# The Upper Atmospheric Research Collaboratory

**Gary M. Olson, Daniel E. Atkins, Robert Clauer, Thomas A. Finholt, Farnam Jahanian, Timothy L. Killeen, Atul Prakash, and Terry Weymouth**

Upper atmospheric physics focuses on the study of the earth's ionosphere, looking particularly at the interactions of the solar wind, the earth's magnetic field, and the characteristics of the upper atmosphere. Observations of these phenomena are made with ground-based instruments, satellites, and rockets. In recent years, a series of computational models of the entire upper atmospheric system has emerged, aided by progress in supercomputing.

Although observations of upper atmospheric phenomena are made at all latitudes, because of the characteristics of the earth's magnetic field, most ground-based instruments are concentrated at high latitudes, particularly in the Arctic. Many of these facilities are in remote areas and are relatively difficult to reach. With the ending of the Cold War in the early 1990s, inexpensive military flights to many of these remote areas ended, making access more difficult and more expensive.

Fortunately, these changing circumstances coincided with the emergence of the Internet. It occurred to a number of scientists in the field that network connections to these remote facilities could improve access and have a beneficial effect on the practice of science. Obviously network access to remote facilities would ameliorate the transportation difficulties, but in addition, it would provide greater access to such facilities for scientists and students at all kinds of institutions and would offer great flexibility for the scheduling of observations to coincide with scientifically important events. For instance, spontaneous coordinated scientific campaigns in response to events such as solar flares would be possible.

These problems and opportunities fit very nicely the vision of a collaboratory [1]. In 1992, a group of space scientists, computer scientists, and behavioral scientists at the University of Michigan obtained funding from the National Science Foundation to launch the Upper Atmospheric Research Collaboratory (UARC). UARC is a 6-year project to design, develop, deploy, and evaluate a testbed collaboratory.

Our initial goal was to provide a distributed community in this field with real-time access to remote instrumentation and to provide collaborative tools that would allow them to interact with each other over real-time data. As our user community gained more experience with early generations of UARC, their horizons widened. New ideas about how their science might be practiced led to the development of new collaboratory capabilities to support these visions. Thus, the technologies we have provided have evolved considerably as we coupled our system development with detailed observation of scientific practices that were resulting from the use of the technology. In this article, we trace the progress of the UARC project, reflecting in particular on the evolution of the practice of science through the use of Internet-enabled collaboratory technologies.

From the outset, the project has been itself highly collaborative. We believe that testbeds of this sort must be developed cooperatively by a multidisciplinary team consisting of domain specialists, computer scientists, and behavioral scientists. We have engaged in a user-centered, iterative rapid prototyping design strategy [4]. The project began in early 1993 with intensive investigation of the user community. We visited laboratories and instrument sites and collected extensive use cases from the scientists about the details of their practice. These were fed into the design of the technology itself. Because we used an object-oriented development environment (NeXTStep), we were able to build initial prototypes quite quickly. By April 1993, 4 months into the project, we had initial versions of the technology in the hands of the scientists. Study of the use of the technology in actual scientific practice led us to revise our design. We have continued this user-centered iterative strategy throughout the project.

UARC has provided us with a number of important opportunities. First, it has been a rich testbed for the exploration of the collaboratory concept itself. Second, it has given us an opportunity to develop and evaluate user-centered development methods. Third, it has allowed us to explore the technical challenges in building collaborative applications over the Internet. These challenges have played a major role in the progress of the project. Finally, UARC is a long-term, longitudinal study of the effects of collaboration technology on scientific practice.

## The UARC Project: A Brief History

The UARC project began by focusing on a community of users of the Sondrestrom Upper Atmospheric Research Facility in Kangerlussuaq, Greenland. This facility is located on the west coast of Greenland north of the Arctic Circle and is jointly operated by NSF and the Danish Meteorological Institute. The principal instrument at this site is an incoherent scatter radar (see Figure 1). Most of the scientists in our initial UARC testbed had been using this radar or participating in scientific campaigns in which the radar played a central role from the time it was established at that site in 1983. In addition, over the years, a number of other instruments had been located at the Sondrestrom facility. We worked with the principals responsible for a total of five instruments at the site: the incoherent scatter radar, a magnetometer (part of an array of magnetometers along the west coast of Greenland), an all-sky camera, an imaging riometer, and a Fabry-Perot interferometer. By the fall of 1994, all of these instruments were accessible over the Internet via UARC.

Figure 1. The incoherent scatter radar at the Sondrestrom Upper Atmospheric Research Facility in Kangerlussuaq, Greenland.

UARC began by focusing on the existing

practices of a subset of the community of Sondrestrom users. We focused on a particular element of their science, namely, the real-time collection of data from the set of instruments at the site. Because the radar is expensive to operate and requires an hour-long warm-up period, it is usually scheduled in advance. Coordinated observations involving several scientists and often multiple instruments are referred to as "campaigns." SRI International provides operational management of the Sondrestrom facility and coordinates the scheduling of observational periods.

The fact that the site is remote introduced a number of constraints on the practice of science before UARC: the cost of going to the remote site, the limited number of people who could be accommodated at the remote site, the lead times needed for scheduling the facilities, and the logistics and travel to the site. Interestingly, the community of UARC users was highly collaborative prior to the appearance of UARC technology. They would regularly schedule joint campaigns at Sondrestrom. They often combined data from multiple instruments into analyses that resulted in joint publications. They interacted by phone, fax, and e-mail and scheduled interactions at periodic scientific meetings, but the frequency of such activities was limited by the constraints of distance and time.

Thus, from the very beginning, the goal of UARC was not to create collaboration where there was none but to allow an existing collaborative community to work together more flexibly. Our users talked about being able to collaborate more on real-time data collection if it did not require extensive travel. They talked about being able to respond in an opportunistic manner to emerging events (e.g., solar flares that produce electrodynamic changes in the earth's upper atmosphere within the next day or two). They talked about being able to interleave different kinds of observational protocols depending on the current conditions above Sondrestrom. All of these would reduce the cost of collaborative observation, with a corresponding increase in the kind of collaborative analysis and publication that everyone agreed would advance the science.

Another perceived benefit of UARC at the beginning was the impact on junior scientists, graduate students, and possibly even undergraduates. Prior to UARC, real-time scientific campaigns were mostly the province of relatively established scientists at major research institutions. Junior scientists at smaller institutions and students had little opportunity to experience firsthand the interactive aspects of real-time data collection. We expected Internet-based access to have a big impact on this.

During the first 3 years of the UARC project, we developed a series of versions of a collaborative environment in the NeXTStep programming environment. We had selected this technology platform because it provided a rich set of tools for doing rapid prototyping in an object-oriented environment. Because the early user community was modest in size, we provided them with either NeXT or Intel platforms for running the NeXTStep software. The software during these early years of the project provided viewers with the ability to look at real-time displays of "quick look" data from the five instruments listed earlier here. Because the visualizations presented to users were calculated locally from these data, UARC participation was possible for users with a wide variety of network connections. Each user could select their own parameters for the display, though one collaborative feature we offered the users was the ability to share windows, including a shared telepointer [6]. We built a chat facility that allowed users to enter a text message and broadcast it to all other users who were logged on. Later we added a separate chat window for communication with the site crew in Greenland. This NeXT-based system evolved through our iterative design strategy from 1993 to 1995. There was a gradual expansion of the community of users over this period, but the scientific use of the system was still mostly confined to providing real-time access to data sources in Sondrestrom. The principal use of the technology was to support coordinated, multiuser campaigns, most typically with a small set of users (2–5). Essentially, this was an Internet-based version of earlier practices that was developed through visits to the facility.

During these early years, we were able to document a number of changes in practice. Participation by a broader community of scientists increased, including a new mix of empirical and theoretical scientists. Younger scientists and students were able to participate, and we saw many examples of distributed mentoring where senior scientists at remote sites helped younger scientists and students. We saw greater flexibility in the scheduling of campaigns, including in the spring of 1995 two interleaved campaigns involving participants at a number of sites. In short, the UARC technology had a number of specific positive effects on scientific practice.

In the 1995–1996 period, three major changes occurred in the project that contributed to a true transformation in the practice of the science. First, we completely redesigned the UARC system. Whereas the early generations of UARC provided our users with a useful system for attaining more flexible access to the instruments at Sondrestrom, the client-server architecture we used had limited scalability. As the number of instruments supported increased, as more users joined real-time sessions, and as Internet traffic grew, performance degraded. Extensive testing revealed that the communication architecture did not work with growing Internet congestion when the number of users approached a dozen or the number of live instruments approached a half dozen. We needed a new architecture for UARC system, both to support the modes of collaborative work that we had observed so far and to support the additional functions, users, and data sources that we envisioned [2].

Second, the appearance of new World Wide Web technology in the midst of our project led us to switch our core technology strategy to Java applets accessed through Web browsers for our end users. This would give us (1) modularity, a key to our ultimate goal of providing a suite of tools that different communities of science users could shape to their own needs, (2) interoperability across all platforms, and (3) integration with other emerging Web tools and data sites. Building collaborative applications with this technology has been difficult, and the vision of platform independence has not yet been realized. The integration with other activities on the Web, however, has proved to be a significant boon and has helped us increase the scope of UARC very quickly.

Third, having seen what might be possible with our early generations of UARC prototypes, our growing community of users began suggesting to us entirely new kinds of activities the technology might support. In particular, three broad classes of new uses were suggested: (1) expansion of the data sources to cover a global field of view in real-time campaigns, (2) inclusion of computational models in real time, and (3) use of the UARC technology to support distributed, retrospective workshops in which the goal is to generate publications based on interesting events that have occurred and that the community has had some time to digest. During the 1996–1997 observational season, we supported two real-time global campaigns that addressed Items 1 and 2. We also experimented with some "replay campaigns" as an early version of 3, and during the 1997–1998 season, we will explore the support of a distributed workshop using UARC.

## UARC Use Transformed

In the past year, we have held several campaigns that bring together multiple data sources and supercomputer models in real time. Figure 2 shows the configuration of sites used in April 1997. In addition to the original Sondrestrom site, we were able to add additional radar facilities in Norway (EISCAT, at Tromso), Millstone Hill (Massachusetts), and an array of high-frequency radar across Canada called SuperDARN (Saskatoon, Kapuskasing, and GooseBay). The radar at Arecibo would have been included in April if it had been operational and will be added to future campaigns. We are also exploring the option of adding the radar at Jicamarca. As is evident from Figure 2, these sites span a wide range of latitudes in the Northern Hemisphere, providing a broad view of the upper atmosphere (albeit at a restricted set of longitudes).

We also had real-time data feeds from two satellites. Polar is a satellite over the North Pole that gives a broad view of the Northern Hemisphere. The screen dump in Figure 3 shows a Polar view (along with a theoretical

Gary M. Olson
University of Michigan
School of Information
304 West Hall
550 East University Avenue
Ann Arbor, MI 48109-1092
gmo@umich.edu
http://www.si.umich.edu

Daniel E. Atkins
University of Michigan

Robert Clauer
University of Michigan

Thomas A. Finholt
University of Michigan

Farnam Jahanian
University of Michigan

Timothy L. Killeen
University of Michigan

Atul Prakash
University of Michigan

Terry Weymouth
University of Michigan

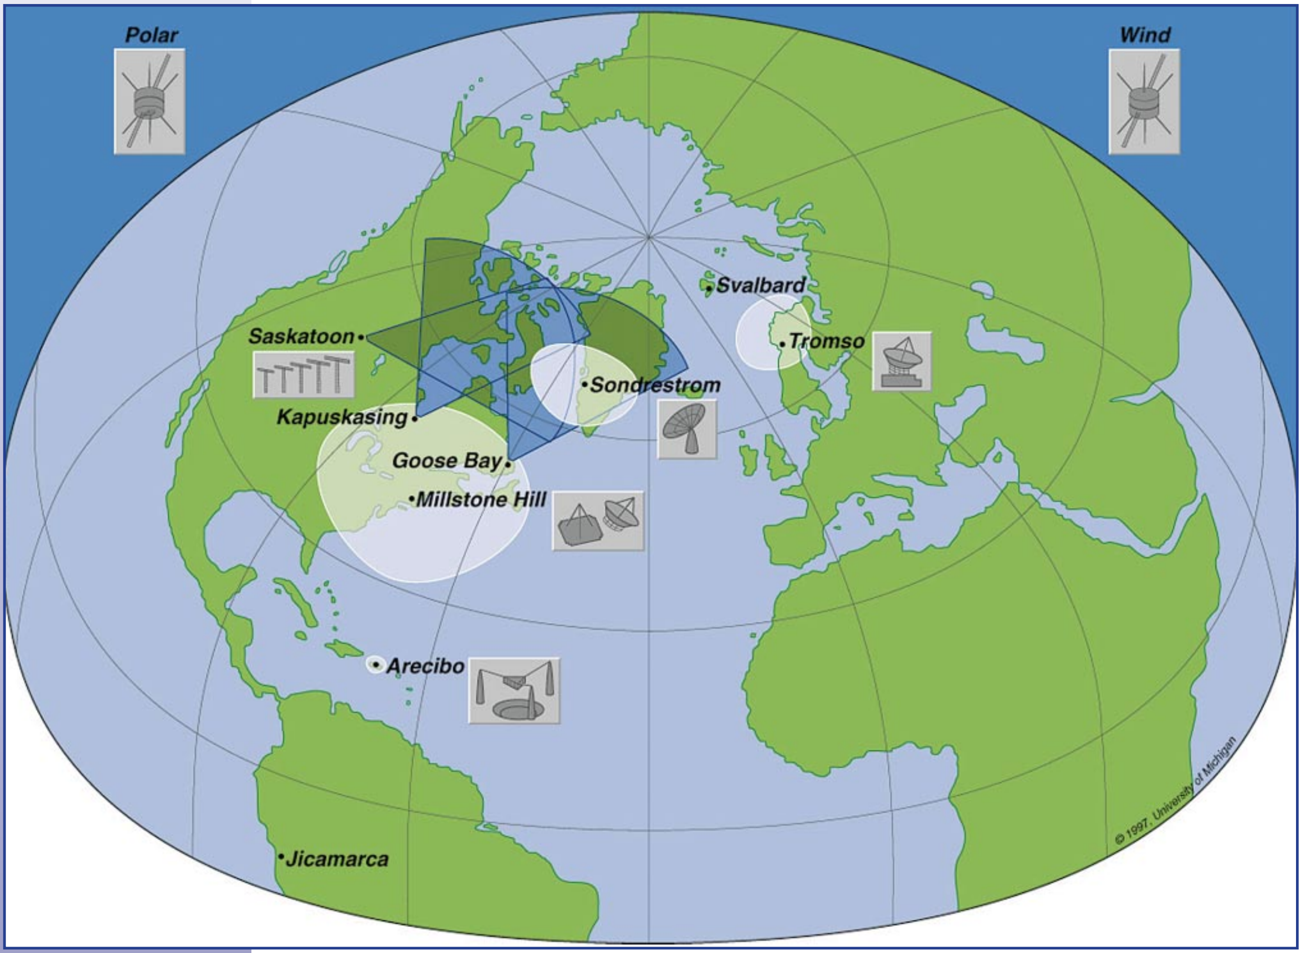


**Figure 2. The worldwide distribution of instrument sites used in the April 1997 Upper Atmospheric Research Collaboratory campaign.**

prediction of what this view should look like). Wind is a satellite that sits at the Lagrange point between the earth and the sun and provides early warning of characteristics of the solar wind.

Finally, the outputs of several supercomputer models of the earth's upper atmosphere were available. These outputs were prepared to look like the data feeds from the various instruments. In the screen dump in Figure 3, several such comparisons are shown. The globe in the upper right is a theoretical prediction for the Polar view shown in the upper middle. In the lower right are both data and theory for the radar at EISCAT in Norway. The supercomputer predictions were developed as much as 20–30 minutes into the future, allowing the observers to anticipate what might be happening. Furthermore, the models could be updated with revised parameters from the actual observations. This allowed the scientists to "close the data/theory loop," in the words of one of our users.

Figure 3 is a screen dump of the UARC technology as it appeared in the April 1997 campaign. Each user is free to configure their screen as they wish, though there are also tools in UARC that allow users to share exact copies of windows with their colleagues if they need to coordinate their displays for communication purposes [6]. The various data displays show both actual data and model predictions, selected from a wide range of options (many data sources, multiple views of each data source, and corresponding model predictions). On the left are two coordination and communication tools. In the upper left is the session manager, which provided overall organization for the campaign [3]. Functional clusterings of capabilities were grouped into rooms. Each room allowed as many Java applets and general URLs as were needed by the users. Rooms represented functional clusterings by science purpose, as well as places for the developers to interact and a help room. The numbers after a room indicated how many people were in it, a simple awareness mechanism. When one opened a room, the addresses of the people in the room were made available. This metaphor, borrowed from the MUD/MOO world, provided a useful way for the scientists to organize themselves.

In any particular room, most communication was through a multiparty chat facility. An example is shown in the lower right. Throughout UARC, these chat windows have provided the principal form of communication and have proved to be quite useful. We have analyzed these conversations in some detail, comparing computer-mediated conversations with comparable ones from face-to-face science campaigns at the Sondrestrom site [5]. The chat provides a persistent record of the conversation, a feature that is particularly useful when our users are scattered across so many time zones. Skimming through the chat log can be an easy way to establish context for a new user signing on in the midst of a campaign session.

These capabilities were initially explored in October 1996, and a full-blown campaign with broad participation was conducted in April 1997. During this April campaign, approximately 50 scientists participated, and their locations ranged from Alaska to the former Soviet Union. Performance with this many users was satisfactory, a marked contrast to the severely degraded performance of the old architecture when the number of users approached a dozen. Figure 4 shows communication patterns among the most active users on 1 day of this campaign. These data were obtained from logs of the chat facility and represent communications that were directed

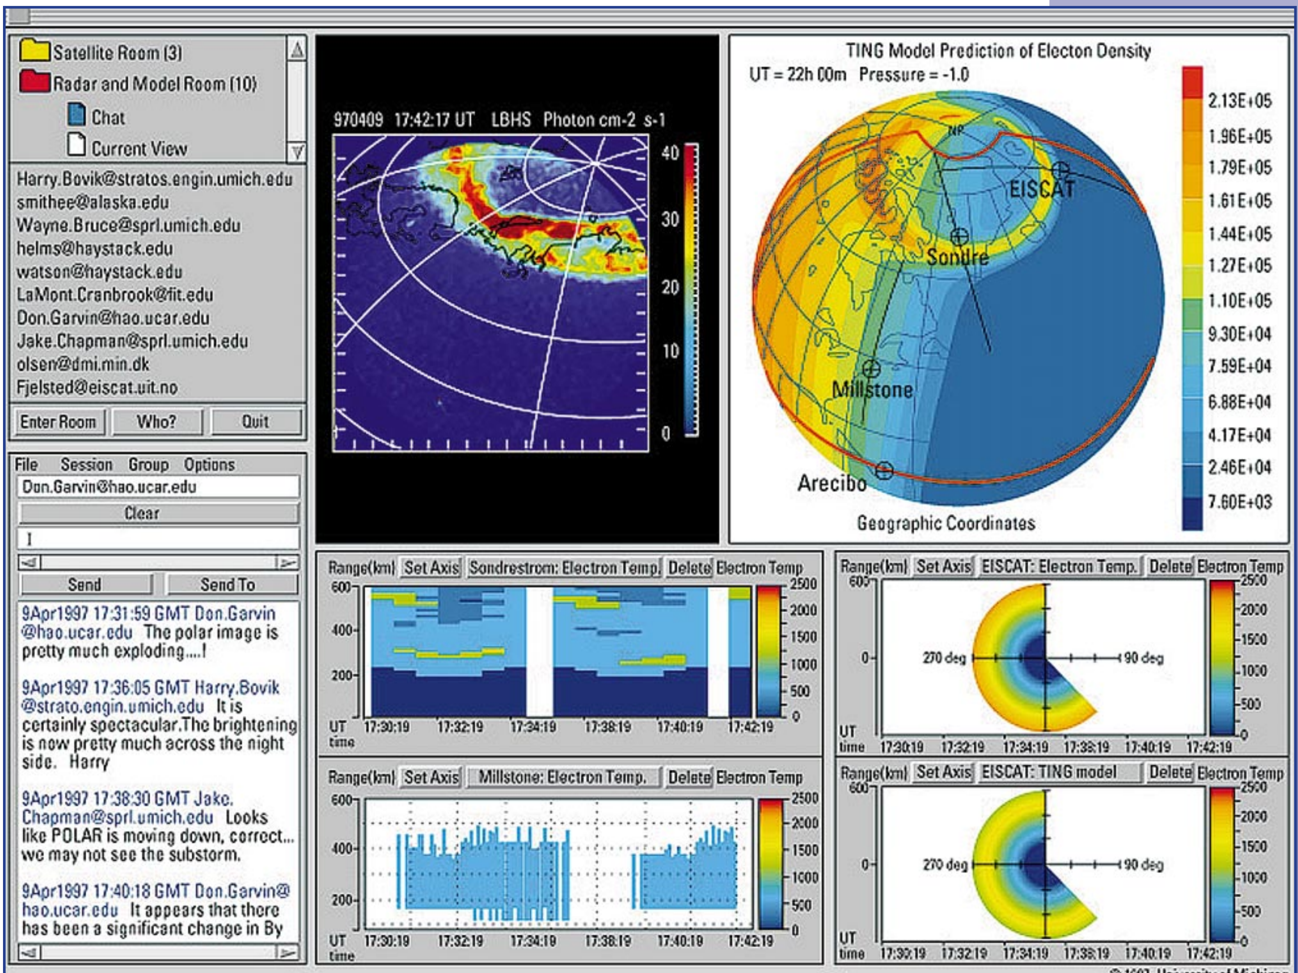


**Figure 3. Screen dump of the Upper Atmospheric Research Collaboratory user software during the April 1997 campaign.**

Pattern of Communication, UARC Campaign, April 9, 1997

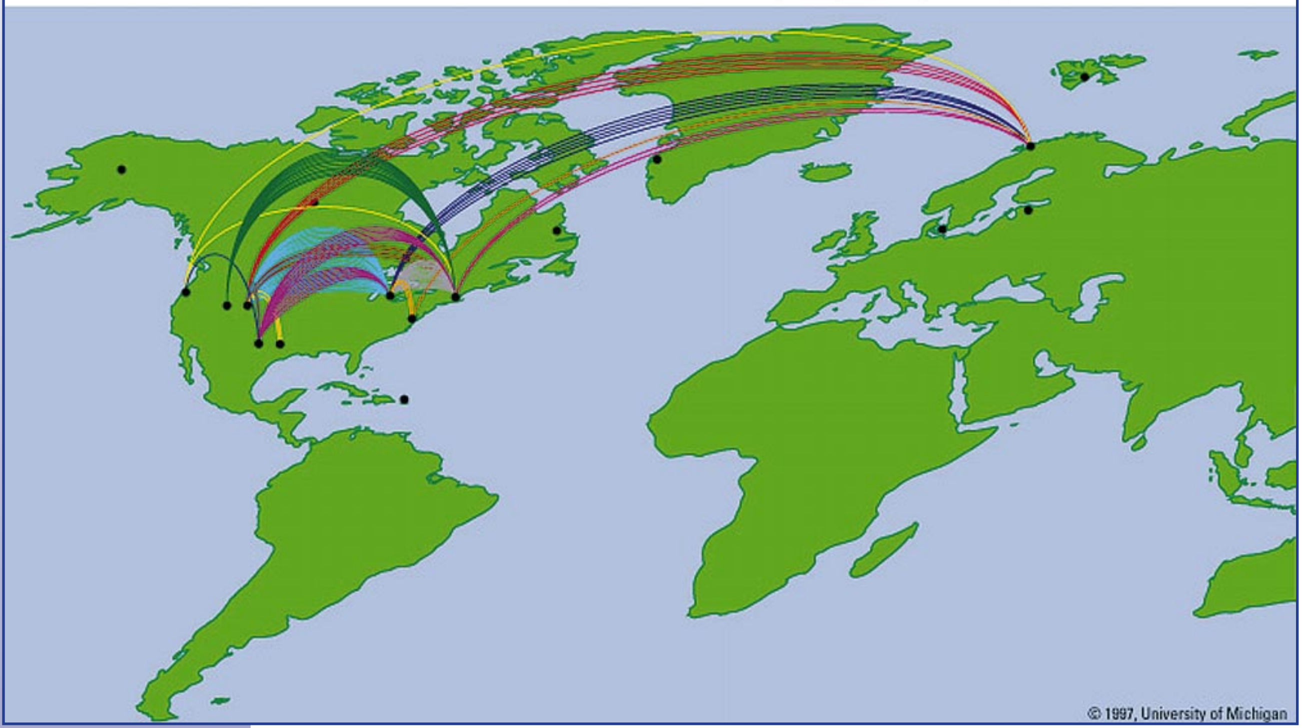


**Figure 4. Communication patterns among Upper Atmospheric Research Collaboratory users on April 9, 1997.**



from one person/site to another. The density of lines corresponds to the frequency of communication. This figure shows that the vision of a spatially distributed community of scientists conducting scientific conversations over phenomena of interest has clearly been established through UARC.

Interesting science took place during the April 1997 campaign. On April 9, an unexpected solar coronal mass ejection occurred, sending a major stream of charged particles earthward. There was extensive conversation in UARC about this. The ultimate major effects of this mass ejection on the earth's upper atmosphere occurred out of range of the UARC instruments (mostly over eastern Russia), but some effects were observed within the UARC field of view. This event in April is of such interest that the first UARC electronic workshops will focus on a retrospective analysis of this event.

## Summary

UARC has emerged as a mature collaboratory and has begun the process of transforming scientific practice in our user community. We like to think of what has happened as revolution through evolution. We began slowly, providing Internet-based ways of doing what had been done before. Gradually, as the community became familiar with this way of working, their vision grew. The most recent versions of UARC support a scale and character of real-time interaction scarcely imagined at the beginning of the project.

Real-time interactions over data collection, however, are only a very small part of the practice of these scientists. Using UARC to support interactive workshops is an obvious next step. These collaboratory workshops represent the kind of reflective deliberations over theory and data that move the observational work toward insight and ultimately scientific publication. We believe that modifications of our current UARC technology that seamlessly incorporates archival data, presentation software, and informal workspaces such as electronic whiteboards will facilitate such collaboratory workshops.

User-centered design with distributed communities of users is very difficult, particularly as we have attempted to work quite closely with our user community. We have had to work with considerable human-resource constraints and, as always happens in design, have had to make numerous trade-offs. The difficulties of building applications to support real-time collaboration over today's Internet have consumed large amounts of our development effort. It is likely that the commercial world is going to provide a number of generic capabilities to support collaboration in all of its forms, including real time. Advanced networks, such as Internet 2, may provide some relief to the technical challenges, though for a truly global effort such as UARC, there will always be considerable heterogeneity in the quality of network performance for different users. Some level of custom development of tools will also be required, as scientists adapt generic capabilities to the particulars of their science. Easy-to-use toolkits that allow such adaptation will be a requirement for functional collaboratories, and we aim to provide such a toolkit as a legacy of the UARC project.

### Acknowledgments

UARC is supported by a cooperative agreement between the National Science Foundation and the University of Michigan (IRI 9216848).

### References

1. Finholt, T. A., & Olson, G. M. From laboratories to collaboratories: A new organizational form for scientific collaboration. *Psychological Science, 9,* 1(1997), 28–36.

2. Hall, R. W., Mathur, A., Jahanian, F., Prakash, A., & Rasmussen, C. Corona: A communication service for scalable, reliable group collaboration systems. In *Proceedings of CSCW '96* (Boston, 1996). ACM, pp. 140–149.

3. Lee, J. H., Prakash, A., Jaeger, T., & Wu, G. Supporting multi-user, multi-applet workspaces in CBE. In *Proceedings of CSCW '96* (Boston, 1996). ACM, pp. 344–353.

4. McDaniel, S. E., Olson, G. M., & Olson, J. S. Methods in search of methodology: Combining HCI and object orientation. In *Proceedings of CHI '94* (Boston, 1994). ACM, pp. 145–151.

5. McDaniel, S. E., Olson, G. M., & McGee, J. Identifying and analyzing multiple threads in computer-mediated and face-to-face conversations. In *Proceedings of CSCW '96* (Boston, 1996). ACM, pp. 39–47.

6. Prakash, A., & Shim, H. S. DistView: Support for building efficient collaborative applications using replicated objects. In *Proceedings of CSCW '94* (Chapel Hill, NC, 1994). ACM, pp. 153–164.