic” one for Bert, the shifters could realize many of the benefits of voice-mediated group communication. The auditory stream, in the form of Bert speaking the alert, tremendously cut down the time to focus attention where it was needed. Response times of several seconds to even a minute or longer were cut to essentially zero, according to scientist feedback.

4. General awareness. The automated instrument control program (AIC) was the first program to capitalize on Bert’s presence in the chat room. AIC is an astronomical schedule-execution program implemented by the SNfactory collaboration. AIC processes observational “events” in the schedule, slewing the telescope to each target, running acquisition software, setting up for the exposure, running it, and so on. AIC uses a client script to send messages for Bert to relay to the chat room.

Each event is announced to the chat, along with its position in the schedule. This helps the shifter to keep track of where he or she is in the schedule. If AIC stops for some reason, this is reported to the chat. Bert announces other events, in particular time to sunrise and sunset, in the chat. During the night, time to sunrise is announced at intervals of one hour. In the last hour before sunrise this rate increases. Within 15 minutes of sunrise it is announced every minute. This practice induces a sense of urgency in the shifter to complete time-critical calibration observations with the impending approach of sunrise.

5. Reminders. Certain operations need to be performed by shifters at certain times. In particular, a few setup operations for a night require some level of coordination with a telescope operator in Hawaii. Bert reminds the shifter when to perform these interactions (telephone calls). The program also can remind the shifter to check focus, or to stand by to examine target images and approve or veto further followup.

6. Knowledge management. The summit (telescope control) computers are equipped with an SQL database that keeps track of target coordinates, associated finding charts, and the list of exposures made by the telescope. There are a large number of targets, including active supernovae and standard stars, which can be conveniently queried by asking Bert in the chat. Also, Bert can trigger an SMS (Short Message Service, or text message protocol) message to alert an “on-call” expert, who can then come to the VNC client and chat to resolve problems. This greatly reduces the amount of information a shifter needs to carry around in his or her head, or keep at the ready on a webpage or notebook.

Implementation

Bert was written in Perl using a variety of modules available from the Comprehensive Perl Archive Network (CPAN), and communicates with the chat network via the OSCAR protocol employed by the AOL Instant Messenger system.

At start-up, Bert connects to the chat network and finds a preset chat room. The program forks a “listener” process that accepts connections from processes on the summit machine and forwards those messages to Bert for the chat. There are two primary interaction scenarios with Bert. A message may be sent by the system (other programs) through Bert to the chat room for shifters to act upon, or shifters may query Bert for information.

Programs send messages through Bert via a simple UNIX file socket. In particular,

- AIC announces when the shifter needs to focus on an incoming acquisition image to possibly veto followup,
- AIC announces when a schedule stops because it ended or encountered an error,
- ultrafocus (an automated telescope focusing program) reports the focus results to the chat.

By extension, the UNIX “at” command can be used to trigger a Bert message at a pre-specified time. Some of such triggered messages prompt the shifter

- to assess telescope focus one hour after the previous focus was done,
- to call the summit to ensure that the dome is empty (of maintenance personnel, tours, etc.) before starting pre-sunset calibrations, and
- to contact the telescope operator 1/2 hour before sunset to set up.

A shifter communicates with Bert by including its name in a message. If such a message is detected, it is parsed to determine an appropriate response. The collaboration considered implementing a more elaborate grammar, but found that a primitive syntax also simplified the interface for the user. Supported functionality includes database-related queries, simple astronomical computations, and telescope status reporting. Some examples:

“bert, what is humidity?”

“bert, where is the moon?”

“bert, where is the telescope?” (telescope pointing coordinates)

“bert, where is (TARGET)” (returns target altitude/azimuth, airmass and compass direction)

“bert, NAME info” (gives contact information for person NAME)

The most common question asked is “bert, when is sunrise.” Another is “bert, what standard at (airmass)” which suggests possible standard stars to observe near that airmass. When needed, Bert can trigger SMS messages to members of a group of “on-call” observing experts available to help resolve incidents, and can send short messages from the sender describing the need for assistance.

Bert’s interface was built through a process of cross-cultural user-centered design. For example, originally the French shifters acted offended when Bert did not respond when they

greeted him or when they entered the chat, so Bert's interface was modified to respond and to greet each person by name as they entered. Other modifications in Bert's interface were made over time, based on scientist feedback.

Bert also has a lighter side; the CPAN Eliza "psychiatrist" module can be triggered by typing, "Hi, Bert." [14, 20]

USER EVALUATION

Interviews were conducted with four of the current shifters to evaluate the utility of Bert during a shift. Three of these scientists have been shifting since before Bert's integration.

A major motivation for creating Bert was the need to provide the scientists with a way to keep track of the time left until sunrise. The scientists all felt that the periodic announcements of and the ability to query sunrise times are Bert's most useful functions.

One scientist remarked, "to be honest, the end of night / sunrise is not a stress at all today as it was in the past." Additionally, these features now serve as the primary methods for keeping track of this information. The scientists found the information on 12-degree twilight particularly useful, because all science data must be taken before then to be usable. The scientists often use this function to determine whether there is sufficient time to observe an event.

Another scientist pointed out that these announcements give the shifter the sense that time is passing by and that the shifter should focus on data taking rather than problem solving. He recalls that, prior to Bert, it was not uncommon for the chat participants to lose track of time when trouble shooting a particular object.

The scientists also felt that Bert's AIC announcements play a crucial role in helping them focus throughout the night. One scientist explained that the hours for a shift are very long, and the shifter is typically performing multiple tasks (viewing images, talking to someone else, etc.). The Bert AIC announcements signal to the shifters that certain processes are complete and need attention. One scientist stated that this functionality improves reaction time and is "nice support for the long hours of observation." For those using the speech synthesis capabilities of iChatter, Bert announcements allowed them to work on another window without losing important events. Additionally, scientists commented that the announcements serve as a way to keep track of where they are in the schedule, functioning as a "secretary for the shifter."

Bert's mechanism for knowledge management is a time saver for the scientists, and all the scientists regularly use this function. One scientist added that the querying capabilities are particularly useful when trouble shooting, explaining, "you don't want to recompute your ephemerides by hand!"

One scientist noted that Bert doesn't overload the chat with announcements and provides the right amount of basic information to orient shifters on the state of the night. Another scientist confirmed this, stating that Bert "doesn't

talk too much." However, a few comments were made that, at certain points in the night, the periodic time announcements were not useful and usually ignored.

Each scientist personified Bert at some point during the interviews, and sometimes even insulted Bert in a joking fashion. For example, one scientist referred to Bert as a "helpful little guy," while another joked that Bert doesn't make for great company for a single shifter. He explained that the very fact that shifters "love insulting Bert" is proof that it's "worth having."

LOG ANALYSIS

The Bert chat archive spans over three years of logs, consisting of a total of 449,684 comments. This includes Bert, users, and system notices (e.g. as a new user joins the chat). The total number of Bert comments was 145,667, or about 32%.

Methodology

Feedback from the interviews indicated that Bert was able to solve some of the problems it was designed to address and others that had not been recognized. In analyzing the logs, we were not only interested in finding evidence to further support the interview findings, but also to discover if Bert was being used to augment chat communications and collaborative work.

Because a large groupware effort was integrated into the nightly data taking starting April 2007, we decided to focus our analysis on the April 2007 – December 2007 logs, as we believed these logs would best represent the current state of operations. These logs included a total of 150,644 comments, of which about 43%, or 64,260, were Bert's. There were 21 unique users logged during this period. We manually read through the first 3 months' worth of logs to get a sense of usage patterns over a continuous period of time. We then searched through our 2007 subset for specific uses of Bert (the "where is" and "what is" type questions).

Although chat logs from pre-Bert integration would likely have yielded additional insight into understanding how Bert fit into broader communication patterns, very few logs of this type existed. (As previously stated, the lack of such logs was a motivation for Bert's introduction.) The logs that were available, however, proved valuable in exemplifying some of the confusions and difficulties Bert was designed to address.

Scientist feedback indicated that there were asynchronous (in addition to synchronous) uses of the chat, for monitoring and training purposes. In order to understand some of these uses, we searched through the email archives of the group mailing list for mentions of chat and referenced back to the original chat archive. We focused our study on instances that occurred during our subset time span.

For readability of the logs, messages generated by Bert are reproduced here in italics (although in the operational chat no special typography was used). We occasionally elide text where it does not affect the overall sense of the

messages, or insert explanatory notes in square brackets. Spelling and grammar mistakes are left in the text as originally written. All names are pseudonyms to protect privacy. Informal followup interviews with chat participants were conducted to verify our findings.

Pre-Bert Chat Logs

Prior to Bert's deployment, the chat was used heavily to help with coordination and communication among the observers. Common categories of discussion include negotiating the use of shared resources, real time trouble shooting, and ascertaining the current state of operations. We were interested in finding samples of the types of problems that the scientists expressed existed prior to the integration of Bert.

Prior to Bert, there was no centralized resource for determining the current time and time until sunrise, so each participant would have to individually track such information. One shifter mentioned he used an external website to look up sunrise information, but before he learned about it, had a difficult time determining the time of sunrise.

The extract below depicts a typical instance where one of the participants is confused about the current time in HST (Hawaii standard time).

Art: So sunrise is at 6:00 am. And now it's 5.
Art: We are on spectrum 2 of 4, about 1/2 way through, which is 2500 s left.
John: No, it's 4
Art: Eh? Hawaii is 2 h behind...
Art: It's 7 here
John: read director ...
Art: Wow I am really just crazy.
John: yes ...
Art: I have no sense of time or distance.

During another night, technical problems prompted a collaborative trouble shooting effort among chat participants. The participant below is trying to coordinate with the shifter to solve the problem and is trying to determine whether AIC has been stopped or not. Since AIC often would end silently, this information could only be passed through direct communication from the shifter.

Bob: Marcel, ready with ctrl-c
Bob: ?
Bob: in the aic window?
Bob: or did you do that already?

Synchronous Communication using Bert

In reading the logs, we found a great deal of evidence for the use of Bert to gain the general awareness and real-time feedback needed to complete the schedule. At times, particularly days when no weather or technical problems occurred during the night, there was very little interaction among the chat participants, although the shifter would typically still interact with Bert, asking common questions such as "when sunrise" or "what standard." One shifter stated that, on a smooth night, getting through the schedule is usually fairly easy and sometimes even boring. However,

he noted that in practice, it is rare to experience a night without instrument or weather problems.

Examining the logs of nights with problems, patterns where Bert is part of broader group communications start to emerge. The following are the common categories of conversations involving Bert and two or more chat participants.

1. Collective Decision Making. Bert often plays a crucial role in certain types of time-critical decision making. Bert serves as a central knowledge management system, enabling scientists to determine information such as how much time is left until sunrise or to give options for which standard to observe.

Here, a participant asks Bert how much time left until sunrise towards the end of the night, and chat members collectively decide there is enough time to run another event.

Stef: Rene, did the sky spectra of the other day been useful?
Rene: I haven't looked at it yet Stef
Stef: ok
Rene: are you wondering if it would be useful to take an other one?
Stef: yes
Rene: I would say yes if we have time for it and if it is before 18 deg twilight
Stef: bert when is sunrise
BERT: 18deg twilight is at 2007-04-13 14:51:00 UTC (in 30 min).
BERT: 12deg twilight is at 2007-04-13 15:17:00 UTC (in 56 min).
BERT: Sunrise is at 2007-04-13 15:57:00 UTC (in 96 min).
Stef: yep
Rene: good
Rene: lets' go for it then ☺

This example shows how Bert is used to determine the next standard star to observe. (Note: "std" = "standard" and "fchart" = "finding chart," where standard stars with known brightness and spectrum are observed to calibrate the spectrograph, and a finding chart is used to exactly locate the target coordinates of a faint object.)

Gary: I suggest we try with a std star with an fchart
Rob: bert what standard
BERT: There is HR3454, a 4.30 magnitude B3V star, at secz 1.046.
BERT: Or HD93521, a 7.04 magnitude O9Vp star, at secz 1.143.
Rob: heh
Rob: bert what standard at 1.15
BERT: There is HD93521, a 7.04 magnitude O9Vp star, at secz 1.143.
BERT: Or Feige34, a 11.18 magnitude DO star, at secz 1.165.
Rob: There's Feige
Gary: let's try Feige

Follow up interviews confirmed that these use cases are very common. Obtaining the information from Bert allows for quick access of information needed to make time-critical decisions.

2. Collective Trouble Shooting. Bert will make an announcement in the chat if a fatal error occurs while processing an event. These announcements are very clear signals to the chat participants that trouble shooting must be performed. In the below case, Bert announces a fatal error. Chi is the primary shifter, but Stef is more experienced, so he steps in to help.

BERT: aic: ERROR: command [point_object] returned an error.
BERT: aic: ERROR: event 18 terminating on fatal error at 2007-05-03 10:40:22 UTC
Chi: yes
Chi: Stef, can you log in?
Stef: yes coming in
Chi: seen Fatal error
Stef: ok I am in
Chi: don't match?
Stef: weird: there are stars
Stef: just try it again
Chi: OK
BERT: aic: New aic session started 2007-05-03 10:43:13 UTC with schedule file aicSched/sch.070503.tcl.
BERT: aic: Starting event 18 (SNF20070326-012, type Supernova)
Stef: should be far from moon
Stef: Dec is off by 4 degrees
Stef: 4 minutes
Chi: yes
Stef: that's a lot
Stef: I think telescope is stuck

Here, a participant in the chat works with the shifter, and together, they come to the conclusion that the telescope is stuck. We later see that this is actually an incorrect assessment.

During follow up interviews, scientists added that, when using iChatter, they would "hear" the notices of fatal errors and would use that cue to pay attention to the shift and help out.

3. Keeping Track of Events. Chat participants use Bert announcements to keep track of the status of various processes. There were several instances when chat participants would give a brief response to these announcements without engaging in further conversation, primarily serving as indicators that each participant is paying attention to the shift.

BERT: ultrafocus: 2007-04-13T07:06:07 UTC, focus is 4971, seeing is 0.83, temp was 0.833, fmin was 4893, fmax was 5093, roule ma poule!
Stef: nice
Cy: yes
Gary: good!

4. Analysis Discussions. Bert announcements often signal the completion of certain processes, many of which require analysis by chat participants. During such cases, these Bert announcements can lead to focused analysis discussions.

BERT: 6 hrs 0 min to 12deg twilight (at 15:15 UTC)
BERT: (sunrise at 15:55 UTC)
BERT: aic: Starting event 18 (SNF20070413-010, type Screen)
BERT: aic: Starting event 19 (Night standard star 2 of 3, type Custom)
BERT: stdstar_factory: Observing Feige67.
Rene: there is something there
Chi: yes
Rene: i don't see a host though
Chi: but not center
Chi: isolated

New supernova targets are evaluated one final time just before a first spectrum is taken. In 45 seconds, a shifter has the opportunity to bring the acquisition to a halt before committing to a spectrum. Bert announces to the shifter to stand by to make such a determination. In the following example, two

shifters are alerted to stand by and decide to interrupt the sequence (resulting in the "fatal error" where AIC stops).

BERT: Prepare for check_match target confirmation, be prepared with stop_script!
Chi: ok
Dom: don't worry, you get another window just like it in Analysis.
Chi: nothing in ...
Dom: bummer
Chi: i stop !
BERT: ERROR: command [do_object] returned an error to aic.
BERT: aic: ERROR: child: terminating on fatal error at 2007-07-27 10:41:22 UTC

5. Humor. During problem-free nights, a shift can seem somewhat boring to the shifter and chat participants. AIC can process events mostly automatically, and each event can take some time to process. Chat participants have commented that they can usually tell if the night is running smoothly based on the volume of non-work/"humorous" conversations. One scientist notes that these types of conversations are an important relationship building practice. Bert has served as a tool to mediate these types of conversations.

Chat participants have invoked Bert's Eliza mode in order to amuse other chat members. One shifter remarked, "I always enjoy introducing new observers to his Rogerian therapy side." The sequence below was extracted from a longer conversation between two participants that contained technical exchanges, but also plenty of casual banter.

Dave: hi bert, what do you think of pascal
BERT: You're not really talking about me -- are you? (If you don't like my questions, just say "stop it.")

Below is another example where Bert is involved in a humorous conversation. In this case, one participant (Stef) is "speaking" for Bert by feeding him lines to output in the chat.

Rob: bert what is FLIGHTS
BERT: [Fluorescents]
Rob: ok
Stef: fluorescents are on
Rob: I believe you then
Stef: (I have snifscam on)
Stef: man of little faith
Rob: I trust bert more than you
Chi: :)
Rob: I programmed bert, not you
BERT: Stef is more trustfull though
Rob: bert, you got a weird french accent
Rob: is more trustful?
BERT: I am only a creation of a mentally deranged person
Rob: interesting

Asynchronous Uses of Bert

The usability interviews revealed that Bert's archiving feature led to asynchronous communication practices by the group. Although these techniques were not new to the group, they were not common practice due to the fact that no standardized archiving system was in place. As one shifter noted, "you had to either hope you left your chat client open over night, or that someone did." To add to this, the chat alone did not usually provide enough context to analyze the logs.

Bert not only provides a standardized archiving mechanism, but its notifications provide the context necessary to gather state information and understand the shifter's frame of mind.

Orientation

Scientists commonly join and leave the chat several times in a night. These leaves can span minutes, such as a coffee break, to hours, as is often the case for U.S. scientists monitoring the shift.

When shifters experience technical problems, they may call on an expert to help out. The scientists have heavily used the Bert logs in order to orient themselves to the current state of operations. By reviewing the history of the night, the scientists can get a quick summary to determine if the schedule is running, what was the last event, did that event crash with an error, and how much time is left.

This summary gives them the context necessary to know how to proceed with the trouble shooting effort. When there are two or more participants in the logs (a single shifter tends to be "quiet"), the incoming scientist can become more informed about the frame of mind of the shifters.

Other participants, such as the scientists not on shift, telescope operators, or even software developers, may have a low information requirement to orient themselves. Because a standard setup for observing involves logging into several VNC's, each loaded with information, an incoming participant will often skip the full setup entirely and mainly use the Bert logs to orient themselves.

Training/Increasing Group Knowledge

Prior to Bert logs, it was not common practice to review the logs of the previous night's shift. Even when chat archives were available for a night, it was difficult to link the chat comments with outputs from AIC and notes from the shifter.

Bert announcements of basic AIC outputs interleaved with shifter chat comments not only provide the context necessary to get a summary of a given moment in the night, but also furnish a linkage for investigating AIC's verbose output.

By evaluating the previous night's log, experts are able to pinpoint areas where deep misunderstandings exist in the minds of the shifters and chat participants. In the previous interaction sample in the "Collective Trouble Shooting" section (where it was concluded that the telescope was stuck), one of the lead scientists, in reviewing the logs, discovered that the chat participants had incorrectly assessed the situation. In an email the next day, the scientist discussed why the telescope being off by 4 arcmin did not indicate a problem, and the proper checks that should have been performed in that case.

The scientist, in a follow-up interview, describes this as a basic problem, where the "experts did not at first realize the steep learning curve that non-experts and beginner shifters had to contend with." Further, experts located in the U.S. could not monitor the chat all night, and non-expert shifters often could not or did not articulate such problems over

email. Therefore, the chat logs were often the only way to "unmask the training problems."

These types of analysis led not only to increased discussion over email of better practices, but also improvements in training documents. Analysis of the chat log also informed design and development of groupware, to eliminate situations where shifters did not react well.

DISCUSSION

Bert announcements within the shared communication space enable group awareness of task information, such as AIC notifications, and context information, for instance the time until sunrise. We refer to this as context-linked information. The case study we present exemplifies how context-linked information aids in communication and collaborative work.

Context information within the shared communication space provides the scientists with continuous environmental data. This data gives important feedback that informs a group understanding of time and how "on schedule" the night is, which is crucial in coordinating time-critical tasks.

Dourish and Bellotti [5] introduce the idea of shared feedback, where information on an individual's activities is passively collected and presented within the shared workspace, allowing individuals to organize collaborative work without the effort of explicit exchanges of task updates. Similarly, Bert announcements present task information about an individual's work within a shared space, providing awareness information to all chat participants. One benefit of this type of group awareness is a reduction of coordination efforts to establish common ground during group decision-making and trouble shooting. In other words, less time is spent trying to establish the state of the tasks, which is particularly important when working under time pressure.

Bert's notification mechanism differs from shared feedback in that the context-linked information is presented within a persistent, shared communication space, not just within the shared workspace. Our case study shows that this enables rich, asynchronous analysis of the shared work. Chat communication lets a scientist reviewing the Bert log see the shifters' frame of mind, while Bert's notifications allow the scientist to easily link these communications with actual events and work products. In addition, context information about the environmental state provides useful information on the potential urgency of tasks, which again, is directly linked to shifters' communications. In our case study, this high level of context to the archived work and communication provided the information necessary to assess problems in current work practices and led to improved training efforts.

CONCLUSIONS AND FURTHER WORK

Our case study demonstrated that context-linked tools have the potential to facilitate common ground in distributed teams, leading to increased team effectiveness under time pressure. Further, software tools to aid collaborative work in time-critical domains should integrate seamlessly with existing collaboration processes and software. Effective

software is likely to be lightweight and perform only functions necessary to the coordination of the work. Collaboration scientists accepted Bert because it provided several critical functions – knowledge management, key event notification, and enhancement of general awareness – yet did not require elaborate or time-consuming setup. The voice synthesis interface was a convenient side channel for providing event notifications and other information, and avoided overloading the user’s visual bandwidth during the demanding visual task of telescope operation. In short, as a virtual assistant in the chat during shifts, Bert was helpful and not annoying.

User evaluation and chat log analysis revealed a number of ways that Bert facilitated scientist coordination and communication during telescope operation. The case study also confirmed the utility of persistent chat logs for newcomer orientation, training, and increasing group knowledge. Additionally, Bert’s event notifications provided helpful context for these modes of asynchronous communication.

Bert’s interface could be improved if we could detect the degree of a shifter’s cognitive overload and interruptibility, and modify Bert’s verbosity accordingly. Additionally, we would like to perform a more in-depth analysis of the nearly half a million lines of chat logs. It would be interesting to adapt Bert for use in another, similar environment, and perform further studies. Finally, we hope to conduct controlled studies on the effect of providing context-linked information to distributed teams.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers, whose insightful comments significantly improved this paper, and the scientists of the SNfactory collaboration for their time and detailed feedback. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain. This work was supported in part by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and by the Director, Office of Science, Office of High Energy Physics, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and by a grant from the Gordon & Betty Moore Foundation. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

1. Aldering, G., Adam, G., Antilogus, P., Astier, P., Bacon, R. et al., Overview of the Nearby Supernova Factory. *Proceedings of the SPIE* (2002), 61-72.
2. Aragon, C. and Hearst, M., Improving Aviation Safety with Information Visualization: A Flight Simulation Study. *CHI 2005: ACM Conference on Human Factors in Computing Systems*, Portland, OR, 2005.
3. Bimholtz, J., Finholt, T.A., Horn, D.B. and Bae, S., Grounding needs: achieving common ground via lightweight chat in large,

- distributed, ad-hoc groups. *Proceedings of SIGCHI Conf. on Human Factors in Computing Systems, CHI 05* (2005), 21-30.
4. Chapanis, A., Johns Hopkins Univ. report, Interactive Human Communication: Some Lessons Learned from Laboratory Experiments (1976).
5. Dourish, P. and Bellotti, V., Awareness and Coordination in Shared Workspaces. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work CSCW92*, Toronto, Ontario, 1992, 107-114.
6. Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M., Richards, J.T. and Bradner, E., Socially translucent conversations: social proxies, persistent conversation, and the design of "Babble." *Proceedings of CHI 1999* (1999), 72-79.
7. Finholt, T.A., Evaluation of electronic work: research on laboratories at the University of Michigan. *SIGOIS Bull.*, 16 (2), (1995), 49-51.
8. Fitzpatrick, G., Marshall, P. and Phillips, A., CVS integration with notification and chat: lightweight software team collaboration. *Proceedings of CSCW 2006*, (2006), 49-58.
9. Fono, D. and Baecker, R., Structuring and Supporting Persistent Chat Conversations. *Proceedings of CSCW 2006*, (2006), 455-458.
10. Geyer, W., Witt, A., Wilcox, E., Muller, M., Kerr, B., Brownholtz, B. and Millen, D., Chat spaces. *Proceedings of DIS 2004*, (2004), 333-336.
11. Halverson, C.A., The value of persistence: a study of the creation, ordering and use of conversation archives by a knowledge worker. *Proceedings of HICSS 2004*, (2004), 108-117.
12. Lantz, B., SNIFS: a wideband integral field spectrograph with microlens arrays. *Proceedings of the SPIE*, (2004), 146.
13. McCann, R.S. and Foyle, D.C., Scene-linked symbology to improve situation awareness. *Aerospace Medical Panel Conference on Situation Awareness*, 1995, 16:11 - 16:11.
14. Nolan, J., A clone of the classic Eliza program, 2003, <http://search.cpan.org/~jnolan/Chatbot-Eliza-1.04/Chatbot/Eliza.pm>.
15. Olson, G.M., Atkins, D., Clauer, R., Weymouth, T., Prakash, A., Finholt, T.A., Jahanian, F. and Rasmussen, C. Technology to support distributed team science: The first phase of the Upper Atmospheric Research Collaboratory (UARC). in Olson, G.M., Malone, T. and Smith, T. eds. *Coordination theory and collaboration technology*, Lawrence Erlbaum Associates, Hillsdale, NJ, 2001.
16. Perlmutter, S., Aldering, G., Goldhaber, G. et al., Measurements of Omega and Lambda from 42 High-Redshift Supernovae. *Astrophysical Journal*, 1999 (517), (1999), 565-586.
17. Ribak, A., Jacovi, M. and Soroka, V., "Ask before you search": peer support and community building with ReachOut. *Proceedings of CSCW 2002*, (2002), 126-135.
18. Riess, A.G., Filippenko, A.V. et al., Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *Astrophysical Journal*, 1998 (116), (1998), 1009-1038.
19. Smith, M., Cadiz, J.J. and Burkhalter, B., Conversation trees and threaded chats. *Proceedings of CSCW 2000*, (2000), 97-105.
20. Weizenbaum, J., ELIZA - A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 1966 (9), (1966), 36-45.
21. Wulf, W. The collaboratory opportunity. *Science* (261), 1993, 854-855.