White Paper for the Virtual School of Computational Science & Engineering Planning Meeting, organized by the Great Lakes Consortium for Petascale Computation, July 1-2, 2008, Univ. of Michigan

Global Cyber-Commons: Collaborative, Instructional Environments for the Virtual School of Computational Science and Engineering

Jason Leigh, Maxine Brown , Andrew Johnson, Luc Renambot,
Steve Jones, Tom Moher
University of Illinois at Chicago, Electronic Visualization Laboratory

June 27, 2008

Global Cyber-Commons: Collaborative, Instructional Cyberinfrastructure Facilities

The University of Illinois at Chicago (UIC) Electronic Visualization Laboratory (EVL) and several of its collaborators are *already* putting in place a Global Cyber-Commons, a networked set of instructional spaces with access to the same cyberinfrastructure (CI) based resources as research laboratories, as well as additional synchronous and asynchronous tools, called Cyber-Mashups, to enable and encourage computer-science and computational-science graduate and undergraduate student use. Participating faculty use Cyber-Commons to teach CI-based classes, and students use CI to implement class projects.

From EVL's discussions with the research community, it is important – albeit necessary – for instructional laboratories to have the same CI infrastructure as research laboratories – to educate the next-generation workforce, in the workplace of the future. This echoes the goals of the NSF report *Cyberinfrastructure Vision for 21st Century Discovery* [CIC07], in which Arden Bement states, "At the heart of the CI vision is the development of a cultural community that supports peer-to-peer collaboration and new modes of education based upon broad and open access to leadership computing; data and information resources; online instruments and observatories; and visualization and collaboration services. CI enables distributed knowledge communities that collaborate and communicate across disciplines, distances and cultures. These research and education communities extend beyond traditional brick-and-mortar facilities, becoming virtual organizations that transcend geographic and institutional boundaries."

Today's Grand Challenge problems – from biodiversity, to global climate change, to health issues, to the origin of life itself – can only be solved by collaborations among virtual organizations of international teams of scientists and engineers. In order for American students to remain competitive in the global workforce, they need to learn to work in trans-disciplinary, international teams, and to understand how to leverage high-performance computing and communications (HPCC) technologies to work over distance.

In addition, to remain competitive in the global workforce, American students need to learn to think and solve problems on a grand scale. For example, Google can find students anywhere in the world who can sort a few dozen database entries, but what the company needs are students who can develop algorithms to sort trillions of database entries. Google needs students who know how to perform numerical simulations, not on a single computer but on multiple cores, multiple CPUs, and multiple computers. These are things that are currently not taught in core undergraduate CS programs in the United States.

We believe that Cyber-Commons [Leigh08] is a blueprint for enabling faculty and students to tackle both Grand Challenge problems and HPCC problems by serving as the *instrument* for interacting with, and interpreting, the data generated by NSF's large-scale computing facilities, such as the TeraGrid, the upcoming Petascale facility, and eventually exascale machines. We also believe that Cyber-Commons is a blueprint for facilitating collaboration on a global scale.

The NSF 2007 report *Cyberinfrastructure Vision for 21st Century Discovery* states that CI moves us beyond the old-school model of teachers/students and classrooms/labs to enable powerful opportunities to collaborate, to model and visualize complex concepts, to create and discover scientific and educational resources, to assess learning gains, and to personalize learning environments [Ainsworth05]. Today, most students are not prepared for work in the globally distributed environments that characterize contemporary scientific research and development.

Cyber-Commons is a model for transforming the historical CS instructional facilities that are often filled with aging computer terminals, into HPCC environments equipped to teach students to solve emerging technological challenges – such as scalability, security, ubiquity, and interoperability. Cyber-Commons also enables remote collaboration for solving Grand Challenge problems, enabling students to learn about the interaction, negotiation, and management hurdles associated with the absence of face-to-face meetings, time-zone differences, language barriers, and cultural differences among distributed team members. Given UIC's 16-year experience in tele-immersion, telepresence, amplified collaboration environments, distance teaching, and international networking [Leigh97, Leigh98, Leigh99a, Leigh99b, Leigh00, Leigh03, Leigh06, Park01], we feel we can bring much understanding to this area.

Whereas Cyber-Commons provides infrastructure in which to develop new methodologies for tele-mediated education, it alone is not enough. There needs to be ways to gradually evolve the CS curriculum that is responsive to the needs of the global marketplace. *This is the mission of the Virtual School of Computational Science and Engineering – to address the unique opportunities and challenges associated with petascale computing and petascale-enabled science and engineering, to create courses, curricula and certificate programs tailored to the educational needs of 21st Century scientists and engineers, and to further facilitate the education of students in computational science and engineering through workshops, conferences, summer schools, and seminars.*

Tools for Global Cyber-Commons

We advocate the need for instructional laboratories to have the same CI as research laboratories, but with new sets of tools to encourage and support the use of these spaces in both formal and, especially, informal learning settings by computer science and computational science students, both graduate and undergraduate. EVL is interested in designing, developing, and evaluating these new sets of tools to make these spaces effective, efficient, replicable, and scalable.

High-definition (HD) conferencing alone is not enough to support the full range of activities that occurs during distance collaboration. A *mashup* is a web application that combines data from multiple sources into an integrated experience. This is one of the most exciting outcomes of the Web, as it enables people to use online databases in ways not envisioned by the original data providers. A good example of a visual mashup can be found at <www.housingmaps.com>, where housing information is mashed onto Google Maps. Another is <www.scipionus.com>, which enables users to post hurricane information; it was used in the aftermath of Katrina to post information on the condition of homes and neighborhoods in New Orleans and became a primary source of information for many displaced residents. Development tools are becoming widely available to facilitate the construction of mashups, such as Yahoo Pipes, Google Mashup Editor, and Microsoft Popfly.

A *Cyber-Mashup* takes this idea one step further. It can be thought of as the integration of an electronic journal or notebook, a conference poster, an interactive instrument panel (such as Apple's Dashboard), and a visualization environment (such as Google Earth). Cyber-Mashups will not only enable the integration of data from public databases on the Internet, but will enable access to NSF's Major Research Equipment & Facility Construction (MREFC) data repositories, including EarthScope, Network for Earthquake Engineering Simulation (NEES), Atacama Large Millimeter Array (ALMA), Large Hadron Collider (LHC), National Ecological Observatory Network (NEON), the Ocean Observatories Initiative (OOI), as well as the simulations running on NSF's TeraGrid and future Petascale facility.

Cyber-Mashup toolkits are analogous to science kits that many children in the US grew up with. Cyber-

Mashup kits are software environments that can run on laptops or can take advantage of high-resolution displays and high-performance networks that are in the Cyber-Commons. Cyber-Mashup toolkits will enable students to access data and create visualizations using CI, and will contain the tools to keep a scientific notebook, and to present their results in class and on a Cyber-Mashup server. In essence, they allow students to *do what scientists do,* except in an encapsulated form.

Hence, Cyber-Mashups can become real-time monitoring and analysis tools for day-to-day research activities. Cyber-Mashups can contain advanced visualization plug-ins to provide data visualization capabilities beyond those provided by Google Earth and others. As an electronic journal, students and researchers can use Cyber-Mashups to collect and organize data during the course of scientific investigations. This can become the persistent outcome of the work they conduct in the Cyber-Commons. Text, images, graphs, audio, video, and visualizations can be captured within a Cyber-Mashup. The journal could operate from a laptop or take advantage of high-resolution tiled display walls. Journal entries could be organized and printed as high-resolution paper posters for display at professional conferences, or shared on the web for communities of viewers to critique and/or extend.

One outcome of the OptIPuter project is EVL's SAGE (Scalable Adaptive Graphics Environment) middleware that simultaneously enables human-to-human communication and data-sharing communication on variable-sized tiled displays connected via optical networks [SAGE]. SAGE enables the real-time streaming of extremely high-resolution content − such as ultra-high-resolution 2D and 3D computer graphics from remote rendering and compute clusters and storage devices, as well as high-definition video camera output − to tiled display walls of variable size over high-speed networks. SAGE serves as a window manager; allowing users to move, resize, and overlap windows as easily as on standard desktop computers. SAGE is just one middleware system that can be used to develop plug-ins that provide access to NSF cyber-observatory data repositories and to display the information on high-resolution tiled display environments. Note, too, that SAGE has the ability to perform low-latency per-frame image compression, enabling HD video to be streamed with as little as 100Mbps.

Building Cyber-Commons Spaces

Left: EVL's high-end Cyber-Commons contains the 100-Megapixel LambdaVision tiled display. Faculty and students use SAGE to *push* their laptop screens onto the wall for weekly research meetings.
Right: UIC computer science students collaborate in UIC Cyber-Commons.

Cyber-Commons is a physical environment that is designed to be both economical and easy to deploy, enabling faculty and students to interact with colleagues over distance and to access CI resources, such as cyber-data repositories and supercomputing facilities, via high-speed networks. During class time, a room equipped with a large tiled-display wall, such as LambdaVision, can be used for media-rich HD lectures[1].

[1] It has been our experience that HD provides enough resolution to display a lecturer and either a Powerpoint screen or whiteboard on which the lecturer is writing at the same time, making the interaction with remote audiences more engaging and natural. This was not previously possible with standard-definition video conferencing – lecturers would download Powerpoints prior to a class and assume students could follow, and

Outside of class, a more compact Cyber-Commons system (using a few high-resolution monitors and a single computer), such as developed by the UIC CS department, can be used to support small-group activities. The UIC CS Cyber-Commons space is equipped and networked to run HD conferencing and cabled to allow students to plug personal laptops into HD displays [EVL08]. We envision virtual teams of students working in their respective Cyber-Commons using HD conferencing to not only see and talk with distant collaborators, but to also point at physical objects, such as their laptop screens, whiteboards, mechanical parts, geological rock samples, etc.

The components of the Cyber-Commons can be altered and enhanced as needed; increasing network bandwidth and the number of HD displays can enable more participants and more interaction with high-resolution scientific images stored in distributed databases. In fact, EVL has been experimenting with several new layout designs for media-rich environments that support traditional lecture mode, full-class discussions, small-group discussions, and localized audio/video setups (from laptops, computers or game consoles) for private groups.

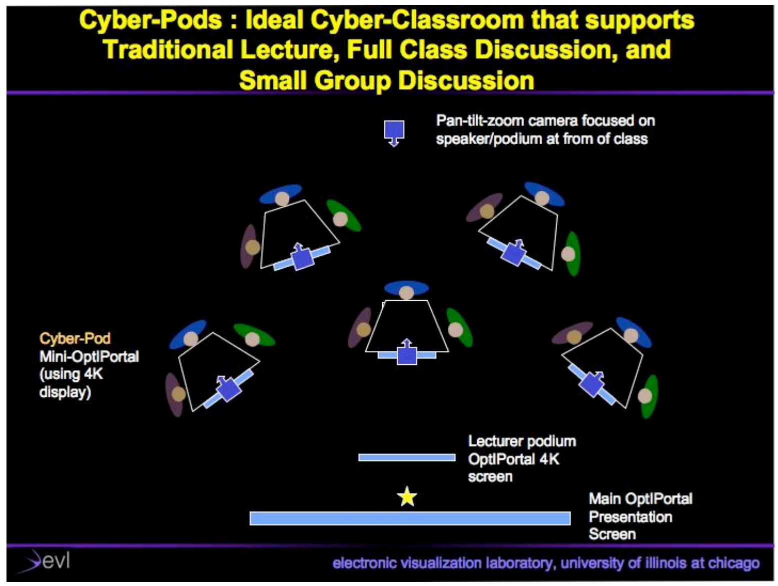

EVL is designing a *Cyber-Pod,* which is a single user of small-group collaborative system that sits in a scientist's office or in a shared lab or meeting space; the Cyber-Pod consists of a 4K LCD panel (or a 2x2 tiled HD display wall, plus HD conferencing equipment. Alternatively, one could have a room of many Cyber-Pods, intended for large-groups collaborations and/or a cyber-classroom.

- Pods facilitate small-group collaboration in team-based projects with local team members and remote pods; also facilitate audio localization for small group conversations.
- SAGE distributes video to pods and podium, which can be toggled on/off.
- Rear-classroom camera points at lecturer podium.
- Lecturer has podium screen to look at remote video as well as to control large classroom screen.
- Large classroom screen is a totally shared wall. Lecturer can toggle who has permission to post to the screen.
- All content (wireless laptops, cameras, remote apps) are considered sources that can be subscribed to by any number of pods, podium and wall screens.
- Challenge is making the interaction effortless.
- Overall scheme should be able to support the full variety of collaborations; i.e., lecture centered, multi-lecturer, discussion-centered, small group discussions.

Cyber-Commons spaces can scale. They can begin as a low-cost investment for immediate participation, but can be expanded to leverage the Petascale facility when it comes online. A low-cost system can be something on the order of what EVL and its collaborator, Louisiana State University (LSU), are putting in

whiteboard notations were only possible if a camera was pointed away from the speaker and towards a section of the whiteboard. HD is a necessary component of Cyber-Commons, enabling students to see and talk with distant collaborators, but by itself, it is not enough. Cyber-Commons provides a meeting and collaboration space based on advanced CI technologies that individual students cannot afford on their own, but university departments can.

place. Each site is building a Cyber-Commons space with one LifeSize Express conferencing system (~$7000), plus a projector or LCD display and audio system. A large-scale system can come in several flavors – a LambdaVision-type tiled-display wall, a personal Cyber-Pod, or a room of many Cyber-Pods. All these environments can be connected over advanced national and international dedicated networks made available by National LambdaRail (NLR), the Internet2 Dynamic Circuit Network (Internet2-DCN), or the Global Lambda Integrated Facility (GLIF); they can also be connected to slower-speed, best-effort routed networks, such as Internet2.

Using Cyber-Commons to Teach

There has been considerable work involving distance teaching over many generations of conferencing and broadcast technologies. The AccessGrid community is a good example of one. While introducing students to HD conferencing in a seamlessly integrated Cyber-Commons environment will pose some technical challenges, we believe that the use of, and the study of the use of CI among young users, will yield a greater return on investment if it is focused on out-of-class usage (i.e., the time when students must get together to do the actual work for their classes) rather than simply in-class usage.

For two semesters now, UIC/EVL professor Jason Leigh has taught a videogame design course at UIC and remotely at LSU using HD video streaming over the 10 Gbps Louisiana Optical Network Initiative (LONI) between Chicago and Baton Rouge [EVL07].

Left: EVL students listen to instructor Jason Leigh in this undergraduate/graduate videogame design class. The HD videoconference system provides feedback to Leigh, showing him what students at LSU see (so he can make sure he and/or his PPTs are properly displayed) as well as letting him see and hear LSU students in the remote class.
Right: A picture of the LSU class environment.

At LSU, Leigh's high-resolution image filled a large HD projection display equipped with a two-way audio channel, where students, overseen by LSU associate professor of computer science Gabrielle Allen and professor of composition and digital music Stephen Beck, followed the same syllabus and course requirements as their UIC classmates. At UIC, Leigh and his students saw and interacted with LSU students.

Leigh accepted LSU's invitation to teach his class remotely, as he saw an opportunity to understand how to offer a distributed course where classroom activity and content are extremely dynamic and require the use of high-resolution media. Leigh's videogame design course involved creative teamwork, high-resolution graphics, guest lectures from industry experts, and interactivity. HD captured the subtle nuances of interaction with high-resolution graphics over distance, especially when the interaction requires use of physical devices such as a special game controller.

Currently, student teams use online discussion groups, though future classes will experiment with more dynamic environments; LSU and UIC CS department are building low-cost Cyber-Commons' spaces. These Cyber-Commons will link students in virtual space much like many popular social networking sites, but with advanced technologies that allow virtual presence and interactive tools that foster high-quality collaboration and projects over high-speed networks.

Today's schools are collaborating in order to compete – to export and/or import renowned talent from other schools to make their students more competitive when searching for jobs. LSU sought out Leigh, a visualization and video gaming expert, to teach his game-oriented graphics class, to support an initiative for a new campus-wide curricula in digital media.

Last spring, LSU professor Thomas Sterling, an expert on high-performance (HPC) computing, taught the first HD video-broadcast, distributed HPC course both locally and to other schools in the US, lecturing students in Louisiana, Arkansas, North Carolina and the Czech Republic. The HD-enhanced streaming provided a highly engaging experience where students could easily listen to the professor's lectures and clearly see supplemental materials being broadcast.

Leigh's game class transmission goes out from a classroom at EVL over I-WIRE, a dedicated 10 Gbps network that connects universities and research institutions within the state of Illinois. I-WIRE connects EVL to the StarLight optical exchange in Chicago, where EVL connects to LONI, which is deployed on the NLR infrastructure, as well as to other regional, national and international networks. Network access and HD capability at LSU is provided by the LSU Center for Computation & Technology.

Acknowledgments

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago specializes in the design and development of high-resolution visualization and virtual-reality display systems, collaboration software for use on multi-gigabit networks, and advanced networking infrastructure. This material is based upon work supported by the National Science Foundation (NSF), awards CNS-0420477, CNS-0703916, OCI-0441094, OCI-0225642, OCE-0602117 and DRL-0426328; and, the NASA ASTEP Program, award NNX07AM88G. EVL also receives funding from the State of Illinois, Adler Planetarium, Science Museum of Minnesota, Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC), Sharp Laboratories of America, and Pacific Interface on behalf of NTT Network Innovation Laboratories in Japan. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies and companies.

References

[Ainsworth05] Shaaron Ainsworth, Margaret Honey, W. Lewis Johnson, Kenneth Koedinger, Brandon Muramatsu, Roy Pea, Mimi Recker, and Stephen Weimar, Cyberinfrastructure for Education and Learning for the Future (CELF), Computing Research Association, 2005 www.cra.org/reports/cyberinfrastructure.pdf

[CIC07] Cyberinfrastructure Council (CIC), Cyberinfrastructure Vision for 21st Century Discovery, National Science Foundation (NSF 07-28), March 2007 www.nsf.gov/pubs/2007/nsf0728/

[EVL08] High-Tech, Networked Meeting Space Links CS Students to Project Collaborators," April 14, 2008, www.evl.uic.edu/core.php?mod=4&type=4&indi=564

[EVL07] "Broadcasting Live from Chicago in High-Definition, It's Videogame Design Class!" November 7, 2007, www.evl.uic.edu/core.php?mod=4&type=4&indi=532

[Leigh08] J. Leigh, M. D. Brown, Cyber-Commons: Merging Real and Virtual Worlds, Communications of the ACM, vol. 51, no. 1 01/01/2008 - 01/31/2008, pp. 82-85.

[Leigh06] J. Leigh, L. Renambot, A. Johnson, B. Jeong, R. Jagodic, N. Schwarz, D. Svistula, R. Singh, J. Aguilera, X. Wang, V. Vishwanath, B. Lopez, D.J. Sandin, T. Peterka, J.

Girado, R. Kooima, J. Ge, L. Long, A. Verlo, T.A. DeFanti, M. Brown, D. Cox, R. Patterson, P. Dorn, P. Wefel, S. Levy, J. Talandis, J. Reitzer, T. Prudhomme, T. Coffin, B. Davis, P. Wielinga, B. Stolk, G.B. Koo, J.Y. Kim, S.W. Han, J.W. Kim, B. Corrie, T. Zimmerman, P. Boulanger, M. Garcia, "The Global Lambda Visualization Facility: An International Ultra-High-Definition Wide-Area Visualization Collaboratory," Future Generation Computer Systems, Volume 22, Issue 8, Elsevier, October 2006, pp. 964-971.

[Leigh03] Leigh, J., Johnson, A., Park, K.S., Renambot, L. "Impact of Display-rich Environments for Enhancing Task Parallelism and Group Awareness in Advanced Collaborative Environments," Proceedings of The Global Grid Forum, GGF8 : Workshop on Advanced Collaboration Environments, Seattle, WA, 06/24/2003 - 06/27/2003.

[Leigh00] Leigh, J., Rawlings, M., Girado, J., Dawe, G., Fang, R., Verlo, A., Khan, M., Cruz, A., Plepys, D., Sandin, D., DeFanti, T., "AccessBot: An Enabling Technology for Telepresence," CDROM Proceedings of INET 2000, The 10th Annual Internet Society Conference, Yokahama, Japan, 07/18/2000 - 07/21/2000.

[Leigh99] Jason Leigh, Andrew Johnson, Tom DeFanti, Maxine Brown, Mohammed Ali, Stuart Bailey, Andy Banerjee, Pat Banerjee, Jim Chen, Kevin Curry, Jim Curtis, Fred Dech, Brian Dodds, Ian Foster, Sara Fraser, Kartik Ganeshan, Dennis Glen, Robert Grossman, Randy Heiland, John Hicks, Alan Hudson, Tomoko Imai, Mohammed Khan, Abhinav Kapoor, Robert Kenyon, John Kelso, Ron Kriz, Cathy Lascara, Xiaoyan Liu, Yalu Lin, Theodore Mason, Alan Millman, Kukimoto Nobuyuki, Kyoung Park, Bill Parod, Paul Rajlich, Mary Rasmussen, Maggie Rawlings, Daniel Robertson, Samroeng Thongrong, Robert Stein, Kent Swartz, Steve Tuecke, Harlan Wallach, Hong Yee Wong, Glen Wheless, "A Review of Tele-Immersive Applications in the CAVE Research Network," Proceedings of IEEE VR'99, Houston TX, March 13-17, 1999, pp. 180-187.

[Leigh99b] J. Leigh, A. Johnson, M. Brown, D. Sandin, T. DeFanti, "Tele-Immersion: Collaborative Visualization in Immersive Environments," IEEE Computer, December 1999, pp. 66-73.

[Leigh98] J. Leigh, K. Park, R.V. Kenyon, A.E. Johnson, T.A. DeFanti, H. Wong, "Preliminary STAR TAP Tele-Immersion Experiments between Chicago and Singapore," Proceedings of HPCAsia '98, Singapore, September 22-25, 1998, pp. 687-693.

[Leigh97] Leigh, J., DeFanti, T., Johnson, A., Brown, M., Sandin, D. "Global Tele-Immersion: Better than Being There," Proceedings of 7th International Conference on Artificial Reality and Tele-Existence, Tokyo, Japan, 12/03/1997 - 12/05/1997

[Park01] Park, K., Leigh, J., Johnson, A., Carter, B., Brody, J., Sosnoski, J. "Distance Learning Classroom Using Virtual Harlem," In the Proceedings of VSMM 2001, Berkeley, CA, 10/25/2001 - 10/27/2001.

[SAGE] www.evl.uic.edu/cavern/sage/community.php