Response to the OptIPuter Year 2 Site Visit Report

NSF ITR Cooperative Agreement SCI-0225642

Larry Smarr, Principal Investigator
Calit2, University of California, San Diego
smarr@cs.ucsd.edu

SUBMITTED February 16, 2005

This is the OptIPuter team's response to the NSF Project Review Committee's report following its visit June 30, 2004. The section headers and bold italicized text represent excerpts from the Review's report, which we felt warranted further discussion and/or response by the OptIPuter team.

Executive Summary

The assessment of the site review committee is that the OptIPuter project team is making excellent progress in its efforts to develop, deploy, and promote a new architecture for the grid and new information technology facilities for e-science. To be fully successful, a project of this magnitude must produce fundamental scholarly contributions and do a better job of identifying and refining the fundamental research contributions present in their work, positioning their work relative to the state-of-the-art in the various subfields, and targeting publications to the top-tier forums, particularly in the networking subfields.

Overall, the OptIPuter project received a very positive review. The Review Committee recognized the high level of technical skill, collaboration, and management, and the importance of the research domains being investigated. This executive summary tempers the critique of research at the end of the Report, which emphasizes publication placement and better articulation of research advances. This is excellent guidance for the next Site Visit, as well as for our next Annual Report.

This is essentially the same feedback we received from last year's Site Visit committee, only much more clearly and concisely expressed. One way to better articulate our activities would be to have each OptIPuter team address its contributions relative to a specific subfield; another way would be to structure the next Site Visit to cover specific "topic areas" and present deeper highlights in those areas rather than try to cover everything.

While the design and development of the OptIPuter architecture is our primary goal and contribution, fundamental scholarly contributions in subsidiary areas are required in order for the project to be truly successful. We will:

- Encourage "distilling out" nuggets of domain-specific contributions, particularly networking
- Organize demonstrations around new capabilities, and specifically highlight those capabilities/contributions

We could create a tighter organization around lambda management and control protocol, and around transport protocol development efforts, either by reallocating our funding or by not presenting those activities that are perceived as not contributing toward our goals. However, while network hardware, including DWDM, optical fiber, and high-speed routers and switches provide raw capabilities, the OptIPuter project aims at the re-optimization of the entire Grid stack of software abstractions, so is focused on hardware and software integration across the physical, transport, middleware and application levels, rather than on one specific level and/or domain.

As the Site Visit Report noted in its Executive Summary: "An important milestone for year two was publication of papers in a Fall 2003 issue of *Communications of the ACM,* one of the most visible forums in advanced information technology and computer science. The OptIPuter project was also "all over Supercomputing 2003," [and subsequently SC 2004] with demonstrations and presentations at this leading and highly visible forum."

A common problem faced by projects of this scale is integration across subprojects. The OptIPuter review team saw good evidence of a strong coupling between visualization, networking, cluster computing, and domain scientists.

We will continue to do experiments that involve infrastructure, system software, visualization, and applications. We are also adopting the suggestion of the OptIPuter's Frontier Advisory Board that at future Site Visits or conference demonstrations, we showcase a few (~4) experiments that have deep cross-team integration and limit the non-integrative demos in number (i.e., 2-3 maximum) to balance the perception. However, some non-integrative demonstrations may be required to illuminate fundamental research contributions in a particular subfield or domain.

Criterion 2 / Broader Impact. The site review team saw evidence of outreach, integration of research and education, and creative efforts to broaden and accelerate the impact from the project. The site review team hopes that outreach activities will expand further in later stages of the project with the arrival of Debi Kilb as outreach coordinator. The project and driving application areas offer outstanding opportunities for interactive eduware, as proposed in the original proposal.

We would like guidance on how, if possible, we can "expand" outreach further. We currently have good coverage with respect to grade level (i.e., K-5, 6-8, 9-12, 13-16 are all represented). This leaves content development and dissemination. We will continue with our current level of content development (both neuroscience and geoscience), as well as discuss regional or national distribution. However, these areas do have budget ramifications.

Note that the OptIPuter's outreach efforts extend beyond school children and teachers. The OptIPuter has also successfully done outreach to other agencies and international institutions, transferring research results to USGS, NASA Goddard Space Flight Center, the Joint Oceanographic Institutions, and partners of BIRN and BIRN2.

I. Research Plan

The project is proceeding mostly according to plan, with significant progress to date. Their work diverges from the proposal in some respects that are significant, but we view these as minor relative to the overall scope of the project. For example, the original proposal detailed an integration of the OptIPuter backplane with Infiniband cluster interconnects, but most of the activities apparently use Ethernet in the clusters. We expect that the research team will revisit this question in later deployment phases based on experience with initial Ethernet-based prototypes.

We have seen good evidence of adaptation to changing circumstances in the project to date, and we have faith in the PI team to make the right choices to realize the OptIPuter vision without being unduly constrained by their earlier thinking.

Infiniband is still an active topic of discussion among OptIPuter partners. The OptIPuter Infrastructure Team is actively researching this area, keeping its options open since, as of January 2005, there is no clear answer, and members do not want to spend their equipment money on mass replication of a poor choice.

It is noteworthy that the impact of OptIPuter on computer science has already led others to propose and receive funding to study SAN-WAN performance and architecture issues[1]. These research questions are at the heart of the proper role for Infiniband, and we are pleased to see a larger community engaging in exploration of these questions.

II. Research Accomplishments

Research accomplishments in optical networking integration, system software integration and deployment, Grid architecture and resource/data management, network protocols/transport, optical networking architecture and visualization were enumerated.

Network protocols and transport. The work on RBUDP and SABUL is also interesting in the sense that real application scenarios in the OptIPuter project motivate them. However, more thorough, rigorous, and scientific approach is needed to understand the fundamental tradeoffs and benefits of these approaches as compared to the state-of-art.

It should be noted that other grants are funding the research and development of protocols SABUL/UDT[2], RBUDP[3]

---

[1] At a recent CISE PI meeting, Andrew Chien discovered a new award to OSU that is focusing on this exact question and has framed it with an OptIPuter context.

[2] SABUL/UDT received major funding from NSF award ANI-0129609, called "Tera Mining: A Testbed for Distributed Data Mining over High Performance SONET and Lambda Networks," Robert Grossman, PI.

[3] RBUDP received major funding from NSF award ANI-0129527, called "Quanta," Jason Leigh, PI.

and LambdaStream; OptIPuter funding is being used to tailor these protocols to its bioscience and geoscience application drivers and integrate them into the OptIPuter's networking and cluster infrastructure.

For RBUDP, there is a published paper that describes the formulas used to reliably predict performance[4]. For SABUL/UDT, there is a published paper that provides a theoretical approach to understanding UDP-based protocols[5]. The more recently developed GTP protocol has produced publications that both document the performance of GTP and compare it to a set of rate-based protocols, including SABUL and RBUDP[6,7]. Papers comparing SABUL/UDT to other protocols have also been published[8,9,10]. In addition, the Explicit Control Protocol (XCP) results were presented in a CACM survey article[11] and at broadly attended networking forums, including the Workshops on Protocols for Fast Long-Distance Networks (PFLDnet) 2004 and 2005. XCP was "jump started by the OptIPuter project at its onset to address the fundamental problem of achieving superior transport performance in networks with a large bandwidth-delay product, and was subsequently funded by NSF as a separate effort; the protocol is also documented in an Internet Engineering Task Force (IETF) work-in-progress[12].

We expect that as our SABUL/UDT, RBUDP/LambdaStream and GTP efforts mature, more extensive modeling, analysis, and evaluation will emerge. These efforts will result in publications which characterize the protocol in greater detail, and as suggested by the review committee we will target a range of high-profile venues in the networking community.

An important contribution of the OptIPuter architecture is the collection of these varied protocols under a shared interface. This shared interface will enable applications to easily select and use the appropriate protocol for their needs, and facilitate a systematic comparative study which builds on Chien's work[7] and other work by OptIPuter partners[13]. This comparison, which both documents and distills the contribution of the protocol work, will accelerate its acceptance by the broader networking community.

As we stated in the Executive Summary, we could create a tighter organization around lambda management and control protocol and transport protocol development efforts, either by reallocating our funding or by not presenting those activities that are perceived as not contributing toward our goals.

---

4 For RBUDP, see: E. He, J. Alimohideen, J. Eliason, N.K. Krishnaprasad, J. Leigh, O. Yu, T.A. DeFanti, "Quanta: a toolkit for high performance data delivery over photonic networks Future Generation Computer Systems," Special Issue: iGrid 2002, Volume 19, Number 6, August 2003, Elsevier bv, The Netherlands.

5 Yunhong Gu, Xinwei Hong and Robert Grossman, "An Analysis of AIMD Algorithms with Decreasing Increases," Proceedings of GridNets 2004, IEEE Press, 2004.

6 Ryan Wu and Andrew Chien, "GTP: Group Transport Protocol for Lambda-Grids," Proceedings of the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid2004), Chicago, Illinois, April 2004.

7 Ryan Wu and Andrew Chien, "Evaluation of Rate-based Transport Protocols for Lambda-Grids," Proceedings of the Thirteenth IEEE International Symposium on High-Performance Distributed Computing, Honolulu, Hawaii, June 2004.

8 Yunhong Gu and Robert L. Grossman, "SABUL: A Transport Protocol for Grid Computing," Journal of Grid Computing, Volume 1, 2003, pp. 377-386.

9 Robert L. Grossman, Yunhong Gu, Xinwei Hong, Antony Antony, Johan Blom, Freek Dijkstra, and Cees de Laat, "Teraflows over Gigabit WANs with UDT," Journal of Future Computer Systems, Elsevier Press, 2004.

10 Robert L. Grossman, Yunhong Gu, Dave Hanley, Xinwei Hong and Parthasarathy Krishnaswamy, "Experimental Studies of Data Transport and Data Access of Earth Science Data over Networks with High Bandwidth Delay Products," Computer Networks, Volume 46, 2004, pp. 411-421.

11 Aaron Falk, Ted Faber, Joseph Bannister, Andrew Chien, Robert Grossman, Jason Leigh, "Transport protocols for high performance," Communications of the ACM, Volume 46, Number 11, November 2003, pp. 42-49.

12 draft-falk-xcp-spec-00.txt

13 Antony Antony, Johan Blom, Cees de Laat, Jason Lee, "Exploring practical limitations of TCP over TransAtlantic networks," accepted for publication into the FGCS special issue (section) on DATATAG.

Visualization. The majority of the visualization research has been concerned with issues relating to the problem of mapping very large 2D, 3D and 4D images across the OptIPuter network to very large arrays of display pixels, especially on tiled displays. The application tools JuxtaView, Vol-a-Tile, LambdaRAM and TeraVision have been developed rapidly and are providing basic solutions to visualization needs. The work done thus far should provide a solid foundation both for more application building and for research into the more fundamental issues. The Grid Visualization Utility (GVU) point-rendering project addresses the challenging problem of large 4D data (volumes + time) with an interesting combination of point-based rendering together with filtering.

The new SAGE (Scalable Adaptive Graphics Environment) architecture provides the opportunity for tackling these issues in a focused and systematic way. SAGE provides a foundation for concurrent visualization development efforts, such as Thiébaux's GVU and Cox's high-definition renderings, to integrate seamlessly on future-generation ultra-high-resolution displays. SAGE's data transport requirements address hard networking problems, such as how to provide low-cost switching at end hosts for pixel routing. Solutions for GVU may have relevance to data mining applications as well.

Is the research original and/or innovative? Many of the key activities deal with "putting the pieces together" to realize the OptIPuter vision. The key question is whether the OptIPuter work in each subfield is competitive with the best work in those subfields. The review team is convinced that the work is visible, credible, and competitive in each of the subfields, but that opportunity exists to improve in key areas. In particular, the project should seek to direct each piece of work to the community best suited to evaluate it.

The OptIPuter is a broad-based project which is making contributions to many subfields, but a major element of OptIPuter's contribution comes from the integration of knowledge across domains, which aids both understanding of the value of particular knowledge and problems in a larger context. In this sense, the sum is greater than the parts; we believe the best communities for evaluation are the applications ones, as well as computer science research groups who can evaluate the information technology research. To the extent that subfield communities are interested in any phase of our integration efforts, we are open to advice, and, moreover, we actively seek collaboration. Note strongly, though, in the OptIPuter case, each subfield piece must work in a supportable way with the whole, which means each piece needs to be mature and therefore likely to be considered somewhat of a novelty by the leading-edge researchers in that subfield.

As a specific example, partner UCI/Smyth has published papers in the "sub-area" of data mining/data analysis on spatio-temporal data mining techniques at top conferences, notably ACM SIGKDD, NIPS, UAI, as well as publishing "applications" of these techniques in the science literature, such as the 2004 Journal of Climate. Smyth's work is viewed as novel in the data mining community, primarily because it deals with large complex spatio-temporal data. While many science problems (e.g., brain imaging, earth sciences) inherently focus on data with both spatial and temporal characteristics, there are relatively few data mining methodologies that have been developed specifically for such data. The fact that OptIPuter datasets will be "streaming past" at very high rates also makes the research novel: there has been a fair bit of recent work in the data mining research community on "data mining of streams" but nothing on the scale of OptIPuter and very little with scientific data. So, we can safely say that the data mining aspects of OptIPuter research are original, almost by definition.

Partners have published papers and participated through posters and live demonstrations in several conferences, including IEEE Computer Graphics and Applications, American Geophysical Union, Geological Society of America, IEEE Visualization, IEEE Cluster Computing, HPDC, SPIE Electronic Imaging Proceedings, ACM/IEEE SC, PFLDnet, Global Internet, and the NASA Optical Network Testbeds Workshop.

Early project work on XCP has also been presented to and published through IETF.

Grid architecture and resource management. In its purest form, the OptIPuter Distributed Virtual Computer [DVC] vision embodies a new direction in grid management that is similar to the PlanetLab (e.g., SHARP) vision: grids should support dynamic allocation of bundles (or "slice") of raw physical resources sized to meet the performance demands of a particular use, with software custom-configured specifically for that use (e.g., HP UDC, Duke Cluster-on-Demand, or Stanford Collective). This vision is innovative in that it departs from the dominant vision in the grid community. However, much work remains to be done in the DVC team to go beyond the plumbing and deal with fundamental issues such as secure federation and allocation policy. These other

projects may offer ideas and technology that are applicable to the OptIPuter context, and the DVC team should consider adopting them, and if not, to fully understand and explain the architectural choices and assumptions that lead them to a different design. Similarly, the LambdaRAM team should look at research in the mid-late 1990s on network memory and prefetching (e.g., GMS and PGMS see Voelker).

The Distributed Virtual Computer (DVC) provides a number of capabilities which are distinct from the vision currently articulated for PlanetLab (though it is a moving target). First, the DVC integrates control of configurable optical networks and communication capabilities as a key element of the OptIPuter project. Resource management structures such as SHARP focus on end resources, and provide novel capabilities in being able to represent them and trade them. More recent work on SWORD allows expression of resource needs and communication constraints. However, these communication constraints are solved against a model of FIXED connectivity (a single set of attributes for dynamic properties of these attributes, such as latency, bandwidth). In contrast, the DVC provides the capability to select communication and end resources in a coordinated fashion, and produce a prescribed network configuration which can then drive underlying optical network configuration software. In this, the presentations at the Site Visit last June did not fully represent the depth of comparisons that have been made to date.

DVC goals differ from the other mentioned systems, and it may well be the case that as the DVC becomes more complete, some of these other technologies will be incorporated. It is important to note that Andrew Chien and the OptIPuter's system software team spent much of the first year of the OptIPuter project understanding how to create an integrating architecture. The success in finding a promising model, the DVC, and its progress to date, should be viewed as a major success that supports the future impact and success of the project as a whole. Note that the DVC "model of use" and architecture are likely *major* contributions of the project, as there has only been a single student working with Chien on DVC for nine months, and another MS student was recently hired to work with him.

The LambdaRAM team has extensively studied the aforementioned areas of research, including work by Voelker.

Optical networking architecture. This work is original and innovative, but it is also being pursued by industry. Nevertheless, to deploy operational systems, the team cannot wait for industry products to become available, so it is quite appropriate for them to develop their own protocols and algorithms. However, this sort of work may also be quite hard to publish in scholarly scientific outlets.

A key objective of the OptIPuter project is to create a large-scale distributed architecture that supports new capabilities for data-intensive applications. Within this architecture, optical transport is considered less as a classical network service and more as a type of distributed backplane function – subordinate to, and controlled by, other virtual system processes.

This approach requires new considerations of particularly close integration between core network processes and various edge processes, including those based on computational clusters. Whereas, there is significant industry and academic research development in all areas of optical networking, few of these activities envision the type of seamless integration between optical networking and multiple differentiated edge resources that is being explored by the OptIPuter project. For example, the Grid community has long been addressing issues related to distributed architecture, but has only recently been examining the potential of integrating resources with dynamic optical networks. In part, these activities have been motivated by the OptIPuter project.

These new architectural concepts have only recently been brought to the attention of industry, again, in part through the OptIPuter project. Certainly, as noted, OptIPuter researchers are using commercial optical products tailored to new requirements, but are also using non-commercial prototypes in order to investigate new architectural concepts and protocols not being addressed by industry.

In pursuing its research agenda related to optical networking, OptIPuter investigators work closely with other researchers who are shaping the future of advanced optical networking architecture. We are tracking new methodologies and new technologies in advanced photonics, next-generation devices (e.g., high-performance tunable lasers and amplifiers, optical subsystems, novel switch design), new methods for management and control, techniques for optical burst switching and investigations of optical packet switching. Notably, independent of the OptIPuter, team members have partnerships with industry to work on basic research and pre-commercial equipment. We also work within key standards organizations, including the IETF, the OIF, the ITU, the GGF, and the IEEE, on shaping new optical network architecture standards.

OptIPuter development is influenced by a wide spectrum of perspectives related to optical networking architecture that range from leading-edge research to standards development to prototype design and experimentation. This perspective has allowed OptIPuter researchers to distinguish between concepts that are truly innovative with a potential for imparting a major impact and those that merely recast familiar ideas or that can only provide for incremental improvement of classical architectural approaches.

In addition, OptIPuter research has been presented by invitation at major international photonics conferences. These activities enable the OptIPuter project to stay informed about research trends and commercial development.

What impacts are project accomplishments likely to have on the state-of-the art? As previously stated, the key role of the OptIPuter project is to integrate technology to demonstrate and promote the OptIPuter vision, developing missing pieces as necessary. This is a significant intellectual activity in itself, and integration of the effort with experts in driving application areas (microscopy, bioscience/BIRN, geoscience) is central to the effort and critical to its success and technology transfer. Engagement of commercial technology providers is also important. The project is on track to have several significant impacts…..[one of those impacts cited was to] provide powerful new tools for and other areas that interact with massive data and/or real-time data, e.g., homeland security.

UIC partner Grossman and UCI partner Smyth are both involved in homeland security projects, so are obvious advocates of 'tech transfer' to other data-intensive problems. In particular, Grossman has Army funding and Smyth is part of the UCSD/UCI RESCUE ITR (information technologies for crisis response). For the latter, one obvious area is situation assessment: large-scale tiled displays that allow visualization and integration of real-time (e.g., satellite, video, traffic, weather) and archival (e.g., GIS) data.

Alternatively, USGS Earth Resources Observation Systems (EROS) Data Center (EDC) became an OptIPuter partner to leverage new OptIPuter technologies to better utilize high-resolution (0.3-meter) ortho-imagery of the 133 most-populated US metropolitan areas in support of Homeland Security initiatives. The OptIPuter's large-scale visualization capabilities are currently the only potential method USGS has for interactive display and analysis of high-resolution data covering vast geographic areas.

Regarding network security research, researchers are looking at several problems with respect to TCP and IP vulnerabilities that arise exclusively in very fast networks. Given the site visit's emphasis on publications, we envision disseminating the results of these thrusts at IEEE International Conference on Network Protocols (ICNP), IEEE INFOCOM, and the RSA Security conferences, with more thorough reporting in IEEE Network, Computer Networks, and ACM SIGCOMM Computer Communication Review. The important task of technology maturation will be pursued in IETF, especially in the emerging working group on Better Than Nothing Security (BTNS), where the FastSec work has been presented.

- The heightened likelihood of a successful sequence number attack on a TCP connection when high link speeds cause the sequence number space to turn over rapidly during the attack timeframe and admit the possibility of a reset catching the connection at a particularly inopportune point. OptIPuter researcher Joe Touch of ISI will apply network security protocol IPSec in an unconventional way to affect very weak authentication that is sufficient to protect against sequence number attacks.
- Encryption can introduce a full packet's worth of latency, which warrants research into latency reduction in network security protocols (such as IPSec). Joe Touch has an approach (called FastSec) that modifies IPSec modes to support microblock ciphers that suffer negligible start-up latency during encryption.
- A third networking topical thrust is to adapt XBone virtual network overlay technology developed by Joe Touch of USC/ISI to the problems of lambda management and topology configuration. XBone employs rich, agile tunneling to build dynamic IP overlays well suited for security, resource management, traffic control, and automated management and configuration. XBone has a well-tested open-software base and is used in production systems.

With respect to technology transfer, it should be noted that although this project is not designing or developing new optical hardware, it is unique in orienting its research activities toward innovative optical networking architecture, in some cases, using methods that are not being addressed by other research projects. The orientation of this research to a future architecture, as opposed to that which is current or recently emerging, has attracted considerable interest from other research communities as well as industry. The OptIPuter is designing data-intensive services that are not being handled by typical data communication services, which address general consumer and enterprise requirements, or by most forms of emerging large-volume data-transport services.

Project's current and planned research outputs (publications, presentations, etc.)? The committee is hopeful that the project's publications will be targeted and accepted in the top-tier forums for each sub area. Some core elements of the work (e.g., optical network implementation and integration) are not publication-friendly, so it is not expected to lead to scholarly publications. This does not detract from the intellectual significance and potential impact of the activities in those areas.

We appreciate that the site visitors recognize that the OptIPuter is a systems integration project, not a theoretical research project, so we are limited in the number of scientific journals in which we can publish. However, we have had success with several mainstream publications and conference proceedings, as the list of publications in our Annual Report attest. Papers have appeared, and continue to appear, in "top-tier forums" in several sub-areas. Notably, in the sub-area of high-performance computing and communications and grid technologies, we have published in the following "top-tier" conferences over the past year (to name a few):

- SC 2003 and 2004 (best meeting in HPC and Grids)
- HPDC 2004 (the best meeting in Grids and High Performance Distributed Computing, single track, acceptance 15% this year)
- CCGrid 2003 and 2004 (strong forum, probably the third best meeting in grids after HPDC and SC)
- IEEE Visualization (top meeting for visualization)
- Workshop on Protocols for Fast Long-Distance Networks (PFLDNet) (premier workshop for this area)
- Have had posters and presentations accepted by American Geophysical Union and Geological Society of America (which together covers the entire geoscience community).
- Workshop on Optical Networking for Grid Services at the European Conference on Optical Communications in Stockholm, Sweden (invitational paper, published in proceedings).
- International Conference on Photonic Technology 2005 in Tokyo, Japan (invitational paper, published in proceedings).
- Workshop on Networks for Grid Applications (GridNets 2004), co-located with BroadNets 2004, sponsored by the IEEE and ACM among other organizations (invitational paper, published in proceedings).

We understand that there are other venues, such as SIGCOMM and INFOCOM (for networking) and OSDI/SOSP (for resource management), and we will make every effort to target them for publication to increase the communication of key results from the OptIPuter project to these communities. As the project progresses in the scale of testbeds and integration, which enable easier large-scale experimentation, we expect that the number of publications in top-tier venues will continue to grow, including the IEEE International Conference on Network Protocols (ICNP) and RSA Security conferences, as well as the periodicals IEEE Network, Computer Networks, and ACM SIGCOMM Computer Communication Review (CCR). Continued collaboration within IETF will be pursued. We are considering the possibility of a special issue on topics in optical networking for IEEE Network.

Does the whole exceed the sum of the parts? (Is the project truly multidisciplinary?)…The project is truly interdisciplinary, and strong and convincing evidence to this effect has been provided…As evidence of this collaboration the site review team saw an unusual set of integrated, professional, state-of-the art demos – a whole that most definitely exceeded the sum of the parts. The Project Manager noted that their travel funds are inadequate to sustain the current level of interaction and collaboration in the project. This may be an issue for NSF to address.

While NSF has informed us that it is highly unlikely we will receive additional travel funds this year, and need to budget accordingly in future years, face-to-face meetings have been a major factor in cultivating a strong working relationship between the discipline scientists and the computer scientists. In addition to travel, we also make extensive use of video conferencing, without which our travel budget needs would be even higher.

III. Education and Outreach

How appropriate are the project's future goals and plans with respect to educational and outreach activities?... The OptIPuter team is to be commended on their education and outreach activities. The GeoWall and GeoWall2

are an outstanding example of University-based research being disseminated into a large community… The GeoWall2 project has the potential to be providing the dissemination of scientific research to the broader public. A plan to use GeoWall2 displays to make very large ortho-images of cities available to the public is also commendable and likely to increase public awareness of urban environmental issues.

For the past year, a portable GeoWall2 has been on loan to OptIPuter affiliate partner USGS; it is used to display complex (0.3-meter resolution) maps of major US cities at homeland defense workshops and conferences.

OptIPuter partner EVL also developed the Personal GeoWall2 (a 2x2 tiled display) as a way to put tiled displays on individual researcher's desktops. This may dramatically expand outreach in much the same way that the original low-cost GeoWall (a passive stereo one-screen projection-based system) was adapted by museums and classrooms.

In addition, Joe Touch will offer the USC Computer Science CS-599 course on Virtual Overlay Networks in Spring 2006, one of the first courses of this type in academia. Touch has also designed and agreed to present tutorials on Virtual Overlay Networks at ICNP, INFOCOM, and SIGCOMM Asia. He has also been appointed IEEE lecturer on the topic. Clearly, the promise of wide outreach and broad dissemination of project innovations is strong.

For year 3 the OptIPuter group plans to continue and extend this work. A GeoWall will be installed at the Birch Aquarium, SIO. An additional project is planned for the Preuss Middle School that will greatly increase the capabilities and the range of materials that can be presented by integrating with ARC GIS.

The Birch Aquarium, part of Scripps Institution of Oceanography, has an installed GeoWall that is used for a number of programs, and is seen by thousands annually. As for the Preuss School, we have switched from focusing on ARC GIS to using OptIPuter-generated data and programs and/or freeware programs whenever possible.

IV. Project Management

Does current project organization and management actively support collaboration across sites and disciplines? The site review team saw evidence of strong, proactive, and effective central coordination through Larry Smarr and project manager Maxine Brown. PI Smarr, the co-PIs, and many of the graduate students showed familiarity with all aspects of the project…The project team has…annual all-hands meetings and regular meetings…The review panel does feel that these meetings are an important opportunity for participation by graduate students, and that their participation should be further increased to further leverage this opportunity. Continuing engagement with the advisory board is also important to the project's success.

Graduate students attend and participate in all meetings, and will continue to do so. OptIPuter outreach activities extend to elementary, middle and high schools, and undergraduate and graduate programs. Independent of, but related to OptIPuter, SIO has an extensive outreach program and holds an annual SIO Graduate Student Visualization contest, which helps attract interest in OptIPuter goals <www.siovizcenter.ucsd.edu/news_events/comp2004/index.html>.

V. Future Research Goals and Activities

Opportunities for new directions in the research that have not yet been pursued? As stated in the executive summary, the PIs have an opportunity to further enhance the impact of their work on the state-of-art within the discipline. To be fully successful, a project of this magnitude must produce fundamental scholarly contributions.

We addressed this issue above. Suffice it to say that technical contributions may well be recognized in the future, but not at the instant of first demonstration. Recognition has many forms (invited lectures, participation in global networking planning, etc.) which do not all lend themselves to journal publications.

The PIs can do a better job of identifying and refining the fundamental research contributions present in their work, positioning their work relative to the state-of-the-art in the various subfields, and targeting publications to the top-tier forums, particularly in the networking subfields.

We addressed this issue above. Note, too, that this is something all the graduate students need to do and are doing anyway as part of their thesis; that is, they survey the state of the art and compare it with their new proposed methods. We could take some of this information and present it at future Site Visits.

Network protocol and transport. The tight interaction of the transport protocol research and the overall OptIPuter system and applications is a strength and should be continued in the future. The proposed research directions on new Multi-Endpoint communication abstractions/protocols and RobuStore (high-performance service from geographically distributed storage) are promising.

It would be good to see more interaction among the multiple transport research teams within the OptIPuter project. For example, it would be interesting to compare XCP, GTP, RBUDP, and SABUL in different network environments (shared Internet, dedicated wavelength). The insights from the study should be of general interest to the networking community. To gain the maximum feedback and impact, publications resulted from the future research in the transport protocol areas should be targeted to mainstream networking research venues such as SIGCOMM, INFOCOM, and Transactions on Networking.

As indicated above, the OptIPuter team has been building an integration framework to enable more extensive and systematic comparisons, though comparison papers have already been published[7,11]. These transport protocols have been developed to address different things (Internet *versus* dedicated bandwidth, streaming video *versus* data movement, etc.), which is the basis of the evaluation and enables selection of the appropriate protocol based on application needs.

We hope to report exciting new results on Multi-Endpoint communication and Robustore in the near future.

Visualization. A major thrust in future work is to determine how best to allocate distributed computational resources to the classic visualization pipeline [Data → Filtering → Rendering → Visualization]. The goal is to create a rendering abstraction to support the flexible distribution of rendering given available resources and the work will be carried out through the construction of prototype applications. This is an extremely challenging problem and the success or lack thereof will be a major issue in determining the ultimate impact of OptIPuter-based visualization. A good test application will be the proposed development of an application in spatio-temporal data mining.

We are in complete agreement. Different sub-teams will be tackling sub-portions of the problem simultaneously. SAGE <www.evl.uic.edu/cavern/sage> is intended as the unifier for all these efforts.

Optical networking architecture. In terms of the optical networking component of the project, the focus so far has been on development of control-plane and lambda-provisioning algorithms and protocols. The team has stated future research goals and plans, but they (namely, the research component, not the development component) are not very convincing. The reviewers felt that the research team's responses to questions on research on the state of the art on lightpath provisioning models and network survivability issues were inadequate. Supplementing the current networking team (which has excellent development skills) with researchers who are well versed in optical networking research (and the associated research literature) will be helpful.

As noted, OptIPuter investigators remain current on major research projects that are influencing the advancement of optical architecture, related to methods, technologies, early prototype development and standards activities. These researchers are providing a special focus on control and management plane architecture and new methods and protocols for dynamic lambda provisioning, survivability and restoration. The majority of current academic research and industry development in this area is currently focused on centralized techniques for implementation within a single domain (intra-domain). From this perspective, inter-domain provisioning architecture is considered most appropriately designed by carrier requirements and related mutual agreements. In contrast, OptIPuter research is focused on investigating new methods for highly distributed management and control methods that can be provisioned by edge processes across domains. As indicated, these control processes are subordinate to others, so that collectively they all comprise a single integrated system.

OptIPuter optical networking researchers (notably Joe Mambretti of Northwestern University, Oliver Yu of EVL/UIC and Cees de Laat of University of Amsterdam) represent OptIPuter as active participants in forums organized by the research community, non-profit industrial developers, and standard bodies, such as the IETF, OIF, ITU, IEEE and GGF. For example, recently, a draft on Lightpath Control Protocol was developed by Northwestern researchers and submitted to the IETF. OptIPuter researchers were also co-authors of a GGF standards draft on high-performance networking for grids. In addition, we have participated in a number of workshops and conferences presenting research on this topic, including international photonics workshops, Grid networking and computing workshops, such as GridNets 2004 and CCGrid, and the International Optical Control Plane workshops organized by Gigi Karmous-Edwards of MCNC. Joe Mambretti was one of the organizers of GridNets 2004 and is currently assisting with GridNets 2005. We are also involved in activities directed at implementing these new techniques, such as the Global Lambda Integrated Facility (GLIF). Hence, OptIPuter research is shared with, and influenced by, a wider audience, including those in the research community, industry developers, participants in standards bodies and advanced optical networking designers and implementers.

A disadvantage of introducing "pure" researchers into the OptIPuter's developments is that they must speak the same language as the bioscience and geoscience people – a problem that we have encountered. It is only by working alongside the application people that one learns to "translate" between groups.

Expanding the networking research team, as stated above, would be very useful for the project. Additionally, the research team stated several times about "optimization" of various issues, e.g., when to use existing Internet infrastructure for a data transfer vs. setting up a new lambda path. Can these decision-making processes be formally compared and the optimal choice made? Another example is at what locations should a data stream be duplicated for a multicast application.

The framework defined by a DVC will allow applications simultaneous access to both shared Internet and dedicated lambda environments. As such, it will provide an excellent opportunity for well-defined scientific experiments that explore these questions. We expect that as integrated experiments involving applications become easier and more frequent over the next year or two of the project, we will be able to make comparisons that could inform an automated decision-making capability.

In addition, Joe Touch of the USC ISI Postel Center for Experimental Networking recently joined the OptIPuter project, to contribute new ideas in high-speed IP and TCP. His experience in experimental high-performance networks will supplement the superb systems skills of the current team.

An important issue of integration of packet and optical networks is the bandwidth granularity of circuits. The general consensus in industry is that wavelength may be a too coarse granularity and sub-wavelength such as SONET/SDH piles (STS) maybe more appropriate. Therefore, provisioning of sub-lambda connections (possibly using SONET/TDM transport and cross-connect technologies and latest advances in Virtual Concatenation and GFP) is an important area to pursue, both from engineering and research perspective. Right now, it was given only "lip service" to.

OptIPuter researchers agree that providing options for sub-lambda granularity is an important objective, as the OptIPuter environment should not address only requirements for high-performance large-scale data flows. Consequently, we are investigating various methods for provisioning, controlling, and managing sub-wavelength channels. We have strongly encouraged the development of proposed architectures for digital framing being developed by the ITU (e.g., G.709), that will assist in providing this type of flexibility. We are also tracking related standards, Link Capacity Adjustment Scheme (e.g. G.7042/Y.1305), High Order Virtual Concatenation (ANSI T1.105, ITU G.707), Resilient Packet Ring (IEEE 802.17), and the Generic Framing Procedure (ITU G.7041, which is complementary to G.7042). In addition, the OptIPuter project is exploring a partnership with another group to create an optical testbed that supports GFP. However, these standards have received considerable industry support and they are migrating to implementation in commercial equipment, either as part of a "plan of record" or actually as part of product development.

The primary focus of OptIPuter research initiatives will be on topics that are not yet being addressed by industry. The OptIPuter is interested in sub-lambda capabilities for large-scale, deterministic data flows, with highly distributed control and management processes, as one component of a larger integrated system rather than as a form of common carrier service. Providing these capabilities is a major challenge that requires investigative research and experimentation in areas that are not yet being directly addressed by industry, because commercial service providers are not planning to implement these types of capabilities for many years.