

Collaboration Tools for the Global Accelerator Network Workshop Report

Deborah Agarwal¹
Lawrence Berkeley
Laboratory

Gary Olson, Judy Olson²
University of Michigan

Introduction

The concept of a “Global Accelerator Network” (GAN) has been put forward as a means for inter-regional collaboration in the operation of internationally constructed and operated frontier accelerator facilities. An initial workshop to explore the topic of building a GAN was held in March 2002 at Cornell. This first meeting looked primarily at the accelerator issues involved with building a GAN. On September 17-20, 2002, there will be a follow-up to the Cornell Workshop at Shelter Island, New York. The Shelter Island Workshop will focus on possible experiments in remote control that are now being widely discussed.

During the time between the Cornell and Shelter Island workshops, a workshop was held to allow representatives of the accelerator community and of the collaborative development community to meet and discuss collaboration tools for the GAN environment. This workshop, called the Collaboration Tools for the Global Accelerator Network (GAN) Workshop, was held on August 26, 2002 at Lawrence Berkeley National Laboratory. The goal was to provide input about collaboration tools in general for the Shelter Island Workshop and to provide a strawman for the GAN collaborative tools environment.

The participants at the workshop represented accelerator physicists, high-energy physicists, operations, technology tool developers, and social scientists that study scientific collaboration (attendance list in appendix). The day began with an introduction, followed by an overview of the user-centered design process. We then had presentations by various experts in accelerator physics, control, and operations. They described some of their work practices and possible collaboration capabilities needed to support it. These talks were followed by presentations describing collaboration technologies and related issues.

The afternoon was spent engaged in the user-centered design process. We chose a few representative scenarios to study the communication flow and coordination activities, followed by a discussion of possible collaboration technologies to support this work. For this session, we drew flow diagrams showing locations of participants, annotated by possible technologies. We then had a general discussion about cultural issues, training

¹ This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Sciences Division, U.S. Department of Energy under Contract No. DE-AC03-76SF00098. LBNL Report Number LBNL-53253.

² This work was supported in part by grant IIS-0085951 from NSF to the University of Michigan.

and adoption. The presentations given at the workshop are posted on the meeting web site at <http://www-itg.lbl.gov/Collaboratories/GANMtg>. This report serves as an overview of the results of the workshop.

Collaboration Technologies

Collaboration technologies can be designed to facilitate interactions in several dimensions. One dimension involves interpersonal communication, access to stored information or access to the real world, including real time monitoring of instruments. The second dimension involves the time place at which work takes place, with researchers working at the same or different times and in the same or different places. Work often involves a mixture of these dimensions.

There are several tools that have been designed to support conversations. Text-based chats like those in Instant Messaging (IM), Multi-User Dungeons (MUD), and Internet Relay Chats (IRC) provide mostly synchronous interaction but can also allow participants to interact intermittently. Videoconferencing tools including among others Polycom Video Conferencing, NetMeeting, Virtual Rooms Videoconferencing System (VRVS), and Access Grids are designed to support face-to-face meetings at a distance. These tools provide varying degrees of visual and audio interaction capabilities and immersion. It turns out that a variety of subtle technical factors can influence the conversation, such as how far away people look (because of the monitor placement and the zoom of the camera at the far end, and how tall or short they look). The Access Grid developed by ANL provides a very immersive experience with a large video wall to project all the participants, several camera views of participants, and naturalistic audio. Technologies like NetMeeting from Microsoft, Via Video from Polycom, and the VRVS system from Cal Tech support workstation-based and somewhat limited interaction meetings, with video and audio, including meetings over high and low bandwidth.

There are several technologies that can be used to enhance synchronous interactions including voting mechanisms, ways to get brainstormed ideas out, etc. It is also important for the participants to be able to talk about a shared display, either a data stream or a report. A number of technologies allow this: screen sharing, use of electronic whiteboards, presentation software, and remote control panels. Ideas currently under development include workstation docking and peer-to-peer file sharing capabilities, which allow users to share data on any computer or PDA in a relatively ad hoc fashion.

To support asynchronous interaction, technologies such as email with multimedia attachments, conversation databases, newsgroups, etc. help to organize the stream of incoming information into topics, and to include the objects under discussion. Conversation databases add the important property of being persistent, allowing people to return to previous discussions and to review or discuss earlier topics. Data can be archived, and annotated for continuing discussion. Early sessions can also be captured and replayed.

Electronic notebooks allow access to an organized stream of activity, comments, and data. There are several electronic notebooks that have been developed and are in use in various communities. Many accelerators already use some form of electronic logging of information in routine operation. Recent development efforts in electronic notebooks have concentrated on standardizing the data structures, auditing, and signature capabilities. An example of such a notebook is the CORE 2000 electronic notebook from PNNL. Authoring environments allow the creating, commenting, and editing of joint reports in an easy manner.

Some technologies support and span the transition between synchronous and asynchronous work. Calendars allow access to a person's whereabouts, as do pagers and cell phones. Instant messaging systems like MSN IM often display whether someone is working at their computer or not, indicating an aspect of their potential availability. Project management and workflow software allow people to schedule upcoming experiments and simulations, track progress, and coordinate activities. It also allows people to calculate the implication of a delay in some stage of a multi-stage process. The cutting edge technologies are focusing on the awareness that people have of others' work, giving them a visual field that shows the activity (over the shoulder, a fly on the wall) at a remote site. Handhelds and GIS can support people's accessibility globally. Cutting edge workflow systems are also incorporating Grid technologies to provide security and allow submission of compute jobs on Grid enabled machines.

In order for collaboration technologies to be acceptable to end users, we must consider a number of social issues. Security, privacy, interruptions, training etc. are all important factors that must be considered. People have to be able to say who they are in an authentic way, yet be able to consider their interactions private, accessible only to the their colleagues. This is a significant challenge when activities are conducted over the Internet. If people are being seen, they want to see who is seeing them, a concept of reciprocity. Whatever technology is employed will have to be easy to use, and independent of the hardware platform of choice at the site. Applications should be customizable to fit the particular situation. And people have to talk through the "rules of the road," the habits, procedures, and work practices that surround the effective use of the technology.

General GAN observations

This section contains a summary of some of the requirements and observations pertinent to a GAN scenario. It provides an example of some of the operational procedures that might be in place at a typical accelerator, some of the social hurdles that will be encountered, and some general thoughts about the approach that could be taken in building a GAN.

Operations

Focusing on operations per se, there are a number of communication tasks involved:

- The reading of and making entries of detailed technical information in the shift log
- Shift change meeting with the outgoing crew chief talking to the incoming crew. Often this is in a separate room (not the control room) with a lot of Q&A. It also provides management contact with non-standard shift crew, 15 minute with the night crew.
- Daily status update meetings. These meetings provide department heads, maintenance crews, and repair crews with an opportunity to talk about preventative maintenance support and general staff. Weather takes out power grids, so you have to know the weather at remote sites.
- Maintenance coordination. These are meetings to review recent problems and to conduct medium and long term planning with outside groups. They help to assess program impact if the work is authorized (e.g. if a power supply is down, is there other work that can be done).
- Experimenter liaison meetings. Here they look at the background conditions, what just happened, and how long it will be until they have beam.

A lot of these encounters involve building working relationships, and learning to trust another's judgment.

There are also a number of safety critical tasks:

- Control access to the accelerator enclosures and radiation areas. They have to search and secure tunnels after general access. There should be video monitors of doors for controlled access and restricted key control.
- There are periodic checks of safety system interlocks. There is a configuration of devices consistent with the requirements, such as beam stoppers, RF, HV, and physical checks of the devices after work is performed, to insure radiation shielding, etc.
- Monitor accelerator and beam conditions. They are responsible for maintaining that the beam power is within the safety envelope.
- Respond to emergency incidents. Fire alarms, atmosphere alarms, and injuries may impact the accelerator program.

There are accelerator tuning and monitoring phases, where they monitor performance, measure luminosity, energy, current, etc. They compare it with what is expected and change things like the data update rate. When there is a problem, the beam is dumped and there is a great deal of follow-on analysis of multiple signals to determine and make corrective action and then proceed.

There are also plans to do machine development studies. These studies are essentially experiments; they plan, and then alter the plan after trying things out. They take new measurements on the fly and have to determine if the problem is with the accelerator or the measurement device. Often they will devise a new measurement on the fly. In recovering, they have to restore all conditions before they can proceed with a new setup.

Training is a regular and no-going process, it is important to make sure that everyone maintains a working knowledge of the frequently changing control hardware and software. Everyone has to keep up with the changes in procedures and methods as well. There are many idiosyncrasies, such as having to make improvements to an older working system only when they limit performance because budgets are tight. They know that if they had the money, they would redesign many aspects of the current system. But the reality is that you have to learn the idiosyncrasies and move on.

Turnover in the operator role is modest, with most staying 3-4 years and others 20. After three years, people typically decide whether this is what they want to do or not. Training takes about a year. At many sites, there is no clear career path from this position.

Social factors

For the GAN remote operations to be successful, there are a number of social aspects that must be addressed. Because the collaborations are across national and cultural boundaries, the participants will have to relinquish a substantial portion of their natural, national working modes and seek a GAN working mode. Conscious effort will have to be spent on finding best practices and a way to adapt.

The GAN will be a joint venture from the start, as opposed to a particular location being awarded the instrument with others being granted secondary remote access.

Collaborative, remote communication will occur at all stages of development, including Design, Construction and Operation. Social issues of trust and motivation will become evident. We will also have to agree on how to approach the collaborative decision making.

Collaboration capabilities will create a large distributed workload, with people working on different components and developing the remote operation for their component. In order to develop these remote capabilities, there will need to be extra manpower, just to get over the threshold. The plan is to look for pieces that are separable, so that distributed development can proceed without a large amount of coordination. Tools will be developed, however, to help the coordination.

Approach

Collaboration tools are needed for all phases. For the Design phase, they are needed to ensure fit, documentation, etc. There are well-defined procedures for this, which will help in making a remote workflow system. The commissioning phase is perhaps the toughest to do remotely, because nothing works the way you expect it to, and you have to “get into a rhythm.” There is an amorphous point in the commissioning when it is declared complete, yet the anomalies and fixes are still going on, though at a slower rate. In the Operations phase, there are a lot of cycles of experimental settings and requisite changes.

At the beginning, the emphasis of the collaboration tools might be more on the social interaction, the longer term on the more technical. We might recommend starting with a simple scenario and then build from there. Start with simple collaborative tools, like

electronic notebooks, and then support talking around the electronic notebook, etc. During Operations, there are several places on the accelerator where things need to be communicated among and to people. It would be nice to broadcast webscreens, electronic notebooks, and voice annunciators and, we'd want to have remote observation of and participation in a shift-change meeting. We want to minimize the inefficiencies of second-hand information.

Prospective GAN scenarios

This section contains descriptions of a few prospective prototype GAN scenarios. These are scenarios that were put forward by the attendees of the workshop as possible examples of the types of activities that might be pursued in a prototype GAN deployment.

TTF2

During the next year the Tesla Test Facility at DESY in Hamburg Germany will be significantly upgraded in energy, to be called TTF2. TTF2 presents excellent opportunities for possible GAN prototypes. Cornell and DESY are both very interested in pursuing this idea. Cornell has a strong commitment to GAN because it will not be a candidate site for the eventual international Linear Collider (LC), they have an extensive background in accelerator physics and they already have experience with remote control and videoconferencing for the CLEO detector.

Development of the TTF2 data acquisition system ("DAQ") is a possible GAN prototype environment. The DAQ will be a distributed development project, based on experimental detector expertise from Ohio State, Cornell, and DESY. All three locations will be involved in code development. They plan to use a system to coordinate development of code at different locations, probably using CVS, and some mechanism to keep the documentation synchronized with the code management system. They envision also using a distributed whiteboard that allows sketches and diagrams to be seen and manipulated remotely, as well as a remote compilation and testing facility. This facility would first function locally, and then allow remote access as if on site.

A second GAN prototype scenario, requiring use of the DAQ, involving machine development studies, and emittance measurement is a likely subject of study. This is an important and difficult problem that must be solved for the LC. Team building will be critical for making this measurement effectively. It typically will require sufficient time on site, with many groups and individuals involved. This work has important security issues, with a high level of trust developed among the people and in the security of the communication mechanisms. Remote participants will use simulation extensively.

Machine Development, with emittance measurement as an example, involves:

1. Meeting with colleagues to agree on goals and plans of action.

2. Design of any new equipment required, with access to survey data, geometry database, and CAD files, because the instrumentation has to be inserted *in* the accelerator, e.g., inside a vacuum. Everything has to fit. Stringent safety requirements also have to be met.
3. Extensive software development
4. Instrumentation manufactured both at home and at the site.
5. Testing, first done at home, and then installation at the site.
6. On-site testing with the installation team followed by remote testing.
7. Operation of the accelerator, maintaining close communication with the on-site operator and the site shift responsible person.
8. Acquisition of data and transmission to the remote site or
9. Remote access to the on-site data store.
10. Performance of preliminary data scans before ending the session to assure validity.
11. Post-shift procedures and checklist.
12. Analysis of results (both early and over time)
13. Discussion of results with colleagues.

There are considerations of ongoing planning and coordination including the ability to resolve differences when they arise. Team building is essential, including a need for ongoing social interaction.

Note. Although the above scenario is specific to machine development studies, e.g. emittance measurement, it incorporates generic remote collaborative activities, like coordination of meetings, analysis of data with conversation between remote participants, etc. The collaboration technology support will likely generalize over a number of specific scenarios, all relevant to an eventual GAN for the LC.

As for tools for these scenarios, we can see a need for video, application sharing, information portals and distributed program development, including remote control of the DAQ. Security is an important issue since these discussions will be conducted remotely. Notes and data will be kept in an electronic notebook, which is to have easy access and easy entry of information including automatic recording of measurements. They will need access to detailed hardware monitoring and diagnostics.

RHIC/SNS. They are currently trying some remote operations at RHIC, giving control/access to outlying support buildings. This trial is 2-4 months away. CERN is involved with potential collaboration with BNL in 4-18 months, coordinating through RHIC instrumentation. SNS remote diagnosis and the commissioning of the SNS ring will happen in 1-3 years.

A mini scenario: Dumping the collider store. People need to know 10 minutes in advance that the dump is going to happen. It takes about 40 minutes, during which time everyone would like to keep updated on general status and current estimate of completion time. The experimenters want to know what's going on, and want beam ASAP. Nobody

has time to get on the phone to tell others what's going on. Status information should be collected "passively" and made available to the experimenters.

A second mini scenario: Beam study periods. In this, the participants need to compare and discuss screens or scopes. Currently, they place copies in an electronic logbook and discuss them over the phone. They would like to have continuous background voice connectivity with the remote locations.

During the beam study periods, they need ad hoc meetings, where access control is very important. They might have to rewrite code on the fly, with narrow-deep access at the site. It would be good to be able to share files, analysis codes and displays. Sometimes the data is taken and analyzed later, other times it is handed off to the night shifts to analyze. Often in this period, there is improvisation, with high levels of problem solving. These instrument tests often push the boundaries of the accelerators. Perhaps this coordination technology could be first tested with the Phaselock loop tuner commissioning.

A third mini scenario: SNS Startup and Ring Commissioning. In this, the experts are on call 24/7, and require remote access when requested. There is extensive collaboration, with broad/deep controls to access by BNL personnel. For example, the #1 power supply hardware examination/diagnosis involves coordination of onsite and offsite engineers. If at this point there could be "shoulder riding telepresence" and bi-directional audio, there is a great opportunity for training.

LHC. The US LHC Research Program (US LARP) involves commissioning, accelerator physics, beam instrumentation and upgraded Nb₃Sn IR quads R&D. The participating DOE labs are BNL, FNAL and LBNL. In the context of the US LARP, LHC machine development studies undertaken to improve the performance of the accelerator for physics experiments are a possible use of prototype GAN tools. Using GAN tools for remote participation in the LHC machine development studies and logging could help gain familiarity and acceptance of the full-blown GAN. The benefit for the US LARP would be greater participation with less travel and cost and presumably CERN would benefit from this as well. However, it is not likely that GAN tools could be used successfully for LHC or any other accelerator facility until after a team culture has been established. The establishment of the team culture would require that most if not all of the participants work on site at CERN for a significant period of time (six months to a year). Maintaining the team culture would also very likely require periodic on site visits several times a year even after GAN tools were up and running.

We can envision using GAN tools during LHC machine development studies and operation for physics experiments to run simulation codes and compare these with experimental results, for example for the electron cloud effect and beam-beam interactions. We will need flexibility in recording the suite of instrumentation needed for a particular machine design study.

As for collaboration technologies, we do not believe the entire control room needs to be duplicated, merely what's on the screens. They wish two-way communication. If it looks the same to everyone, it is easy for everyone to deal with. Video is needed in order for everyone to see everyone, and it should be accompanied with high quality audio. In addition to the control room, there should be a meeting room for development and discussion of run plans, analysis, information exchange at shift changes, etc. They also need the electronic logbook, whiteboards visible at remote locations, and computers for logging data streams, analyzing and displaying summary data, and running simulations. In addition, they want everyone to be able to plug in their own laptops and participate from their laptops.

User-centered Design

In building (or purchasing) new information technology, it is critical that the decision be based on an understanding of the users' needs. In the field of Human Computer Interaction, there are a collection of methods, called User-centered Design, that extract from the users their needs, based on the users' goals, work, and work setting. When the users' goals and work are understood, the functionality of the new system can be defined, built, and tested with end users. This collection of methods can be applied to the development of collaboration technologies and the associated work practices.

Among the User-Centered Design methods are interviews with end users, asking not only what they want, but also what they do. To do this, we collect simple narratives, draw flow charts, data flow diagrams, or swimming lane diagrams, such as those shown in Figure 1. During the workshop we used a group interview or focus group to draw detailed scenarios for use in determining what kind of collaboration technologies might be useful in the GAN context.

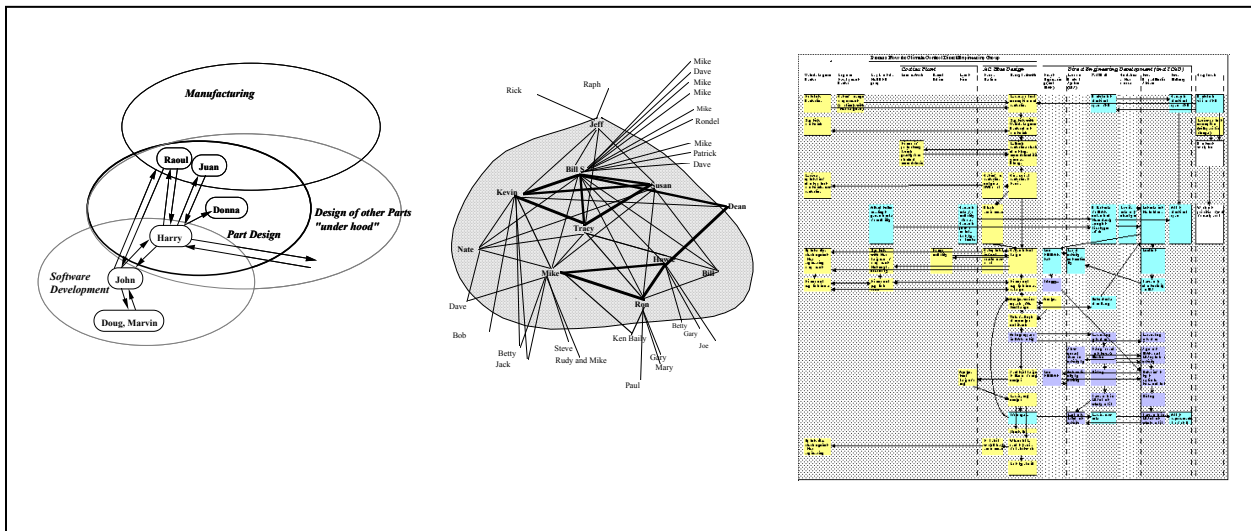


Figure 1. Various representations of the communication and coordination in a collaboration used in User-centered Design. The one on the left represents people and their major contacts in various departments/functions. The one in the middle shows

different levels of communication among people in and outside a department. The one on the right shows individuals and data sources in columns, with the communication flow in time going top down, similar to a flow chart.

Two detailed scenarios

We chose to focus on TTF2 as an example of both design and operational phase scenarios. TTF2 seemed appropriate as a focus and perhaps a pilot for these technologies since it has a high probability of happening, it has a 1 ½ year time frame, and about 3-5 institutions will be involved. Although steady production operation is not envisioned, many aspects of operation are included, such as shift changes and information exchange. Both Cornell and DESY are deeply committed to the idea of remote participation.

Data Acquisition System (DAQ). Figure 2 shows the locations and activities involved in designing, building and implementing a DAQ for TTF2. People at three locations (plus possible additional remote collaborating sites) are involved: DESY (Germany), Cornell (US), and Ohio State (US). The hardware for which the DAQ is designed resides at DESY in Germany. They are also responsible for associated software, API specifications, configuration management, change control, and documentation oversight. People at Cornell and OSU are in charge of the development of the DAQ software, emulation of hardware for early testing, and also deployment of tools for documentation of the development process. The end users of the system will be the accelerator physicists, all over the world.

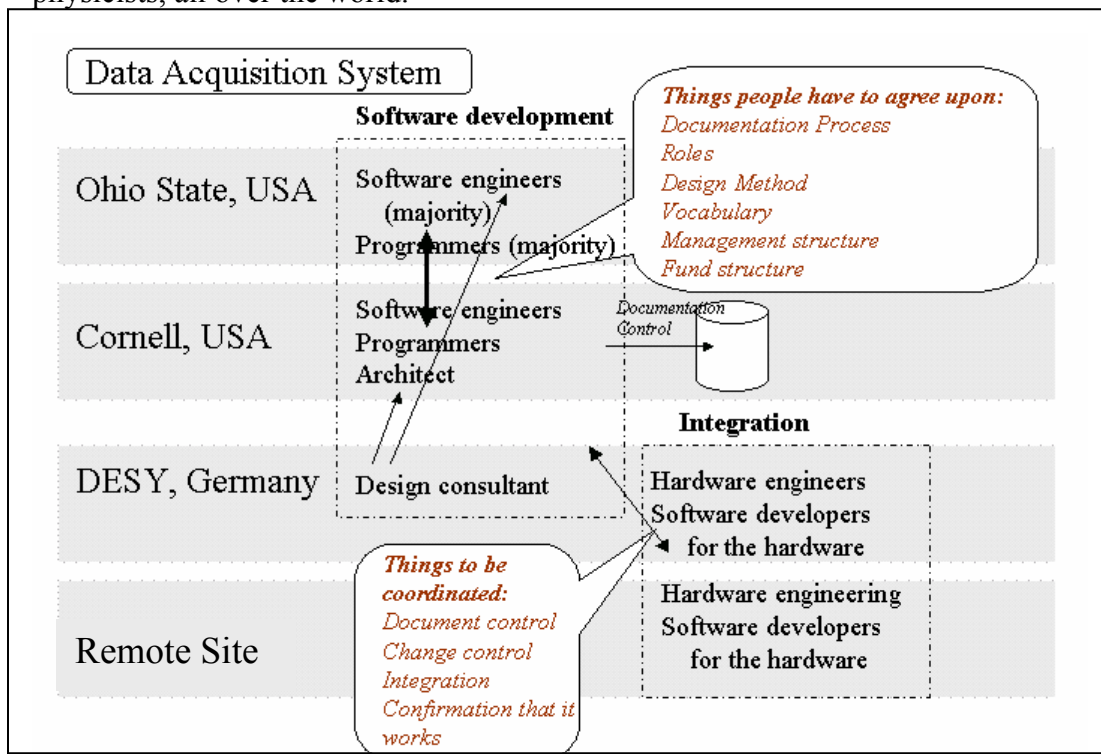


Figure 2. A representation of the locations of various roles in producing the DAQ, annotated with the kinds of issues the players need to coordinate.

They coordinate through a number of face-to-face meetings, especially at the beginning, and then continue with a mix of face-to-face and remote meetings. Early on they have to agree on various ways of managing their work: how documentation will be collected, who will play what role, what management structure, what design method to use, language definition, and funding structure. There will be local and remote testing, the remote being when they are integrating with the actual hardware. The documentation is expected to be web-based and evolve as the software is developed, requiring a high degree of coordination between Cornell, OSU, and DESY.

Machine development studies. Machine development studies involve the same participants as in the DAQ development, Saclay (France), and INFN (Italy), possibly others. Likely there will be a number of studies going on in parallel, making this activity a good prototype setting for GAN.

As shown in Figure 3, the machine operations are in DESY, with a number of remote accelerator researchers (ARs) coordinating on a study. One of the ARs takes the lead on a study, and typically many studies are in the works at the same time.

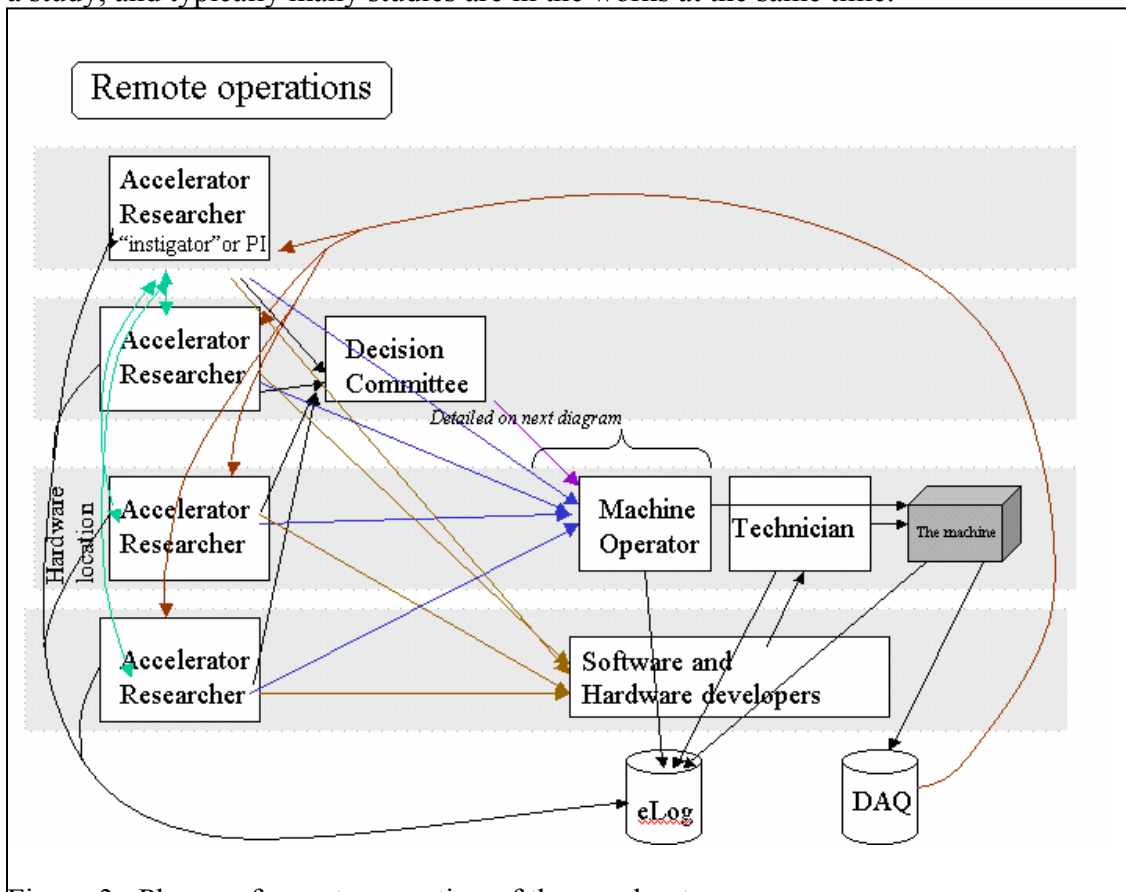


Figure 3. Phases of remote operation of the accelerator.

Requests for time are decided by one person, advised by a committee, whose membership is determined in proportion to their monetary contribution. In order to do the study, various hardware components need to be installed, which is a DESY responsibility.

To carry out the study, the accelerator is run for a number of shifts, which involve changes in personnel. The remote researcher is envisioned to be able to control various accelerator instruments and diagnostic tools. For safety reasons, the beam itself will be controlled by the operator on site (shown in Figure 4). During the study, various entries are put into an electronic logbook. Data collected from the study (from the DAQ) goes to the lead researcher and is shared with various other interested researchers. They collaboratively write up a paper, and store it on the document control system.

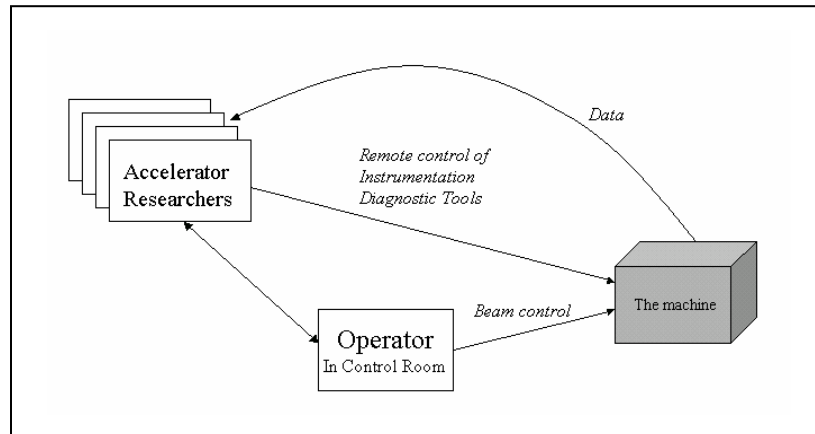


Figure 4. Kinds of remote control during operations.

Additional notes on machine operation and development. There is value in observing both routine operation and the shift change meetings. Observers learn by watching. Regular machine operation is very much like the machine development studies outlined above, with the exception of the time scale. In machine development, there is a lot of on-the-fly problem solving.

Collaboration/Coordination Technologies to support these scenarios.

The coordination that takes place in all these settings is both asynchronous and synchronous, involving both scheduled meetings and informal communication. In the following, we note the functionality that the technology should offer, not a specific product. Decisions about products, whether they are specially built, reused tools from others (e.g. the electronic notebook from EMSL), or off-the-shelf commercial products, will be determined at a later time.

Meeting support. Users want to be able to give presentations remotely with both audio and video of the presenter along with the PowerPoint slides or presentation of other things (like a data stream, simulation results, 3-D CAD tools, video, documents, visualization and the electronic logbook). All participants should be able to see the other participants, as well as the speaker so that they can capture their reactions. People should be able to connect from their offices as well as special video conferencing rooms, and should be able to participate when traveling. When connected at the desktop, it would be desirable to have the stream running in the background, allowing passive participation

and multi-tasking when appropriate. The system should have a control panel that makes it easy to set up a call. There are issues here of how to connect various technologies.

The session should be captured for replay by those who were not able to connect. It would also be good if the replay could be indexed for special events, making search for particular things easy. A scheduling system (like a shared calendar) would help coordinate people's participation and allow others to look-ahead to upcoming activity.

Informal meetings. On occasion, such as when an unexpected result happens, people will want to have immediate access to experts for consultation. For this, an electronic calendar system with appropriate rules about its use (a cultural issue) and an awareness system (for either on-line work or using badges and pagers and cell phones) will support finding and contacting the right person. Once contacted, the kinds of technology that support formal meetings would support these less formal sessions.

Remote operations. Initially, it might suffice to have remote *observation* of control room activities, rather than direct control. This might serve a "looking over the shoulder" service, for learning and discussion. Then the remote researcher might wish to control a subset of parameters, instrumentation or data collection parameters. For safety reasons, prototype remote operation will not include remote control of the beam.

Since the operator's hands and eyes are busy controlling and monitoring the instruments themselves, the primary communication will likely be auditory. For this wireless headsets are appropriate, with nearby video of the remote partner. The remote partner is expected to have video and a remote view of the control/data stream, but only the subset of parameters that are relevant to the study. The same kind of views/communication channels will work for installation, testing and acceptance for those who do not need to be on site.

Asynchronous communication. A primary repository for these projects will be an electronic notebook. The meetings that are captured will also be stored, with annotation capabilities to let people comment on what happened and for others to be able to view this. Email should be organized, like discussion databases, by threads or topics. Documents and data displays should be able to be annotated as well. There should be a common agreement about what goes into the electronic notebook; including agendas, open issues lists, etc. Voice mail is also important, especially for those with mobile phones. It might also be good to have a system that integrates voice and email so that there isn't an extra burden for the user to check various sources.

Related issues

Culture. We recognize that one of the overriding issues in making remote collaboration work is that people from different cultures will be working together. Since remote work does not include the normal social events that accompany travel, where people can begin to understand patterns that distinguish one culture from another, misunderstandings are likely to be more frequent. Clearly, some attention will have to be paid to both educate

people about each other's cultures and come to some explicit agreement if there are differences in expectations about work practice, decision making, and various criteria for prioritizing and making decisions.

Training. There is some concern that remote people will not be properly trained in the use of the remote operations interface. Possible remedies include being trained at the host site, so that training is done face-to-face and practice can take place under the eyes of the trainer/operator or using simulators. There is a legal requirement for certification, so this is an important issue.

Adoption. We recommend that the initial tool set be simple to use and a small step from the way people are working now. Also choose tools at the outset that have applicability in many different settings.

The host site is likely to wonder what the advantages are to providing remote access. Among the advantages are extended manpower and expertise, reduced need to host a person throughout their stay, automatic documentation of the activity, and greater "literacy" with tools that are common in many types of work. Early adoption of these tools will be important for LC. In the future of HEP, there will be limited numbers of sites, so remote access will be a way of life. It will be important to track the patterns of use by local and remote users and to have regular user group meetings.

Acknowledgements

We would like to thank the participants of the workshop (listed below) for all their help with the workshop and this report.

Appendix - List of Attendees

- [Deb Agarwal](#) - LBNL - Host
- [Paul Avery](#) - University of Florida
- [Nathan Bos](#) - University of Michigan
- [George Chin](#) - PNNL
- [Terry Disz](#) - ANL
- [Hans Frese](#) - DESY
- [Miguel Furman](#) - LBNL
- [Ray Helmke](#) - Cornell University (Wilson Lab)
- [Wolfgang Krechlok](#) - DESY
- [Stu Loken](#) - LBNL
- [Gary Olson](#) - University of Michigan
- [Judy Olson](#) - University of Michigan
- [Marcia Perry](#) - LBNL
- [Massimo Placidi](#) - CERN
- [Todd Satogata](#) - BNL
- [Mike Stanek](#) - SLAC

- [Mary Thompson](#) - LBNL
- [Bill Turner](#) - LBNL