



Coast-to-Coast Seminar: Background, History, and Practice



**Jonathan Borwein, David Langstroth, Mason Macklem, and Scott Wilson
D-Drive, Dalhousie University**









**Veselin Jungic
IRMACS, Simon Fraser University**





Contents

1	Coast-to-Coast (C2C) Seminar: Background, History and Practice	1
1	Introduction	1
2	C2C Seminar: Structure and Content	2
3	Technical Components and Issues	10
4	Lessons Learned	25
5	General Advice/Recommendations & How to Join!	26
6	Acknowledgments	28
	Bibliography	29







Coast-to-Coast (C2C) Seminar: Background, History and Practice

Jonathan Borwein, Veselin Jungic, David Langstroth, Mason Macklem, and
Scott Wilson



1 Introduction



The C2C Seminar (short for Coast-to-Coast¹) is a seminar run jointly at universities throughout Canada, from Simon Fraser University in British Columbia, to the University of Calgary and the University of Saskatchewan in the West, to Dalhousie, Memorial and other universities in the Atlantic Provinces. This seminar is simulcast to all sites via video-conferencing software, and each seminar provides opportunities for questions and comments from all of the remote locations.

The concept of the C2C seminar first originated with a large project called WestGrid. Starting in 2002, WestGrid was designed to be a massive parallel computing infrastructure to be shared by eight Western Canadian universities, although this number has since expanded to 14 institutions as the project has gradually expanded eastward from B.C. and Alberta to also include universities in Saskatchewan and Manitoba. At the time, Simon Fraser University (SFU) also had an interactive lab and seminar environment called the CoLab (for Collaborative Lab). This lab included a number of tiled touch-sensitive wall-mounted computer monitors, and was used for running courses and meetings, often remotely in cases where speakers were unable to attend the events personally.

As WestGrid progressed, the goal was to have similar ‘grid-rooms’ at each member university, to serve as local communication points for researchers who were working together on the WestGrid cluster from different institutions. In order to promote the resources that were available via WestGrid, a semi-regular event needed to be organized to show what could be done in terms of communication using this new infrastructure.

¹Along with all the technologies described here, our culture is also on the cusp of ‘text-speak’. Thus many of us call the seminar the “Sea-to-Sea” seminar, the text-speak version of our name. Which ironically produces an unintended alternative semantic - <http://www.cbc.ca/cp/0ddities/061109/K110902U.html>. (NT AL CHNGS R 4 TH BTTR)







In late 2003, as WestGrid was built and began to populate its network with users from each member university, the CoLab research group moved to the Faculty of Computing Science at Dalhousie, to construct a new research environment called D-Drive (for Dalhousie Distributed Research Institute and Virtual Environment), and with an additional goal of assisting ACEnet, a WestGrid-style shared network to connect universities throughout the Atlantic Provinces. During this same period, the CoLab environment at Simon Fraser was replaced by a much larger working environment called IRMACS (for Interdisciplinary Research in the Mathematical and Computational Sciences). Once D-Drive and IRMACS were completed, the potential for a cross-Canada video-conference was obvious, and since 2005 the C2C seminar has enabled audiences from throughout Canada to attend lectures by distinguished speakers from across the country.



In this paper, we will discuss the structure of the C2C seminar, including the technical and organizational components, and the lessons we have learned during the start-up process. We will also provide details for how interested people can connect to this seminar from their own local university.

2 C2C Seminar: Structure and Content

2.1 Structure. The Coast-to-Coast Seminar is an hour-long presentation given on a topic from mathematics or computational science and made accessible to audiences at a number of remote sites through collaboration technology. Seminars are held every two weeks throughout the academic year alternating between the West Coast and the East Coast. Initially the Western and Eastern sites were IRMACS and D-Drive exclusively, but as the series grew, and included other universities, presentations in the series have also come from Edmonton and Calgary in the West, and from Acadia, St. Francis Xavier and Math Resources Inc (a Halifax-based educational mathematics software company) in the East. In the near future presentations are also planned from the University of Lethbridge, Memorial University of Newfoundland and University of New Brunswick, among others.

Audiences for a presentation are located at one or more discrete sites at universities across Canada. The collaboration technology enables two way audio and video communication as well as a shared desktop. Thus a presenter is not only audible and visible to the audience, but can also respond to a raised hand, answer a question or interact with an individual at remote site through a shared application. The number of sites has increased to eight for an average presentation, with the promise of more participants in the future.









To set the stage for the presenter, we describe an outline of what a typical seminar entails. The actual presentation is expected to be of a high quality, yet accessible to a fairly general scientific audience. Accordingly they are widely advertised and attract audiences beyond the realms of mathematics and computer science, depending on the presenter's topic. We emphasize that we may have three or thirty people at one or other of the sites, and that typically perhaps 60 to 80 people hear each of the talks. No one has to come just to ensure a respectable audience as is often the case in a departmental colloquium.

The main goal of the seminar is to give an opportunity to scientific communities from various Canadian universities to collaborate and share their interests. At the same time we aim to achieve several other equally important goals:

- To learn and understand the issues related to the organization and the running of a regular scientific event from several universities in different time zones;
- To set the standards for this type of event for the future;
- To test the available technology;
- To motivate the creation of new technological tools and to encourage the improvement of the existing tools;
- To give a chance to faculties to gain experience with presenting through a still relatively new medium;
- To educate the audience attending the seminar about the protocols and etiquette involved;
- To reduce the costs of inviting distinguished speakers;
- To justify the investment in the technology and in the people involved;
- To build a C2C community.

We aim at an environment which is no less familiar than a new seminar room. As Ron Fitzgerald crisply puts it, "No one has to explain chalk." That said, we follow a fixed protocol each time. Roughly 30 minutes before the seminar starts, designated individuals from each site confirm that all facilities are working at all sites. An introduction of all sites and of the speaker is made from the speaker's site. The speaker's presentation is approximately 45 minutes long and is followed by a question and answer (Q&A) session with all sites. The Q&A session starts with local questions and then rotates through the remote sites. As with a face-to-face seminar, the host determines when to stop—and a good host has a first question to start things off with.









2.2 Past Talks. The presentations to date in the Coast-to-Coast seminar have been a mix of mathematical and computational talks, with a wide variety of focusses within each field. As of Fall of 2006, the C2C seminar has featured the following presentations²:

- Sherry Mantyka (Director, Mathematics Learning Centre, Memorial University), *The Math Plague: Learning Strategies for Under-Achievers in Mathematics* - December 2006, presented from Memorial University
- Gordon E. Swaters (University of Alberta), *Modelling Deep Ocean Currents* - November 2006, presented from the University of Alberta
- Ken Barker (University of Alberta), *Privacy Protection in Large Data Repositories* - November 2006, presented from the University of Alberta
- Laurence T. Yang (St Francis Xavier University), *Scalable integer Factorization for Public Key Cryptosystems* - November 2006, presented from ST. F.X. University
- Jonathan Borwein (Canada Research Chair, Dalhousie University), *Notes from the Digital Trenches* - October 2006, presented from D-Drive
- Steve Thompson (Shrum Chair in Science, Professor in Statistics, Simon Fraser University), *Sampling in Networks* - September 2006, presented from IRMACS
- Ron Fitzgerald (President, Math Resources Inc), *Learning Infrastructures and Content Authoring* - March 2006, presented from the Halifax offices of Math Resources Inc
- Bojan Mohar (Mathematics, Simon Fraser University), *Hadwiger's Conjecture* - March 2006, presented from IRMACS
- Carey Williamson (Computing Science, University of Calgary), *Things that Go Bump on the Net* - March 2006, presented from the University of Calgary
- Jeff Hooper (Mathematics and Statistics, Acadia University), *L-Functions and Arithmetic* - February 2006, presented from Acadia University



²This includes the first full academic year of the Coast-to-Coast seminar, with the second year currently in progress.







- Alejandro Adem (Mathematics, University of British Columbia), *Periodic Complexes and Group Actions* - February 2006, presented from IRMACS
- Przemyslaw Prusinkiewicz (Computer Science, University of Calgary), *Computational Biology of Plants* - February 2006, presented from the University of Calgary
- János Pintér (Pinter Consulting Services), *Teaching OR/MS Using Integrated Computing Systems* - January 2006, presented from D-Drive
- Jonathan Schaeffer (Computing Science, University of Alberta), *Solving Checkers* - January 2006, presented from the University of Alberta
- Andrew Rau-Chaplin (Computing Science, Dalhousie University), *Parallel Applications in Phylogeny* - December 2005, presented from D-Drive
- Arvind Gupta (Computing Science, Simon Fraser University), *The Inverse Protein Folding Problem* - November 2005, presented from IRMACS
- Karl Dilcher (Mathematics and Statistics, Dalhousie University), *Wieferich Primes and Fermat Numbers: Computations and Generalizations* - November 2005, presented from D-Drive
- Uwe Glaesser (Computing Science, Simon Fraser University), *Semantic Blueprints of Discrete Dynamic Systems* - October 2005, presented from IRMACS
- John McHugh (Computer Science, Dalhousie University), *Pyrite or Gold? It takes more than a pick or shovel* - October 2005, presented from D-Drive
- Peter Borwein (Executive Director, IRMACS, Simon Fraser University), *The Riemann Hypothesis* - September 2005, presented from IRMACS
- Jonathan Borwein (D-Drive Director, Computer Science, Dalhousie University), *Mathematical Visualization and Other Learning Tools* - September 2005, presented from D-Drive

During the summer of 2005, two test sessions were held. The first session consisted of several short presentations that ran from both IRMACS and D-Drive and were given by graduate and undergraduate students. The presenters were asked to use various methods and tools in delivering their









talks: Power point presentations, PDF slides, pre-prepared transparencies, writing on a white board, writing on paper and using a docucamera, using Maple applications, and so on. After summarizing the experiences from the test session, a format for the future C2C presentation had emerged. This format was tried during the second test session when Colin Percival from Simon Fraser University gave the presentation with the title *Hyperthreading Considered Vulnerable*.

Following the success of the C2C Seminar Series over the 2005-2006 academic year we hosted a more intensive distributed event, The Coast-to-Coast Miniconference on the Mathematics of Computation. This day-long event consisted of a series of six speakers, alternating between IRMACS, the University of Lethbridge, and D-Drive. The event was attended by audiences in each of these locations as well as in some of the other remote sites, according to interest and availability. Since then we have also experimented with shared open houses and other ways to experience ‘presence-at-a-distance’.

2.3 Organizational Issues. The growth of the C2C Seminar significantly increased the number of people involved in the organization and the running of the seminar. To guarantee full technical support, each site needs a technician present locally during every presentation. Since WestGrid and ACEnet employ most of those people, the seminar heavily depends on the willingness of the two institutions to be involved in the C2C project. In addition, for recruiting speakers and advertising presentations locally, it is important to have C2C liaisons at as many universities as possible. Clearly the liaisons should be faculty members ready to put their valuable time in the process of creating something that is as new and complex as the C2C series. However, as the size of the core group increases as more universities become involved, the overall organization becomes more complex. Thus ideally the number of people involved in organization of the C2C Seminar at all levels should be about two dozen. The size of the group immediately implies the question of the hierarchy and communication within it. Currently (November 2006) we are working on the establishment of a cross-institutional e-mail list that would serve as a forum for all people involved in running the series. Also, in an effort to centralize the process of coordinating all of the people involved, the position of Seminar Coordinator will be introduced starting in the Spring 2007 term. The main goal is to centralize scheduling and announcement distribution, and to simplify communication between the C2C group and the rest of the academic community.

One of the challenges in coordinating a series of distributed events across a large geography like Canada is the need to constantly consider time zones. There are six different time zones from British Columbia on the West Coast to Newfoundland in the East. This restricts scheduling possibilities, for a









seminar scheduled at the reasonable hour of 11:00 am in D-Drive on the east coast would be a little too early for audiences on the West Coast at 7:00 am. Seminars are presented therefore at 3:30 pm Atlantic Time which is 11:30 am Pacific Time and 4:00 pm Newfoundland Time.

The issue of time zones also affects communications and advertising between West and East: there are ample opportunities for misunderstanding unless each stated time is explicitly labeled with the local time zone. To state times in terms of your own local time zone and rely on the recipient of communications to translate into their own local time has worked the best for us. Early attempts to translate communications into every applicable time zone resulted in an increase in errors and misunderstandings.

Another issue that has arisen from having multiple academic institutions involved in this event is the fact that scheduling around the academic timetable becomes potentially more complex. It is difficult to guarantee that a given seminar will not conflict with other activities at one or more universities. We also need to have the local WestGrid and ACEnet nodes be available at the time of our talks. Our solution to these potential difficulties is to simply choose a regular time at the outset and stick to it, with the understanding that there may be conflicts with the local academic schedule, but that the regularity of the event gives some possibility for Math Departments, Computer Science Departments, or other regular attendees to organize their schedule to accommodate the Seminar Series. We are also able to record the entire event and so allow it to be watched by a class or other interested group at another time.

Finally, advertising the seminars is complicated by their distributed nature. It is necessary to make clear on any particular advertisement both the time and the place of an event; however, the time depends upon the local time zone, and there are sometimes several places at which the presentation can be attended locally. Our solution is to distribute communications through a hierarchical system: once the details of an event are determined and agreed between the IRMACS and D-Drive administrators, then each of them distributes the information to their respective coasts. In D-Drive this means the administrator is responsible for passing the information on to St. Francis Xavier University, Acadia University, Memorial University and other interested parties. At each of these final destinations the recipient advertises to their local audience using local time and the local point of attendance. A central listing of Seminars is also maintained on the D-Drive website: www.cs.dal.ca/ddrive/seminars. Similarly, on the west coast, communications move outwards from the IRMACS administrator to the other universities and each advertises locally. The announcements are also posted on the WestGrid website together with the list of all seminar locations in Alberta and British Columbia. As there are only currently two time zones and three regular points of attendance on the East Coast, a





joint poster has been designed by ACEnet's Chief Technical Officer, Greg Lukeman, at St. Francis Xavier University which carries the multiple time zone information and multiple attendance locations. It has the advantage of presenting a consistent image to the whole Eastern community at the expense of being a slightly denser packet of information. A version of the poster for Western Canada is made and distributed by WestGrid.

2.4 Organizing Content. The Coast-to-Coast Seminar Series has arisen from a partnership between the two founding members, IRMACS and D-Drive, and although it has grown to include other participants it still retains that basic organizational hierarchy. Planning decisions are made bilaterally between the two founders and then passed on to the other members. At the beginning of the academic year a timetable is drawn up for the series, with presentations alternating between East and West. For the Eastern series of presentations, the Director of D-Drive, Dr. Jonathan Borwein, then compiles a set of suggested speakers, who are approached by the administrator until a full roster of speakers has been confirmed. The director will also indicate a desired distribution of speakers from the other Eastern Universities and then the task of finding a speaker at a remote site passes to a representative of that site.

A protocol for organizing the process between D-Drive and IRMACS was drawn up at the beginning of 2005 and has served as a useful guideline, although it has not always been possible to follow it to the letter. It includes the following procedures:



1. Two-Three weeks before the lecture:



- Confirm the booking location
- Get the title and abstract of the talk, send them to all remote sites, and post them on the web
- Inform the local technician about the date and time of the lecture

2. One-Two Weeks before the lecture:

- Contact the presenter and find out in which way they will deliver the lecture. Acquire a copy of the presenter's material³ to be used in testing. Forward the presentation to the local technician

³While a lecturer in D-DRIVE could in principle write directly on a *SMART Board*, we discourage this (except say in answering questions or for annotation) as it takes a good deal of experience to do this effectively. Likewise, we do not encourage writing directly on a projector or something similar. Advance preparation is really advisable, and digital transmission is preferable to analogue.





- Check if the technician has confirmed testing times with all remote sites
- Send out the first email notice to all sites to invite people to the lecture

3. Two Days before the lecture:



- Final lecture reminder to listserves and contacts on campus and surrounding institutions
- Final consultation with the technician
- Reminder email to all sites that we start set up one hour in advance
- Final email to presenter for any last minute concerns



In practice, much of the process has become automatic. Sometimes, however it proves difficult to pin down a speaker far enough in advance and the whole process compresses into a shorter time frame. Some speakers have to be reminded quite vigorously to submit a title and abstract well in advance.

In D-Drive, and in IRMACS, the administrative role and the technical role are separate. This has also led to some adjustment in the protocol as it often makes more sense for the technician to communicate directly with a speaker about testing issues or about their presentation materials if there is a technical question to be addressed. Good lines of communication between the administrator and the technician are important for good coordination of these activities, and in D-Drive these two people inhabit the same general lab space to facilitate communication.

2.5 Seminar Facilitation. Although we have been interested in producing a seminar experience for distributed audiences which is as close as possible to a face to face event, there are significant differences of which a local audience needs to be aware. In the first place, with increasing numbers of remote participants displayed on screen it becomes more and more difficult for a remote audience member to attract the speaker's attention with a question. Thus the question period must be handled explicitly, with an active request for questions, usually at the end of the presentation. Even then, the audience member needs to be made aware that if they have been chosen to ask their question, they must wait until a wireless microphone is passed to them so that it will be clearly audible across the network.

The orientation of speaker to audience also requires a shift in the usual expectations. In the D-Drive lab, for example, the audience will see, facing them, the faces of all of the remote audiences displayed on the large monitors. A local speaker may choose to address both the live and remote





audiences at the same time, shifting his attention from one to the other and standing in profile so that he can turn his head to make eye contact with the local audience or turning the other way to face his remote audience on the screen. Or, some speakers have chosen to sit facing the large monitors, speaking directly to the remote audience, but with their back to the local audience. Our experience has been that the former choice is usually the more successful, but we leave each speaker to find the orientation which they are most comfortable with.

In the hour leading up to the beginning of the presentation, conversations from one site to another across the technology link will usually be the business of the technicians, checking connections and levels. When everything is ready and the seminar begins, the microphone will be taken by either the Director or the Administrator at the site which is presenting the talk, in order to welcome audiences and to introduce the speaker.



For almost all of the speakers who have presented at the C2C Seminar, this was the first experience with giving a scientific presentation to an audience that was located in various locations of the country. Talking to the remote audience through an advanced but still very new technology is an additional challenge in communicating advanced scientific topics. One must appreciate the fact that the C2C presenters have been ready to take the risk by pioneering in the C2C experiment.



3 Technical Components and Issues

3.1 Technical Overview. The technology behind the Coast-to-Coast Seminar is a combination of open source software, standard PC hardware, and audio/video components. The structure is a client/server architecture, in which individual sites authenticate to a central coordinating server, with audio, video, and presentation data shared between all sites.

The seminar organizers chose to standardize on Argonne National Labs *Access Grid* (AG) software as the video-conferencing suite [1,2,4,5,7]. This selection was made for three primary reasons: 1) AG is quite flexible in site configuration, allowing full auditoriums with complex audio systems and multiple cameras to conference with individual PCs using webcams and headsets in the same collaborative session, 2) AG is platform independent, with clients available on Windows, Linux/Unix, and MacOS, and 3) AG is a highly scaleable video-conferencing package, allowing up to 30-40 remote sites simultaneously (limited by available bandwidth and processing capabilities). Some examples of the use of Access Grid in remote collaboration can be found in [6,8].

Access Grid sessions are coordinated through an *Access Grid venue server*, which builds upon a “rooms” and “lobbies” analogy to coordinate







collaborative sessions. For the C2C Seminar, the WestGrid AG venue server is used as a meeting place for the virtual attendees, with the sites meeting in a virtual conference room on the server. Audio and video from all client sites are shared through the venue server, and the venue server also provides file sharing, presentation syncing, and limited chat capabilities. Information about the Access Grid layout in WestGrid can be found in [3].



Client sites provide information to the venue server about their audio and video capabilities, and this information is distributed to the other sites via the venue server. In the C2C Seminar, there are two primary client options: *Access Grid* clients and *InSORS* clients. The *Access Grid* client is provided by Argonne Labs as part of the *Access Grid* open source project. *InSORS* is a commercial extension of the *Access Grid* project that provides increased video quality and some additional collaboration tools, but is only available on Windows and MacOS. Each site uses whichever client their site finds most appropriate for the seminar series, and adjusts client settings for compatibility with the other sites. The C2C Seminar is dedicated to remaining platform independent due to the nature of the many different academic sites participating, and therefore the Seminar has standardized to toolsets that all platforms can use.

Desktop sharing from the presenting site (to view lecture notes, slides, and whiteboards) is provided via the open source *Virtual Network Computing* (VNC) software package. The C2C Seminar makes use of a VNC server product called *VNC Reflector* to create consistent connection-point and authentication information for all sites from week to week, with the presenting site delivering presentation information to the *VNC Reflector* via a VNC server. The VNC products are platform independent open source projects, in keeping with the Seminar's technical goal of maximum flexibility for connecting sites. To facilitate whiteboard viewing and presentation markup, the primary hosting sites (Dalhousie and SFU) provide SmartTech *SMART Boards* as drawing surfaces for their lecturers, which are then displayed over the VNC connection.

3.2 Sample Presentation Environment: D-Drive and IRMACS Layouts. The component technologies driving the Coast-to-Coast Seminar can best be shown by giving a description of the layout for one of the site locations for the seminar series, namely the D-Drive lab in the Faculty of Computing Science at Dalhousie University in Halifax, Canada. Although this research lab shows some of this technology in action, it would be a mistake to believe that one needs all of it in order to connect to the C2C seminar series - as mentioned earlier, one can connect simply with a web camera and the proper desktop software installed.

The D-Drive lab consists of five touch-sensitive computer displays, four cameras, several microphones, an echo cancellation unit, and a sound sys-









tem, which together provide a suitable presentation and audience environment for a remote seminar. The displays, shown in Figure 3.2, consist of a centrally located 73" rear-projected monitor, along with four 61" plasma screens, with two each on either side of the rear-projected monitor. These screens are separated into two sets of tiled displays as shown in Figures 3.2-3.3: the two left-most plasma screens are tiled as a single display, and are used for audio/video management and video feeds, while the three remaining screens are also tiled as a single display, and contain the presentation slides (on the central rear-projected display (shown in Figure 3.4) and various video feeds from other remote locations. Each set of tiled displays has a single point of focus, meaning that two users cannot simultaneously control different positions on different boards within the same tiled set of displays; therefore, the five displays were separated into two tiled sets in order to allow for local audio/video management and changes to video feeds (on the left tiled set) without taking control away from the presenter and his slides on the central display.

There are three locations for the audience within D-Drive to view a presentation: a conference-table directly in front of the displays, seating area in the centre of the lab, and tables along the back wall of the lab. The conference-table is generally used for multi-site meetings, where there is no central presentation and where control of the discussion frequently changes between sites; for such events, the conference table has a centrally-located microphone that picks up general discussion by anyone seated at the table. For most C2C seminars, the central seating area will be the primary location of the audience, with the conference-table and the tables along the back wall serving as "overflow" seating when necessary. During the seminar, questions by audience members throughout the lab are asked via a wireless handheld microphone that is passed to them when they indicate that they have a question.

The four cameras are spaced throughout the D-Drive lab, with two mounted on the top of the front displays (shown in Figure 3.4) and two more in the two back corners of the lab (shown in Figure 3.5). The cameras are placed in such a way that they give a sense of room context for remote sites, by providing multiple points of reference for the activities within D-Drive. The two front-mounted cameras provide the view of the local audience to the remote sites, with the left-camera being controllable by remote control in order to allow for panning and zooming when local audience members ask questions to the presenter, while the right-camera has a constant position that shows an overview of the audience seating.

The two rear cameras provide alternate views of the lab, and are primarily useful for when the C2C presenter is speaking from D-Drive, or when someone is interacting with the *Smart Board* screens. In particular, the right-rear camera is also controllable by remote, and can provide a





close-up video feed of the presenter to the remote sites, whereas the left-rear camera provides a wide overview of the screens and conference table. These cameras are not generally used when the presenter is located at another site, though they are quite useful for collaborative meetings in the facility.



In addition to the displays and the seating area for the audience, the D-Drive lab also has a desk located to the right of the display environment, where the local Technical Supervisor is located, as shown in Figure 3.6. This person controls the audio and video management on the left-most plasma screens, and is in contact with the Technical Supervisors at each of the other remote sites via instant messaging software on their respective local desktop computers. In case of technical problems, solutions are determined via discussion between sites, in part so that sites that are new to the seminar can get suggestions from more experienced sites.

The D-Drive presentation environment is just one of many within the C2C Seminar Series. To highlight some of the variety that exists within the various locations for the C2C series, Figures 3.7 and 3.8 show the display/presentation environment at IRMACS. This environment is contained within a lecture theatre, with three touch-sensitive plasma displays and one or two projected displays. The presentation slides are shown on the larger projected display, as is the video feed for the presenter if located at a remote site. Figure 3.8 shows the location of the IRMACS Technical Supervisor within the lecture hall.

Although both D-Drive and IRMACS have extensive technology driving their respective display environments, there is no expectation that universities need to have a similar environment in order to take part in the C2C Seminar Series. In fact, one of the goals of the seminar series is to make it as easy as possible for universities to use their local space as presentation and display environments.

§ **General Technical Issues with Distance Collaboration.** Technical challenges discovered while providing a distance collaboration experience such as the Coast-to-Coast Seminar are often strongly related to policy, organizational, and educational issues in the various institutions. The goal of the C2C Seminar from a technical perspective was to build a collaborative seminar experience in which the technology is a transparent mechanism for communication, rather than a barrier or limitation to communication. With that goal in mind, familiarity with the available technology at each of the sites becomes a critical component to building policies and procedures which result in a technologically seamless seminar experience.

§ **Site Preparation and Testing.** As an introduction to general technological issues, we cannot emphasize strongly enough the importance of testing and







Figure 3.1. The computer displays in the D-Drive research lab, one of the locations for the C2C Seminar series.






Figure 3.2. Left-pair of tiled plasma screens, used for audio and video management (left) and presenter's video-feed (right).

site preparation for the seminar experience. Each institution should have a local technician on-site who is both familiar with the site equipment and has conducted thorough tests with other participating sites to test their equipment. To facilitate site preparation, checklists for site participation should be developed, and reviewed regularly by the on-site technician for









Figure 3.3. Right-pair of tiled plasma screens, used to display video feeds from various remote sites, with the IRMACS lecture hall in SFU (left-most video feed), along with remote feeds from St.F.X. in Antigonish, Nova Scotia (left) and Memorial University of Newfoundland (right). Several video feeds from the D-DRIVE cameras are also shown.

an institution. The heterogeneous nature of site equipment will result in some variance of individual site checklists, but all site checklists should generally cover the following areas:

- Stability of hardware used for the site
- Reliable client connection to the venue server
- Audio quality tests, both receiving and sending
- Video quality tests, both receiving and sending
- Desktop sharing tests, both receiving and sending

Institutions presenting a lecture have an additional set of tests which should be run in advance of their presentation date. The first is that the presentation content should be initially reviewed locally on the presenting equipment to ensure that there are no issues with display or audio anomalies. This gives the speaker an opportunity to amend their presentation, or for local technical staff to make adjustments to the presenting system.

In addition, the presenting site should review with remote sites the quality and speed of the presentation materials running over the desktop sharing software. Current desktop sharing software can suffer from update delays that may be acceptable for a static slideshow, but completely unusable for a video clip inserted into a presentation, a highly graphical presentation, or a local program animation. Advance review of the presentation materials with remote sites will catch technology limitations with desktop sharing software and allow for either:






Figure 3.4. Central rear-projected display, with two mounted front-view cameras (top, cameras circled), with view from behind the left-camera (bottom-left image), and video-feed of left-camera (circled in bottom-right)

- distribution of the material to all sites for local execution of the material,
- use of an alternate remote distribution technology, such as using the *Video Lan Client* (VLC) to broadcast presentation videos, or
- adjustment of the presentation to account for the technology limitation.

Finally, the presenting site should ensure that the speaker has been given an introduction to the presenting system and the technology tools available to them. The interface for items such as whiteboard drawings,





Figure 3.5. Camera views from back-left and back-right corners of the D-Drive research lab (left and right, respectively).



Figure 3.6. Desk of Technical Supervisor for D-Drive, located to the right of the display-environment, with contact with Technical Supervisors for all of the remote sites via messaging-software on the desktop machine.

presentation controls, and how they interact with remote sites should be reviewed with the speaker before the presentation date. Speakers should also be made aware of how to highlight their presentation in such a way that remote sections of their audience can see it. This has been demonstrated







Figure 3.7. The presentation environment at IRMACS, for local (top) and remote (bottom) presentations during the C2C seminar.

vividly to seminar audiences on several occasions in which the speaker pulled out a laser pointer partway through the presentation and proceeded







Figure 3.8. The location of the desk of the Technical Supervisor at IRMACS (left-hand circle), and the rear video-camera, mounted on the back-wall of the IRMACS lecture theatre (right-circle).



to highlight sections of the presentation while lecturing, completely forgetting that the remote audiences would not be able to see their indicator. In such situations, it is a delicate decision as to when it is more disruptive to the seminar to correct the speaker, or to live with the problem. Such decisions are often debated and made by the technical group using an instant messaging client in the background during the presentation.

§ **Live Technical Monitoring.** After thorough site preparation and testing, the next critical factor for technical success is live technical monitoring at each site during the presentation. Adjustments will periodically need to be made to technical equipment throughout a presentation, and in a distributed lecture, we do not have the option of just using the blackboard for the remainder of the talk. A technician should be available at each site to provide feedback to the other sites on technical issues and to intervene locally if there is a significant issue.

3.3 Technical Compatibility Issues between Sites.

§ **Network Compatibility.** The most common issues seen when first getting a site set up for the Coast-to-Coast Seminar are with the site's network. *Access Grid* uses a network protocol called *IP multicast* to communicate







audio and video information to all sites. This protocol allows for a much more efficient communication of audio and video information from a site, greatly reducing the network bandwidth required to send your audio and video information out to multiple sites. If a site's network is not capable of supporting IP multicast, that site will not see video or hear audio, even if they are successfully connected to the AG venue server.



There are two common solutions for this problem: talk to the institutional network engineers about enabling IP multicast on that site's network, or connect to the other sites using an intermediary site (or proxy) that does support IP multicast, called a "unicast bridge". The *Access Grid* client has a built-in utility to connect to unicast bridges for a venue server, and we recommend starting troubleshooting with that utility when a given site cannot see audio and video from other sites. Lack of IP multicast support on the network is the most common problem when first setting up a site, and oftentimes IP multicast will not be an option at all when connecting through a commercial Internet service provider.

The other common connectivity issue for new sites is that a firewall is set to block the network ports used to communicate audio and video information. The ports used to transmit audio and video information are determined by the *Access Grid* venue server. Those ports will need to be "opened" or "unblocked" for audio and video communication. In the case of the Coast-to-Coast Seminar, we have a document created by WestGrid (our venue server hosts) to send to new sites that lists the port ranges which need to be opened for *Access Grid* collaborations.

While not common yet, one network issue that will create additional complexities in the near future is the expanding use of NAT (Network Address Translation) devices. These devices are used on networks to allow many different IP addresses to share a single IP address with an outside network. While these are commonly used now on small home networks (often as part of a "broadband router"), there is an expanding role for such devices in many IT organizations. NAT devices require some complex "port forwarding" scenarios for *Access Grid* clients, and we recommend avoiding such devices on networks that will have *Access Grid* clients.

The final network issue that should be considered is available bandwidth. The average *Access Grid* session will generate approximately 800 Kbps per camera being broadcast. For general planning purposes, we recommend having approximately 15 Mbps of bandwidth available for collaborative seminars such as the Coast-to-Coast Seminar. This is usually not an issue on most university networks, but it can become an issue when sites are connecting over a DSL line or other home broadband connections. In our experience, cable modem services of 5 to 10 Mbps can support collaborative sessions fairly well, but slower connections (such as the average DSL connection) will have issues with bandwidth.





§ **Client Compatibility.** Several technical compatibility issues were discovered while attempting to integrate various *Access Grid* and *InSORS* clients into the Coast-to-Coast Seminar since its inception. Currently, the Coast-to-Coast Seminars are standardized on the *Access Grid 2.4* client, and the recently released *Access Grid 3.0* client is undergoing testing at some sites. Our recommendation for collaborative seminars in general, based on the Coast-to-Coast experience, is simply to watch which versions of video-conferencing software a site is using, and have all sites standardized on the same version of the client software.



When using multiple client types for a seminar (such as *Access Grid* and *InSORS* clients) it is critical to ensure that your audio and video streams are using codecs that all sites can understand. This often means that the highest common standard is selected for audio and video. In the specific case of the Coast-to-Coast Seminar, we have standardized on the H.261 video codec for all sites. A common issue for sites in our seminar series using the *InSORS* Windows client is that *InSORS* video has been set by default to a higher quality codec (H.264) which cannot be decoded by *Access Grid* clients, or *InSORS* clients running on MacOS.



The final issue to address in the area of client compatibility is desktop sharing software. We have already mentioned that the Coast-to-Coast Seminar organizers selected VNC software due to its cross-platform availability. Even with a platform-independent software utility such as VNC, there are compatibility issues to watch. VNC clients are not universally consistent in their features across platforms. For this reason, we have developed standards to address desktop sharing compatibility.

One example of desktop sharing standards is simply to limit the presentation screen resolutions to standard sizes. Some sites have high definition television screens for their presentation systems, while others have standard portable LCD projectors. If a presenting site sends out presentation information in a high-resolution widescreen format, it becomes highly problematic for sites with smaller resolution displays. While some VNC clients can simply scale the display to an appropriate video resolution, VNC clients on other platforms do not have that capability. For that reason, the Coast-to-Coast Seminar has standardized all main presentation displays at 1024x768, and limits applications such as whiteboard areas to standard projector sizes.

3.4 Audio/Video Issues.

§ **Audio Production.** Over the course of the first year of the Coast-to-Coast Seminar, audio clarity was identified as the most critical single component for a remote collaboration experience. Video and desktop presentation quality can fluctuate and even have intermittent interruptions without a





significant loss to the quality of a lecture, but audio quality loss almost immediately affects the effectiveness of the communication.⁴

For that reason, we strongly recommend that some careful thought be given to both audio components and room design in a larger *Access Grid* site. A quiet environment in which outside noises are minimized is critical for any large speaking environment, but in a large video conferencing environment additional consideration must be given to limiting audio noise within the environment. This is simply because the echo-cancelling technologies available today for video conferencing have limitations in their efficiency, and the technology usually involves frequency cancellation and noise reduction algorithms. If too many frequencies are generated by a site due to background noise, echo cancellation units can end up cancelling too many frequencies and interfering with communication.



To minimize this background noise effect, we recommend that boundary microphones only be used for *Access Grid* sites with fewer than 8 people, and that directional handheld or lapel microphones be used for sites with more individuals. We have also established protocols for the Coast-to-Coast Seminar to minimize background noise during presentations - for instance, all remote site microphones must be muted for presentations other than during specific question and answer sessions.



Controlling audio quality is the primary task for technicians monitoring live Coast-to-Coast Seminars. Oftentimes the lecturer's volume or the audio environment will change throughout the presentation, and local technicians are responsible for adjusting audio levels to compensate. In this task, they are serving the same function as an audio technician for any other production, whether music, theater, or speaking - but they have the additional complexity of monitoring issues created by the distributed nature of the communication. In these situations the feedback of remote site technicians becomes crucial, as there are times that audio distortion can occur remotely (due to microphone levels overdriving slightly, combined with packet loss on the network) when it would not be noticed by a local technician.

§ **Video Production.** The Coast-to-Coast Seminar series has also highlighted the usefulness of video production methods as an area of expertise for technicians involved in remote collaboration. Minor production adjustments in the video presentation can make a vast difference in the perceived quality and effectiveness of the seminar experience. These video production adjustments can be items as simple as:

1. **Camera location:** Placement of cameras in such a way that speaker

⁴Our experience certainly is consistent with hearing-disabled advocacy groups frequent assertions that deafness is more isolating than blindness in modern society.





movement is captured from optimal angles and distance, and in a location where it appears the speaker is looking into the camera when facing his presentation materials and remote audiences.

2. **Background and lighting issues:** A common problem when setting up sites is an incorrect lighting environment for the cameras to pick up a speaker well.



One example from the D-DRIVE Lab that illustrates lighting issues well - when lecturers in the D-DRIVE Lab would stand at a *SMART Board* video monitor to write on one of the whiteboards, the brightness of the white video display behind them would cause the cameras to wash out the speaker completely. By simply changing the on-screen “whiteboard” to a “blackboard” (changing the background color of the application from white to black), the video cameras were able to adjust their brightness settings and send out a viewable image of the lecturer writing at the board.⁵



3. **Remote speaker location:** Display of the speaker and their presentation should be oriented so that the speaker, when indicating their local presentation content, is also gesturing in the direction of the content at remote sites. Remote audiences find it highly distracting if the speaker and content windows are placed on displays at remote sites so that the speaker is gesturing in the “wrong direction” at the content. The degree of discomfort experienced by remote audiences in these cases has been surprisingly strong.
4. **Remote audience location:** Displays of the remote audience should be placed in an optimal way for the speaker to see activity and questions at remote sites while still seeing his local audience and presentation materials. This area of video production also affects room design.

3.5 Technical Communication. Communication between site technicians is organized into several established protocols for the Coast-to-Coast Seminar series. Initially, the lead technician at the site hosting the lecture will send a communication to all site technicians giving the following information:

- venue server and room location with instructions for each client;
- presentation details and any special instructions on lecture materials;
- information on connection testing times and the availability of the hosting site;

⁵This did also necessitate providing white virtual chalk.





- “back channel” communication methods for technicians when monitoring the live seminar (usually via a chat server such as *Jabber*).



All sites attempt to schedule a “test run” during a regular meeting time the day before a C2C Seminar, to ensure that all site audio/video systems are working correctly, and that any presentation issues are identified and addressed beforehand.



During the lecture, technicians at remote sites initially give feedback to the hosting site technician about speaker audio levels, presentation material transition times, video framing, and any other feedback that the hosting site technician can adjust to improve the quality of the seminar experience. Technical staff also communicate any production issues to all other site technicians as they occur (such as a disconnect from the VNC desktop sharing software) so that other technicians know to watch for a potential problem.

3.6 Emergency Planning. Technicians involved in the Coast-to-Coast Seminar series are also encouraged to have specific plans to deal with common emergency scenarios. For instance, a good plan for sites hosting lectures is to have a procedure for dealing with battery failure on the speaker’s mike in an unobtrusive way. Sites are also encouraged to be practiced in interventions such as blocking audio feedback loops by immediately muting all sites other than the presenting site. For “worst case scenario” situations, sites are encouraged to have a backup communication method for both the technicians and the audience. This can be as simple as having phone contact information for the hosting site and a speaker phone available.

For instance, in the event of unknown audio interference during the speaker’s lecture, the affected site(s) would immediately mute all audio other than the speaker’s, and send a note to other site technicians describing their problem. Other sites would check their audio output to make sure that they were not inadvertently broadcasting the interference. If one site was identified as the cause of the interference, the other sites would block all audio from the source of the problem until that site’s technician assured them that it was safe to resume normal communications.

An additional communication area for the hosting site’s technician to consider is how to communicate with the speaker if there are technical issues with the presentation/equipment, or the speaker forgets how to navigate a portion of the interface mid-way through the lecture. As with other emergency scenarios, communication with the speaker in emergency situations should be focused on minimizing or resolving the technical issue as quickly as possible, and eliminating the distraction from the lecture for the audience.









4 Lessons Learned

4.1 Technical Considerations and Recommendations.

§ Audio.

- Audio clarity is the most critical component of a successful seminar series, and should be given top priority when selecting audio equipment and designing room layout.
- Site technicians for larger sites should have a technical background in audio technologies and sounds systems as well as computers and networking.
- Site technicians should always have full access to the audio system and software control panels for audio throughout the seminar. This should be considered when designing the site and setting up client services on the site PCs.
- Hosting sites should check audio levels in both the audio control panel and via feedback from remote sites shortly after a lecture starts. It is very common for a lecturer to increase in volume from initial “mike-check” levels when they start actually presenting. This can cause the audio system to overdrive and introduce distortions in the audio at remote sites.
- A directional microphone is strongly encouraged for lecturers (whether lapel or handheld) rather than having the lecturer use boundary microphones. The audio quality of the lecture will be reduced when using boundary microphones since lecturers often turn their heads away from the microphone location.
- Lapel microphones for lecturers should be clipped centered on the presenter’s shirt, not clipped to one jacket edge. If the microphone is on a jacket edge it will have a tendency to create audio fading issues as the speaker changes the direction he is facing, and it is quite common for the jacket to fall in such a way that it covers or muffles the microphone.
- All remote audience sites should have microphones muted during presentations until designated question and answer periods. This will greatly minimize site interruptions and distractions in a distributed seminar, as the distracting sounds picked up by microphones will come from the same speaker system as the lecturer’s voice (rather than from behind or to the side). This will make it harder for audience members to filter the distracting sounds, and therefore should be minimized.





- Any sound errors introduced by a site should be immediately isolated (site muted) until they can be resolved behind the scenes or after the talk. Remote technicians have the capability to mute any individual site from their audio control panel, and this should be their first reaction to an unknown audio issue.
- When using *Access Grid*, the Linux client tends to be more stable for audio services than Windows clients.

§ Video.



- Windows-based video capture devices are often limited to one camera per PC system, due to drivers. Care should be use when selecting video capture devices for Windows systems if you need more than one video camera per system.
- H.261 video is a standard codec to allow video compatibility across platforms (Windows, Linux, Mac) and clients (*Access Grid/InSORS*)
- DV (digital video) recorders with *Firewire* (1394) output are recommended for sites purchasing new cameras. This will allow those sites to start experimenting with high-bandwidth/high quality video streams when using extensions to *Access Grid*, such as *Ultra Grid*.



§ Presentation.

- When using VNC for presentation, it is recommended that the VNC server be set to disallow remote mouse and keyboard control. This will keep a remote site from inadvertently stealing the focus from the lecturer when adjusting a remote display.
- If using a VNC reflector, all VNC clients must be set to “shared”. Clients not set to “shared” will experience disconnects immediately after a connection occurs.
- To ensure that all audiences can see pointer events during a presentation, presenting sites are encouraged to purchase a remote pointer mouse and instruct lecturers in its use and the necessity of using on-screen pointers during a distributed lecture.

5 General Advice/Recommendations & How to Join!

The Coast-to-Coast Seminar has provided stimulus for related uses of the technology, as we had hoped. These have included: distributed thesis defenses in which the examining committee is in more than one geographical





location (for example, Thailand and Canada); shared planning meetings for academic projects; joint seminars (e.g., Halifax/Ottawa/Brisbane for computer-assisted architecture). Many variations are possible - for example, during the 2006 Analysis Days⁶ held at Dalhousie in January, Peter Borwein gave a plenary talk from his office at Simon Fraser University in British Columbia into the main auditorium at the Faculty of Computer Science at Dalhousie. Moreover, Jim Zhu participated in the entire two-day event from the Mathematics Room at Western Michigan University in Kalamazoo, while at the last minute Gabor Pataki gave his presentation from University of North Carolina after he discovered he could not get a visa in time.

The C2C Seminar currently focusses on research at the four Western provinces (British Columbia, Alberta, Saskatchewan, and Manitoba) and the four Eastern provinces (New Brunswick, Nova Scotia, Prince Edward Island, and Newfoundland and Labrador). The current absence of Ontario and Quebec is not intentional, and has arisen in part due to the nature of the presence of the Access Grid technology on the WestGrid and ACEnet networks. As similar networks are set up in Ontario and Quebec, we look forward to participation from universities from these two provinces⁷.

Although the discussion of the C2C seminar has highlighted research collaboration within Canada, the presentation environments on both the east and west coasts are conveniently located for international collaboration as well. Figure 5.9 shows the time-differences from D-Drive to various international locations, including lines demarcating locations within a five-hour time-difference, which is easily sufficient to organize convenient joint meetings that will occur late-afternoon in one location and mid-morning in the other. In the other direction, even the 6-hour time difference from Sydney to Vancouver and the 7-hour difference from Tokyo to Vancouver are possible for remote collaboration within the same working day at both sites.

The D-Drive facilities have proved equally useful for local collaboration, such as for writing this paper, teaching local classes interactively, and for mathematical outreach to schoolchildren, to academics from other disciplines, to decision-makers, and to the general public.⁸

If you are interested in joining or attending our regular Coast-to-Coast Seminar Series please contact:

⁶<http://www.aarms.math.ca/events/atlantic>

⁷The Coast-to-Coast seminar is also not intentionally limited to Canadian universities. We welcome participation from American universities as well, provided that their participation is technically and logistically feasible.

⁸All such uses are only as good as their weakest link, including such mundane technologies as the one controlling the security doors at the D-Drive facility, which have occasionally malfunctioned before seminars.









Figure 5.9. Time differences from D-Drive to various international cities (image courtesy of Andrew Shouldice).

- David Langstroth, D-Drive Administrator, d11@cs.dal.ca
- Veselin Jungic, Associate Director, Research - IRMACS, vjungic@sfu.ca

6 Acknowledgments

Special thanks to the following organizations for their generous assistance and involvement in helping to build the Coast-to-Coast seminar: ACEnet, WestGrid, CFI, CANARIE, IRMACS, D-Drive, the Canada Research Chair programme, NSERC, the Killam Trusts, SmartTech, all of the participating universities (as of November 2006: UBC, University of Calgary, University of Alberta, Simon Fraser University, University of Lethbridge, University of Saskatchewan, St. Francis Xavier University, Acadia University, Memorial University of Newfoundland, and Dalhousie University), and to all of the speakers for the Coast-to-Coast seminar series.



Bibliography

- [1] *Access grid community site*, <http://www.accessgrid.org/>.
- [2] *Argonne national labs access grid site*, <http://www-unix.mcs.anl.gov/fl/research/accessgrid/>.
- [3] *Westgrid access grid information*, <http://www.westgrid.ca/collabvis/accessgrid/accessgrid.php>.
- [4] L. Childers, T. Disz, R. Olson, M. Papka, R. Stevens, and T. Udeshi, *Access grid: Immersive group-to-group collaborative visualization*, Proceedings of the 4th International Immersive Projection Technology Workshop, 2000.
- [5] Alan Greenberg and Russ Colbert, *Navigating the sea of research on video conferencing-based distance education*, http://www.polycom.com/common/pw_cmp_updateDocKeywords/0,1687,2898,00.pdf, 2004.
- [6] Dave Jacqué, *Access grid aids sars patients*, <http://www.anl.gov/Media.Center/logos21-2/grid.htm>.
- [7] Wenjun Liu and Rick Stevens, *The access grid as knowledge technology*, Workshop on Advanced Collaborative Environments, 2004, <http://www-unix.mcs.anl.gov/fl/flevents/wace/wace2004/papers/paper-liu.pdf>.
- [8] Jennifer Mears, *Grid technology helps fight sars*, <http://www.networkworld.com/news/2003/0526sargrid.html>.

