An Inter-Regional Grid Enabled Center for High Energy Physics Research and Educational Outreach (CHEPREO) at Florida International University

in collaboration with California Institute of Technology, Florida State University and the University of Florida

Year 2 Progress Report

May 12, 2005

U.S National Science Foundation

This page intentionally left blank

TABLE OF CONTENTS

1. INTRODUCTION

Florida International University (FIU), together with partners at Florida State University (FSU), the University of Florida (UF), and the California Institute of Technology (CIT), in cooperation with the National Science Foundation, is in the midst of creating and operating an interregional Grid-enabled Center for High-Energy Physics Research and Educational Outreach (CHEPREO) at FIU. The report describes the activities accomplished in the second year (February 1, 2004 – March 30, 2005).

CHEPREO made progress in all areas, from the very general (phase I building renovations on the FIU campus, including laboratories, classrooms, and a tutoring and support center) to the very specific (doubling the number of students in modeling classes, responding to changed CMS experiment needs, or an increase in the number of CPUs in the Tier-3 data center). Of special note were the summer physics modeling workshops, research by FIU undergraduates and the STM-4 network link with Brazil.

CHEPREO also enjoyed site visits with the NSF on July 13 2004 and March 30 2005, with representatives from four NSF directorates offering advice and reviewing progress and plans. FIU institutional commitment was represented by the University President, Provost, University Vice-Presidents for Sponsored Research and Technology, and the Dean of the College of Arts and Sciences. The external advisory committee also met in July 2004 at FIU to give input on the strategies of the collaboration.

2. Progress at Florida International University

2.1 Physics Research

FIU again hosted the 3-day winter CMS HCAL meeting in March 2005. This year the sessions were held at FIU's Wolfsonian Museum in Miami Beach. About 30 colleagues from around the world attended including two FIU undergraduates, with L.Lebolo presenting work done with V.Gaultney.

2.1.1 Contributions to the CMS Experiment

Since inception of the grant, we have accepted responsibility for the following tasks.

1. **Level 3 Manager of HCAL detector control system (DCS).** S. Linn was appointed in June 2004. The position involves budgeting and coordination of the following systems:

- laser calibration (FSU/FIU)
- radioactive source calibration (Iowa/Purdue)
- forward calorimeter radiation monitors (Protvino)
- low voltage (Boston U)
- high voltage (Fermilab/Sofia)
- RBX parameter download (Fermilab/Protvino)
- Implement FW/PVSS hierarchy and FSM (Fermilab/Protvino)
- Integration of DCS with run control (FIU/Fermilab/Protvino)

Effort in year 2 was directed at completion of hardware and communications infrastructure as well as development of server software under a unified system which will fit into the CMS run control framework. The latter uses the JCOOP Framework which includes a state machine package (FSM) and a PVSS interface.

2. **M&O of HCAL laser calibration system**.

The HCAL laser calibration system is the responsibility of FSU. Starting in March 2004 a new system was designed and constructed with help from FIU, and tested at CERN in December 2004 before being returned to FIU for mechanical fine tuning, software testing, and development of monitoring and readout hardware. The system is currently being tested and documented at FIU by undergraduate students V.Gaultney and L.Lebolo under the supervision of S.Linn. In June 2005, the system (and students) will go to CMS. The students will each spend about 2 months at CERN with a week of overlap. The student at FIU will perform detailed data analysis of the system using network connections to Fermilab where all data is transferred.

3. **Installation, commissioning, and testing of HCAL.**

Effort in this area centered on the collection and analysis of HCAL test beam data. In addition to data collection shifts, FIU is responsible for analysis of the laser calibration data that was accumulated along with particle data. Analysis turned up a couple of interesting results

- The filter used to attenuate laser light had systematic non-linearities.
- Decoupling components of the resolution determined the HPD quantum efficiency via a new method.

Moreover, the code developed will be useful for future analysis of the laser system, and will be used this summer for that purpose.

To meet these additional duties, a CERN-based postdoctoral fellow was hired. During a visit to CERN, S.Linn was introduced to German Martinez who was looking for a research position. We were so impressed with his experience that we decided to fill the post-doctoral position a year ahead of schedule. Martinez started Feb 2005 and will reside at CERN and initially work on slice/magnet test preparations. Because of his previous experience, he may take a leading role in DCS/run control implementation. German will also provide day-to-day supervision of students when they are at CERN.

2.1.2 Simulation

With much help from J.MacDonald (FSU), the CMS simulation and analysis package was installed on the NAP cluster. Since Feb 2005 simulation of SuperSymmetric physics events in the CMS detector is under way.

2.1.3 Milestones and Budget

The CHEPREO collaboration has responded to requests by the CMS experiment to adapt our initial work in response to refined needs of the experiment. In particular, the original milestones Yx.5 dealt with the data acquisition slice prototyping. CMS centralized this effort at Fermi National Laboratory, and (at the request of CMS) FIU has instead taken over management duties for the slow controls of the detector (DCS) as the level 3 manager.

Similarly, the high voltage graphical user interfaces (HV GUI) which were the basis of milestones Yx.6 grew into responsibilities for the laser calibration systems maintenance and operations, at the request of CMS. CMS decided to develop an existing Fermilab/Dubna prototype based on Borland C++ Tools of the HV GUI instead of writing a new one.

This same responsiveness to the greater needs of the CMS experiment has resulted in FIU taking on responsibilities in the upcoming summer 2005 slice test and the HCAL commissioning.

Milestone	Description	Person
Y2.1	CMS Physics (TB2004, laser analysis, laser upgrade)	Pete Markowitz
Y2.2	Hire Research Scientist 2- FIU (not done)	Pete Markowitz
Y2.3	Begin HEP simulation with FSU (done)	Vasken Hagopian
Y2.4	FIU Grid work with UF (done with FSU)	Paul Avery
Y2.5	Start DAQ Slice Prototype (done by Fermilab)	Pete Markowitz
Y2.6	Multi-source High-Voltage GUI (done by Fermilab)	Pete Markowitz
Y2.7	Tier 3 DataCenter in NAP	H Alvarez
Y2	Total	
Y3.1	CMS Physics (HCAL slice/magnet test/commissioning)	Pete Markowitz
Y3.2	Hire FIU Postdoc (Martinez hired 1/1/05)	Pete Markowitz
Y3.3	Physics simulation with FSU (started 1/1/05)	Vasken Hagopian

The CHEPREO budget for this physics work has spent just under 90% of the funds. Year 3 summer (i.e., summer 2005) is expected to use the residual funds in support of travel to CERN, salary for the CERN-based postdoc and equipment for the laser calibration.

2.2 Educational and Outreach Efforts

Significant progress has also been made in CHEPREO's education and outreach efforts. Highlights of these efforts include:

- Renovation and inauguration of the Physics Learning Center (Phase I).
- Faculty search for upgraded (tenure-track) Physics Education Research Faculty.
- Completion of our second annual summer 3-week modeling workshops with plans underway for third annual summer workshops.
- Expansion of introductory modeling-based physics classes at FIU.
- Year 2 QuarkNet activities, highlighted by outreach to high school students.
- Matriculation of first physics PER graduate student.
- Selection of external evaluator complete.

These efforts have essentially completed the year 2 milestones identified in the PEP.

Phase I of the Physics Learning Center (PLC) underwent renovation over the past year and was recently opened to student and classroom use. Phase I includes the first modeling classroom (capacity 30), the student lounge, and the conference room, all equipped with video conferencing capability. Laboratory equipment and classroom laptops have been purchased as well as the initial complement of video conferencing / electronic classroom equipment. The remaining video conferencing equipment has been selected and is out for bid. Installation is expected in Summer 2005. Phase II includes an additional classroom and study group space, storage, and a formal entrance. Preliminary design is complete and construction will begin once the space is vacated.

The PLC is crucial to our growth, as it serves as the central gathering space for our growing CHEPREO community. Students are given essentially unlimited access to the space for studying / informal discussions / relaxation. Classroom space and laboratory materials are available outside of class time for student use. High school teachers and students are also encouraged to use the space and together we are exploring how they may utilize the space. The video conferencing capabilities are to provide immersive communication with community members at any site – CERN, high schools, collaborating institutions, or meetings.

In the past year our Physics Education Researcher position has been upgraded to a tenure track faculty line with the search and interview process nearly complete. The CHEPREO PEP listed the position as a researcher line that would convert to a faculty line after three year. This upgrade has allowed us to attract a higher quality candidate and the successful candidate will be able to pursue long term goals from the onset. The faculty member will significantly enhance our E/O activities.The position is expected to be filled by the fall 2005 semester.

In the summer of 2004, we held our second set of two modeling workshops for high school teachers and other community members. Both physics I (mechanics) and physics II (light and waves) workshops were held. The physics I workshop was for those new to modeling and was led by Jeff Steinert and Stan Hutto, both of whom lead a physics I workshop last year. The physics II workshop was to continue our partnership with our 2003 workshop participants. This workshop was lead by Mark Schober (one of the developers of the light and waves materials) and Matt Watson. The workshops were held concurrently with both groups sharing lunches to boost the community building. For 2005, we again will be operating two workshops concurrently.

Implementation of modeling-based introductory physics expanded in the past year. Two sections of physics was offered both in the fall and spring terms covering both physics I and II. Kramer led a physics I class in the Fall semester followed by physics II in the Spring. Narayanan led a physics II class in the fall term and Boeglin led physics I in the Spring term. Anecdotal evidence clearly shows the impact of the method – students in Kramer's class are considering becoming physics majors or minors, most students signed up for physics II after taking physics I with Kramer, and most students in Kramer's class

meet regularly in study groups. In addition, students from last year's implementation have continued to use skills developed in the modeling class as part of their regular study habits.

Our QuarkNet center entered its second year, expanding by 11 associate teachers to a total of 13 with the 2 leaders. We held a 2-week institute to introduce the new teachers to QuarkNet and HEP, followed by 5 meetings during the academic year. The institute included talks and lab tours from the physics department as well as laboratory activities centered around nuclear scattering phenomena. These activities prepared the teachers for the "Catching Some Z's" activity, an inquiry-based team investigation of HEP through analysis of simulated data. The academic year days have been centered on building the community while providing additional activities for teachers and their students. This evolution led to a combined open house / competition for high school students. We held an open house in February 2005 where each teacher invited 5-15 students (and parents) to visit FIU. The open house involved faculty and students through panel discussions, lab tours, and lunch with faculty and undergraduates. Students and teachers also toured the (then under construction) PLC.

Our first PER graduate student in physics was welcomed to the group in Fall 2004. Elliot Brown is becoming a valuable asset to the group, taking the lead on assessing adoption of modeling materials in the high school classroom.

Jeff Saul from University of Central Florida has agreed to serve as external evaluator for the project. He is already well acquainted with the project through informal discussions, his participation in the original Remodeling University Physics (RUP) workshop, and his leadership in developing the Scale-Up project at NCSU and UCF. Conversations are underway to identify the evaluation plan.

We also had Ed McClintock retire at the beginning of the Spring 2005 semester. His duties are being reassigned to Kramer and O'Brien on a temporary basis. Once our PER faculty joins the team, those E/O milestones will be reassessed and reassigned.

2.3 Cyberinfrastructure Update

2.3.1 Overview

The CHEPREO IT infrastructure budget has been revised to reflect a delay in releasing the equipment funds for a 10 Gb equipment upgrade until year 4, or as needed based on the eventuality of a continued presence of a 2.5 Gb bottleneck between Brazil, Miami and points north. While we were justifying the 10 Gb equipment upgrade now to put the AMPATH international exchange point on an even footing with the other US international exchange points in places such as New York, Chicago, Seattle and Los Angeles, we do understand that this will have to be delayed. On a more positive note, discussion occurred with the CISE-SCI program officer, Kevin Thompson, about the possibility of upgrading the 2.5 Gbps circuit between Sao Paulo and Miami to a 10 Gbps circuit if we demonstrate that the research and production traffic flows are not adequately sustained by a shared IP service over the year 3 activities of CHEPREO. Aggregate

production and research traffic would be sufficient reason to discuss CISE-SCI funding for an upgrade to 10Gb. Thompson would like to see experimentation with IP flows, which may prove insufficient to support finely tuned 1 Gbps flows, necessitating the need for additional bandwidth capacity to support a hybrid of shared and deterministic network services.

The overall IT infrastructure project execution plan closely resembles the 2003 proposal with these changes in order to achieve our milestones:

1) Introduction of the concept of AtlanticWave networking services spanning Sao Paulo through Miami and up the eastern seaboard to NY (MANLAN) providing high bandwidth services (up to 2.5Gb Brazil-Miami and 10GB from Miami to US international open exchange points, CERN, StarLight, FedNets, and other international networks. This service was justified in the original proposal and the PEP in year 3 by a link between Miami and StarLight through FLR and NLR to StarLight. Funding for this endeavor has been delayed for reconsideration in Year 4 of the CHEPREO project. The equipment needed is shown in Appendix A with a strike through and zero funding in Year 3, as explained above, reducing this years funding request for this activities by almost $250,000 from the 2003 PEP. ***Please see appendix A for complete details of the budget request*** which as also been submitted to the MPS-EPP program officers via a revised budget and budget impact statement.

2) Reprogram some permanent equipment funding to fund Ibarra for 1 month a year due to the need for his increased involvement with managing the CHEPREO cyberinfrastructure services.

The FIU Center for Internet Augmented Research and Assessment (CIARA), led by Executive Director, Julio Ibarra, Research Director, Heidi Alvarez and Chief Operations Officer, Chip Cox, is providing network engineering and network resources for international networks, grids and collaborative systems. In particular, with Caltech, UF and FSU, CIARA continues to extend grid-based computing to the FIU Physics Department's expansion into CMS physics. The AMPATH international exchange point, a project of CIARA located in the NAP of the Americas, continues to provide the iVDGL connectivity extension to South American countries through the CHEPREO STM-4, rather than the AMPATH research network as reported in year 1. We continue to develop plans to support Latin America's GRID (LAGRID) as a multi-disciplinary international resource in support of US science in Latin America, interconnecting with North America's and Europe's Grids. To this end Alvarez, Ibarra and Cox represented CHEPREO and CIARA at the following conferences:

LISHEP 2004	Heidi and Julio
GNEW2004 meeting at CERN	Julio
GLIF meeting	Heidi, Julio
EU IST conference	Chip
OSG meeting at Harvard	Julio
SC04	Julio
SURA Grid workshop	Julio

GridToday meeting in Costa Rica Heidi, Julio

Harvey Newman and colleagues from Caltech also represented CHEPREO at the following events and more information about their activities can be found in Section 5:
LISHEP2004 (Harvey Newman, Ravot, Nae)
GNEW 2004 (Harvey Newman, Ravot, Nae)
OSG Meeting at Harvard (Harvey Newman)
SC04 (Harvey Newman, Xun Su, Nae, Xia; virtually the Rio and Sao Paulo teams)

The CHEPREO project objective of serving under-represented communities through research, science outreach, Grid-based computing and international networking continued as a major focus of our year two activities.

2.3.2 Tier-3 Data Center in the NAP

The Tier-3 data center was upgraded per the project execution plan with assistance from Jorge Rodriguez at UF.

2.3.3 Hire IT Support Person Phased over 3 years

During year 2 CIARA hired Fabian Alcantara to provide IT Support with Ernesto Rubi. Both positions are now funded 50% from CHEPREO and 50% from AMPATH to insure backup coverage for both projects. ***Please see Appendix A for Leveraged Projects Personnel Support table for all CIARA personnel working on CHEPREO.***

2.3.4 Purchase and Install Active Equipment

The original design of the network link between the U.S. and Brazil was an IP link with routers at the end points. Advancements in optical technology, changes in best practices of technology for inter-regional network links, and the emergence of Sao Paulo as a network hub for state, national and international networks that reach Brazil's and South America's HEP community, motivated a change in design using optical multiplexers at the end points instead of routers.

Two Cisco ONS 15454 optical multiplexers were purchased to interconnect the CHEPREO STM-4 circuit between Miami and Brazil. These muxs provide additional functionality for the evolution of the circuit over routed equipment to support HEP's production and research traffic flows. Cisco provided aggressive discounts and a partial donation to not exceed the year 2 budget allocation.

For year 3, layer 2 and layer 3 ports are needed in Miami to support the traffic flows from the clusters between the partner sites in the U.S. and in Brazil.

2.3.5 Bid, Award and Provision circuit to Sao Paulo Y1 & Y4 15%

An STM-4 circuit was provisioned between Miami and Sao Paulo in September 2004. The circuit is provisioned on STM-16 (2.5Gbps) ports at each end. This enables the rapid allocation of additional bandwidth capacity when requested from the provider. The mechanism for requesting and allocating additional bandwidth was tested for SC04. FIU, Caltech and ANSP engineers, working with Internet2 and industry partners Cisco

Systems, Latin American Nautilus, Terremark, FPL FiberNet, established a 2.5Gbps link from Miami to Sao Paulo, and from Miami to the Atlanta Abilene router, then to the SC04 floor, where Brazil participated in the Caltech Bandwidth challenge (see Caltech's section in this report).

The NSF International Research Network Connections (IRNC) program awarded the international connection for Latin America to FIU and CENIC, with FIU as the lead institution. The WHREN-LILA project plan includes the increase of the capacity on the link from Miami to Sao Paulo to integrate with CHEPREO's cyber requirements and to support the international network requirements of Brazil's HEP community. The emergence of Sao Paulo as a hub for state, national and international R&E networks made it a strategic location for the international connection to Latin America. Sao Paulo is the primary node of the RedCLARA regional network for Latin America to connect to Europe's GEANT network. From Sao Paulo, RedCLARA connects to GEANT in Madrid over an STM-4 (622Mbps) link.

2.3.6 Provision circuit thru FLR to NLR Y2

The Florida Lambda Rail (FLR) network is on schedule to be ready for service in April 2005 in South Florida. As described in the PEP, CHEPREO plans to use the facilities provided by FLR to establish a connection to the National Lambda Rail (NLR). To date, CHEPREO has been using AMPATH's connection to Abilene and providing a routed IP service to CHEPREO while FLR is being completed.

With the availability of FLR, CHEPREO will proceed as described in the PEP to establish a link from Miami to Jacksonville for interconnection with the NLR to integrate CHEPREO and Brazil into the AtlanticWave (AtlanticWave is described below). CHEPREO and the Brazilian HEP community would leverage the AtlanticWave resources that would offer the capability to establish a hybrid of traffic flows from Sao Paulo to AMPATH in Miami, and to MANLAN in New York City. At MANLAN, CHEPREO will be able to leverage the IRNC resources to Europe, providing direct connectivity to CERN, as well as to Russia and China, through GLORIAD. The connection to StarLight described in the PEP for year 3 is to be reprogrammed to use the AtlanticWave instead. The use of the AtlanticWave satisfies all of the requirements that were described in the PEP, and in addition leverages the resources provided by AtlanticWave and the international network resources at MANLAN.

2.3.6.1 The AtlanticWave

The AtlanticWave interconnects the operating exchange points along the Atlantic coast of North and South America to create a distributed exchange and peering service for national and international R&E networks. FIU and CENIC, with support from SURA, described the concept and a commitment to the implementation of the AtlanticWave in their successful proposal to the NSF IRNC program. In particular, the FIU-CENIC led proposal describes the importance of establishing a 10G wave service between Miami, Washington DC, and New York, and having it interconnect with the NSF-funded link from Miami to Sao Paulo to enable a hybrid of scheduled temporary use and permanent use network services to support discipline-specific and general-purpose high performance

computing and networking services between North and South America and Europe. This activity does require additional active equipment which will not be requested in the Year 3 CHEPREO IT infrastructure budget, but will be revisited in Year 4 of the CHEPREO project as mentioned in the opening paragraph of this section (see 2.3).

2.3.7 Establish/Cont. Network Engineering Internship(s)

CHEPREO and the role of the AMPATH international exchange point for R&E networks is the impetus for the recent addition of FIU to the iVDGL, which will now extend to South American countries through the AMPATH international exchange point. We are also helping form Latin America's GRID (LAGRID), interconnecting with North America's and Europe's Grids. The focus of the Year 2 activities has been to establish and advertise internship opportunities with Terremark. One such opportunity has been a ITWomen -Terremark Scholarship[1], and another internship opportunity through the FIU school of computer.

3. Florida State University

Florida State University has a large number of personnel helping CHEPREO:

1. Vasken Hagopian, Faculty
2. Harrison Prosper, Faculty
3. Yuri Gershtein, Faculty
4. W. G.D. Dharmaranta, Visiting Faculty (on sabbatical) from Sri Lanka
5. Jeffrey McDonald, Assistant Scholar/Scientist
6. Kurtis Johnson, Scientist
7. Andrew Askew, Post Doctoral Fellow
8. Sharon Hagopian, Scientist
9. Daniel Duggan, graduate student

Florida State University contributions to CHEPREO have 5 major items:

1. Continue to work with FIU PHYSICISTS IN THE CMS Experiment. Specifically the Hadron Calorimeter calibration (both hardware and software).
2. Create the CMS Software environment both at FSU and FIU and connect to the OPEN SCIENCE Grid with the help of our colleagues at University of Florida.
3. CMS Simulations. Use CMS software and computers at FSU, FIU and CERN to generate SUSY events through the whole chain of software.
4. Electron and photon reconstruction and calibration. Create algorithms for electrons and photons.
5. Help with test beam at CERN, especially with the hadron calorimeter system.

a. Using CHEPREO as a part of the justification we were successful in convincing the Physics Department and the Arts and Science Dean to authorize a new faculty position and after a through search we have employed Assistant Professor Yuri

[1] http://www.itwomen.org/contribute/index.asp

Gershtein, who started in January 2005. He is in charge of the US effort of photon and electron ID of CMS. Using 50% CHEPREO Funds and 50% startup funds we have also hired a Post Doctoral Fellow Dr. Andrew Askew to work on the same CMS project.

1. **Integration of FIU in CMS Experiment.** This task is now successfully completed with Dr. Steve Linn being the lead person on this task. We have worked with FIU in getting a laser calibration system built which is now at CERN. The final laser system is under design and work will continue for the next year or so. The goal is to have 5% accuracy on day 1 and use the system to monitor the operation of the calorimeter. As the hardware projects get concluded, we will work more on the software side. Personnel working on this project are Kurtis Johnson, Daniel Duggan, Maurizio Bertoldi (engineer) and Vasken Hagopian.

2. **CMS Software Environment and CMS GRID.** The off line CMS software is comprised of five major components: **FAMOS** for fast simulation and reconstruction, **OSCAR** for detailed simulation. **ORCA** for digitization, reconstruction and analysis, **COBRA** the CMS framework and **IGUANA** for visualization. These software packages are written by CMS physicists and are continuously being upgraded. We have now ported the CMS software both to FSU and FIU and both systems are on GRIDs and we on of Monte Carlo production for CMS. On the hardware side, we now have 30 CPU's soon to be increased to 45 CPU's at Florida State University. Setting up a working system took about 5 months and we now routinely use the simulation packages for Physics studies. The personnel for this task are as follows: Dr. Jeffrey McDonald who has set up the CMS system both at FIU and FSU. Dr. Kurtis Johnson in helping procure fast computers at about $ 400 per CPU.

3. **SMS Simulations.** So far we have generated about 200,000 SUSY events on computers at FIU and FSU. Each event takes 4 minutes per node of each computer (CPU). We have used simultaneously the 20 CPUs at FIU and up to 20 at FSU. The preliminary conclusion is that the calculations performed by theorists are within a factor of two in estimating the signal. We also use the CERN computers to study the background where is a very large amount of background events have been generated.
   Personnel working on this portion of the task include Prof. Harrison Prosper, Prof. W Dharmaratna (on sabbatical leave from Sri Lanka, a former post doc of FSU), Dr. Jeff McDonald and we expect two graduate students to join this effort in May 2005.

4. **Electron and Photon reconstruction and Calibration.** Prof. Yuri Gershtein and Post Doctoral Fellow Andrew Askew have started writing the algorithms for electron and photon ID using both the electromagnetic calorimeter and tracker. This effort started in January 2005.

5. **Test Beam at CERN.** During the past year several of us were at CERN using the test beam to study and calibrate hadron calorimeter modules. This test has been very successful, which used the laser and LED systems built jointly by FIU and FSU. Personnel working on this task are Dr. Kurtis Johnson and Prof. Vasken Hagopian.

SUMMARY

- Work has already started on the CMS experiment and we expect to contribute substantially during the next decade.
- FSU scientists will visit FIU on a need basis.
- We look forward towards a very fruitful collaboration.

4. University Of Florida Progress

4.1 Summary of Grid Education

The grid education effort is led by Sanjay Ranka of the University of Florida, with help from Paul Avery of UF. We have surveyed the existing literature on Grids and developed a first version of a coherent set of slides that address the following:

1. *Requirements of Distributed Data Intensive Applications*
2. *Web services*: Programming using web services.
3. *Basic Grid concepts*: Available grid toolkits. Good sources of information.
4. *Deploying a Grid*: Using various packages and tools like VDT, Condor, Monalisa.
5. *Programming the grid*: Includes writing applications and using Globus tools, OGSI, .Net and GT3.
6. *Deploying applications*: Maintaining and using the grid. Problems and pitfalls involved.

Pradeep Padala presented this work a 1–2 day Grid tutorial by May 2004. Laukik Chitnis will present this tutorial as part of a summer workshop in Argentina.
This grid material was sent to Chi Zhang at FIU. Chi Zhang is integrating this tutorial as part of a graduate course at FIU. The course "COP 6611 Advanced Operating System" focuses on distributed computing. It will include 6 hours of lectures on Grid, and an exercise on "gridifying" an application.

4.2 Cyberinfrastructure at the University of Florida

UF has been searching for a person who will work with FIU physicists and computer personnel on Grid-related aspects of Cyberinfrastructure. Dr. Jorge Rodriguez has continued working closely with the FIU team including organizing a grid workshop for Rubi and Alcantara. Rodriguez continues to collaborate on the following activities begun in year 1:

1. He consulted closely with FIU engineers and physicists on FIU's equipment expansion for FIU's Tier3 cluster, making sure that the components were compatible with the variety of applications expected to be run in a Grid environment.
2. Jorge worked with sysadmins and the High Energy Physics groups at FSU and FIT (Florida Institute of Technology) in January-February 2004 to configure their clusters for operation on Grid2003.

4.3 Physics Activities at the University of Florida

The University of Florida has one of the largest U.S. university group (23 members) participating in CMS. UF is a designated Tier2 Computing Center in the worldwide CMS Data Grid and takes part in many national and international initiatives related to physics computing. In particular, UF takes part in the worldwide simulations organized at CERN. We are working with AMPATH-FIU to integrate FIU into this physics simulation effort as their Tier3 computing equipment arrives.

UF is also working with Fermilab (the CMS Tier1 Computing Center) to develop automatic procedures and tools to maintain the very complex CMS physics software repository at remote institutions. These activities will improve our ability to provide physics analysis support for FIU, but will not become effective before Fall 2004. UF will work with FSU and Caltech in providing expertise and support for FIU physicists in this early part of the project.

5. California Institute of Technology Collaborative Activities Update

5.1 Overview

Highlights of year two activities focused on the Brazilian HEPGrid plans which are progressing well. Presenting a plenary talk at the inauguration of the UERJ Tier2 and taking the occasion to help develop agreement between ANSP and RNP on connection of UERJ as well as USP at 1 Gbps by mid-year, with a future upgrade to 2 X 1 Gbps, are among the steps putting the Brazilian HEPGrid plans on a firmer footing. The viability of the UERJ Tier2 and its future contributions to the CMS simulation and reconstruction efforts (as soon as the network is in place) was recognized by the CMS Management Board in mid-March, 2005.

A meeting in Rio with Co-PI Harvey Newman, colleagues from Fermilab, and a representative of international CMS management is planned for May12-13, 2005. Meeting attendees will engage in discussions with Roberto Salmeron, current responsible for organizing the funding for High Energy Physics, in order to support a new project for UERJ, USP and other Brazilian university groups to participate in the CMS detector development.

Other Caltech activities detailed below range from cyberinfrastructure activities such as network monitoring research and support, network engineering development and support,

grid and distributed computing development and support, to high energy physics research, as well as network tutorial educational outreach, and VRVS collaborative services.

5.2 Network Monitoring

Xun Su worked with FIU engineers Ernesto Rubi and student Jose Fernandez over the last summer, on an NSF-funded Research Experience for Undergraduate (REU) project. They integrated Cisco NetFlow and MonALISA so as to monitor the traffic going through the AMPATH PoP at Miami. This involved developing software modules/scripts to interface the NetFlow analysis tools (Flowtools) to MonALISA's ApMon module that monitors running applications, and deploying an appropriate MonALISA database module (Postgres instead of MySQL). This system is currently running and providing flow-level traffic reports for AMPATH. Please see Appendix B, Exhibit 1 for a graphical example of The network traffic flows traversing the AMPATH PoP and destined to Internet2's New World Symphony orchestral academy site.

5.2.1 Cisco University Research Program

FIU and Caltech are funded by a Cisco University Research Program (URP) grant to study and develop tools for packet-level network monitoring. In this activity, Xun Su has been working with Professor Chi Zhang at FIU and his student Bin Liu on integrating MonALISA and NLANR's

Passive Measurement and Analysis (PMA) system. PMA is currently installed at AMPATH and managed by NLANR, taking several traces per day. After initial implementation of the software modules based on CAIDA's CoralReef, they will consider deploying a CHEPREO passive measurement instrument. This will not only provide detailed packet-level traffic information for CHEPREO/AMPATH networks, but will also improve the general understanding of the value of flow vs. packet level traffic monitoring systems, and it will help inform a decision on the value of DAG-card based monitoring systems (see http://dag.cs.waikato.ac.nz/dag/dag3-arch.html)for other projects such as Ultralight and WAN-In-Lab.

5.2.2 SLAC PingER

Xun Su installed the SLAC PingER system for the CHEPREO/AMPATH network. PingER is a active measurement system that monitors includes a Traceroute server and a data collection module based on the ubiquitous "ping" utility, which is currently collecting delay and packet loss information from AMPATH to and from other major PingER installation sites. PingER's coverage currently includes 673 sites in 114 countries.

5.3 Network Engineering Development and Support

5.3.1 STM-4 Circuit Miami-Sao Paolo

Xun Su worked with the FIU team on provisioning the STM-4 (622 Mbps) CHEPREO circuit between Miami and Sao Paulo. This circuit was used (and temporarily upgraded to STM-16, or 2.5 Gbps) during the SC2004 Bandwidth Challenge demonstration (see below for more details on the Bandwidth Challenge). Single TCP flows of 1.95 + 0.98 Gbps were observed between Pittsburgh and Sao Paulo - the largest flows ever seen on a Latin American network, utilizing 2 gigabit Ethernet ports (one full-duplex) at near-100% efficiency, without packet loss. The circuit is in the final stages of re-provisioning with double STM-4 capacity using two Cisco ONS 15454s. This will allow one mission of the CHEPREO cyberinfrastructure, that of providing sufficient bandwidth for the Brazilian High Energy Physics community to be achieved. It will also facilitate the integration of Miami/Sao Paulo into the AtlanticWave distributed exchange point.

5.3.2 South America Networking, GIGA project

Harvey Newman attended a planning meeting between A. Santoro, head of the UERJ group, S. Novaes, head of the USP/UNSP group, Luis Lopez, Director of ANSP R&E network for Sao Paulo, and Michael Stanton the head of the GIGA project that links many universities in Rio de Janeiro and Sao Paolo, and head of planning for the RNP national network. This meeting resulted in an agreement to provide GbE access for both the Tier2 in Rio and the Tier3 in Sao Paolo to the CHEPREO link. Caltech and Cisco will donate two Cisco 3750 gigabit Ethernet switches to the ANSP Sao Paolo, for use by the groups at UERJ and USP, and Caltech agreed to transfer its knowledge of high throughput network methods to the Brazilian academic community. Startup of this connectivity is planned realistically for June or July. Plans call for an upgrade to 2 X 1 GbE access across the GIGA project by mid-2006.

5.3.3 Brazilian National and International Networking

In October 2004, Dan Nae traveled to Rio de Janeiro to help prepare for the SC2004 Bandwidth Challenge. The international links and the local servers were extensively tested to ensure optimal operation. By working closely with the Brazilian engineers for two weeks, he was able to get an accurate understanding of the local research and academic network (RNP) status. A map of the network connecting the research community in Brazil to US and Europe is shown in Appendix B, Exhibit 4.

5.3.4 SuperComputing 2004 and the Bandwidth Challenge

In November 2004, the CHEPREO team joined forces with a group of scientists and network engineers from Caltech, SLAC, FNAL, CERN and other international partners at the SuperComputing 2004 Bandwidth Challenge and captured the ”Sustained Bandwidth Award” for their entry entitled “High Speed TeraByte Transfers for Physics”.

The group generated an aggregate rate of 101 Gbps to the show floor at Pittsburgh. This rate is of the same order of magnitude as the total traffic expected in the early phase of LHC operation! The data was sent from servers at Caltech, SLAC, FNAL, CERN, the University of Florida and Florida International University, as well as at other international partner sites in Brazil (Rio de Janeiro State University, the UK (University of Manchester and UKLight), the State University of Sao Paulo and the Academic

Network of Sao Paulo) and Korea (Kyungpook National University). The achieved bandwidth was made possible in part through the use of the FAST[2] TCP protocol and through the use of seven 10 Gbps links to Caltech's Center for Advanced Computing (CACR) booth, and three 10 Gbps links to the SLAC/Fermilab booth as shown in Appendix B, Exhibit 4.

5.4 Computational Development and Support

5.4.1 UERJ Tier2

Michael Thomas made two trips to UERJ (Rio de Janeiro) to assist their HEP department in preparations for becoming a CMS Tier2 center. He worked closely with their software and hardware teams in order to design and implement the cluster architecture for their site. Both MonALISA and Clarens were installed on their new cluster, which enabled them to participate in SC2004 as part of the Grid-enabled Analysis Environment (GAE) demonstration, on distributed job processing and monitoring. These demonstrations used GAE/UltraLight software components such as MonALISA and Clarens, as well as the job submission and tracking tool BOSS. The two site visits were followed up with ongoing, numerous email conversations in an effort to continue to make progress on their involvement in CMS production.

Suresh Singh was involved in the early design phases of the UERJ Tier2 cluster architecture. This included setting up four separate interconnected sub-clusters with a common network providing an external Internet connection. Some of the challenges faced (and overcome) were lack of proper software drivers in the cluster management software for the type of the hardware that were acquired, lack of well-established external network connection and the lack of a common switch that integrated the sub-clusters. NPACI Rocks 3.3 cluster management software was used to build the clusters, each of which has one head node, one storage node and several compute nodes. All the nodes are members of a private VLAN whereas the head node is configured as a multi-homed system that has both of its interfaces connected to private and public networks. By default, all of the UERJ Tier2 nodes run Redhat Enterprise Linux 3 (RHEL3) with the 2.4.21 SMP kernel.

One of the initial objectives of the UERJ Tier2 center was to join Grid3 grid project. The Grid3 project which was the precursor of Open Science Grid (OSG), played a pivotal role in establishing the first inter-continental experimental grid infrastructure. Suresh Singh worked on bringing the UERJ Tier2 into Grid3, which involved working with Grid3 management, and subsequently installing and configuring the Grid3 middleware obtained from the Grid3 software cache. The major component of this cache is the Virtual Data Toolkit (VDT) that includes several software components such as the Globus toolkit, Condor job scheduler, MonALISA monitoring framework, EDG VO management toolkits etc. Suresh installed and configured UERJ's Condor job scheduler. After obtaining the necessary DOEGrids signed host certificate it was configured to run CMS Grid Analysis jobs. This capability was demonstrated at the SC2004 conference.

[2] http://netlab.caltech.edu/FAST/

Harvey Newman has been involved in the development of the UERJ Tier2 concept and design. He gave the keynote at the Brazilian HEPGrid Tier2 inauguration in December 2004, and continues to assist planning for ongoing Tier2 hardware upgrades.

A major meeting among the Brazilian groups, US CMS represented by Caltech and Fermilab, and international CMS management will take place in mid-May in Rio. The startup of sustained simulation-production for CMS using the Brazilian Tier2 (at UERJ) and Tier3 (at USP) facilities as well as the CHEPREO link, will be discussed along with the overall Brazilian CMS program on detector development and physics analysis, with the Brazilian funding agency and network providers (ANSP and RNP). Caltech network engineer Yang Xia and a member of the GAE will attend the meeting to provide onsite technical assistance, tutorials and hands-on training during the meeting. These sessions will cover CMS and GAE software, and optimizing system, server and network performance.

5.4.2 Distributed Processing Demonstration

Michael Thomas prepared and ran a Distributed Processing / Job Monitoring demonstration in collaboration with USP (Sao Paulo, Brazil) and UERJ (Rio de Janeiro). The Job Monitoring demonstration showed how CMS analysis jobs could be submitted to multiple sites on the grid and be monitored from anywhere else on the grid. Multiple CMS analysis jobs were continuously submitted to 3 different sites, including UERJ and USP, using the BOSS tool running as a Clarens Grid service. While the jobs were running, BOSS would report the job state (scheduled, running, cleaning up) to a MonALISA monitoring server. Viewers at multiple locations could connect to the MonALISA monitoring network to view the status of each individual job in realtime. Please see Appendix B, Exhibit 2 for the graphic.

5.5 High Energy Physics Research

5.5.1 Higgs Diphoton Studies with UERJ

In preparation for the CMS physics TDR, due in December 2005, discovery potential estimates for the Higgs decay channel to two photons are being calculated by the Caltech group in collaboration with other institutes including Santoro's UERJ group. In particular, Vladimir Litvin worked with UERJ on developing suitable preselections for the PYTHIA event generation runs.

5.6 Educational Outreach; Networking Tutorials

During the Digital Divide and HEPGrid Workshop that took place in Rio de Janeiro, February 16-20, 2004 (http://www.uerj.br/lishep2004), Caltech engineer Dan Nae presented a series of networking tutorials. Starting with an introductory tutorial, gradually several more advanced topics (WANs/optical networks, TCP performance in WANs, MPLS) were approached.

5.7 Global Collaboration Service (VRVS)

The VRVS team provides operation, Quality Assurance and support in all Time Zones on a daily basis. In this regard, the VRVS Operations sub-team is responsible for all the

aspects concerning operation, maintenance and support of all the critical tasks of the system. Please see Appendix B, Exhibit 3 for detailed accomplishments and graphics

Appendix A: Leveraged Projects Personnel Support & Year3 Budget

CIARA Personnel	CHEPREO	IRNC	AMPATH-STI	OTHER
Year 2				
Heidi Alvarez	25%	0%	25%	50% FIU
Julio Ibarra	0%	0%	0%	100% FIU
Ernesto Rubi	50%	0%	0%	50% AMPATH IXP
Fabian Alcantara	50%	0%	0%	50% AMPATH IXP
Bin Lu	Tuition	0%	0%	50% CISCO URP
Xun Su	100%	0%	0%	0%
Ikam Acosta	0%	0%	50%	50% UltraLight
Year 3	**CHEPREO**	**IRNC**	**AMPATH IXP**	**OTHER**
Heidi Alvarez	25%	25%	10%	40% FIU
Julio Ibarra	8%	8%	44%	40% FIU
Ernesto Rubi	50%	0%	50%	0%
Fabian Alcantara	50%	0%	50%	0%
Bin Lu	Tuition	0%	0%	50% CISCO UPR
Xun Su	100%	0%	0%	0%
Ikam Acosta	0%	0%	50%	50% UltraLight

Table 1 Leveraged Projects Personnel Support

APPENDIX B: Caltech Exhibits

Exhibit 1: Network Monitoring

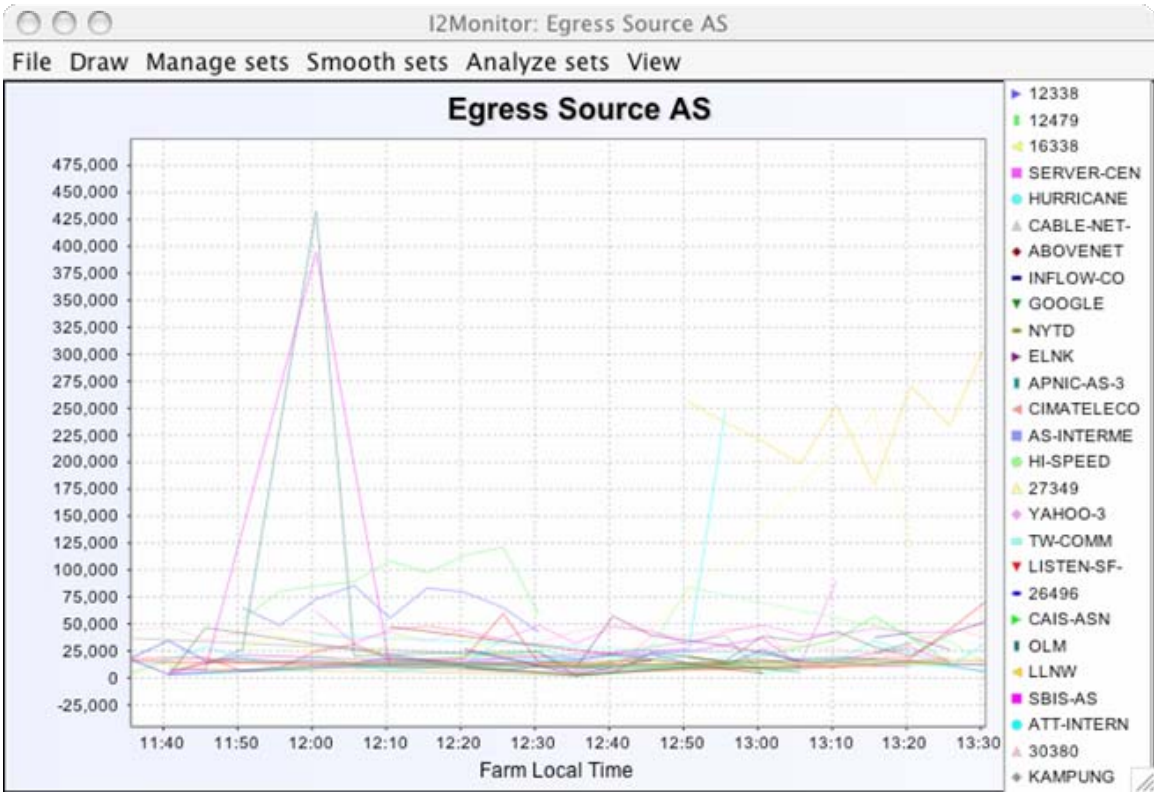

Figure 1 The network traffic flows traversing the AMPATH PoP and destined to Internet2's New World Symphony (NWS) orchestral academy site (see www.nws.edu), categorized according to their source Autonomous System (AS) numbers

Exhibit 2: Job Status

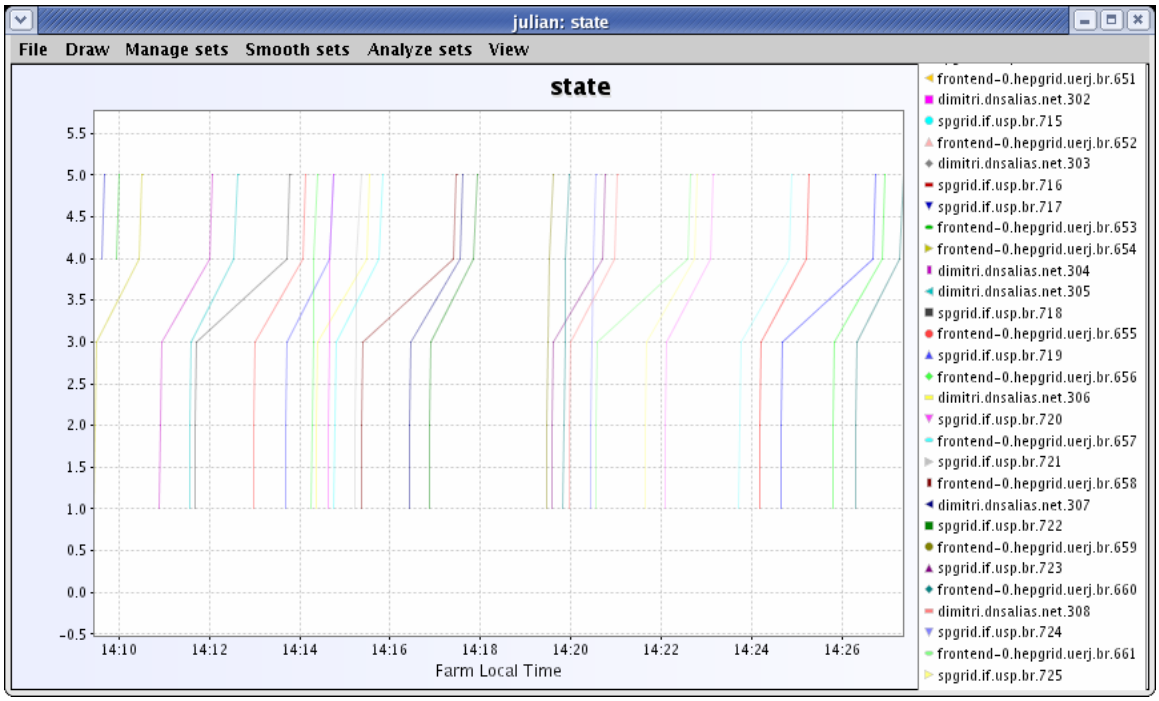

Figure 2 Job status represented as 'lifelines'. The vertical axis represents the state of each job, progressing from job submission (state=1.0), job execution (state=3.0 to 4.0) to job completion (state=5.0).

Exhibit 3; VRVS Accomplishments

- VRVS server management and maintenance in particular critical nodes at Caltech, Starlight and CERN
- Reflector backbone (82 machines) management, maintenance, upgrade and installation
- Worldwide Multilanguage support to VRVS users
- Testing infrastructure of VRVS system, new equipment or software and users setup
- Permanent monitoring of on-going meetings, of servers and general connectivity
- Maintain clear documentation of all the VRVS servers (problems, OS versions, VRVS software versions installed, etc.)
- Maintain all the contacts and interactions with VRVS community administrators and on site reflector administrators

- Maintain, configure and install the MonaLISA Monitoring software in the reflector backbone

The operational activities are thus becoming critical since VRVS has to run in an extremely reliable mode, due to its growing importance for the worldwide Research and Academic Communities. More than 53,000 host computers and 15,000 users from more than 106 countries are currently registered in the system. An average of 1000 worldwide collaborative sessions are held every month, involving more than 3500 users.

The VRVS system has been used regularly for the CHEPREO project as an efficient way to collaborate between the network engineers in Brazil, Caltech and CERN. Several tutorials have been given as well. The Brazilian HEP groups regularly use VRVS to participate in meetings both for the D0 experiment and the CMS experiment, and many other groups throughout Brazil also use the system extensively.

A new 7th reflector that manages the video/audio streaming in multipoint mode has been installed in Brazil recently. This makes Brazil the largest country after the US in terms of reflector deployment. There are currently a total number of 82 reflectors deployed among 28 countries. VRVS uses the AMPATH network to connect to Latin America.

Figure 3 Speaker display mode in VRVS

The VRVS Project offered to each of its partners (initially all the national research and education networks (NRENs), and very large scientific research projects) a dedicated "Community" of virtual rooms hosted in the VRVS Web server. A dedicated VRVS community has been created for RNP (Brazil), REUNA (Chile), REACCIUN2 (Venezuela) and CIARA (the Center for Internet Augmented Research and Assessment, based at FIU), as well as many other US and Europeans NRENs.

Glossary

BOSS: http://boss.bo.infn.it/

NetFlow: NetFlow Services Solution Guide, Cisco Systems Inc.
FlowTools: http://www.splintered.net/sw/flow-tools/ .
MonALISA: http://monalisa.caltech.edu .
ApMon: http://monalisa.caltech.edu .
PMA: http://moat.nlanr.net
CoralReef: http://www.caida.org/tools/measurement/coralreef/
PingER: http://www-iepm.slac.stanford.edu/pinger/

Exhibit 4: Brazilian Networking Connections

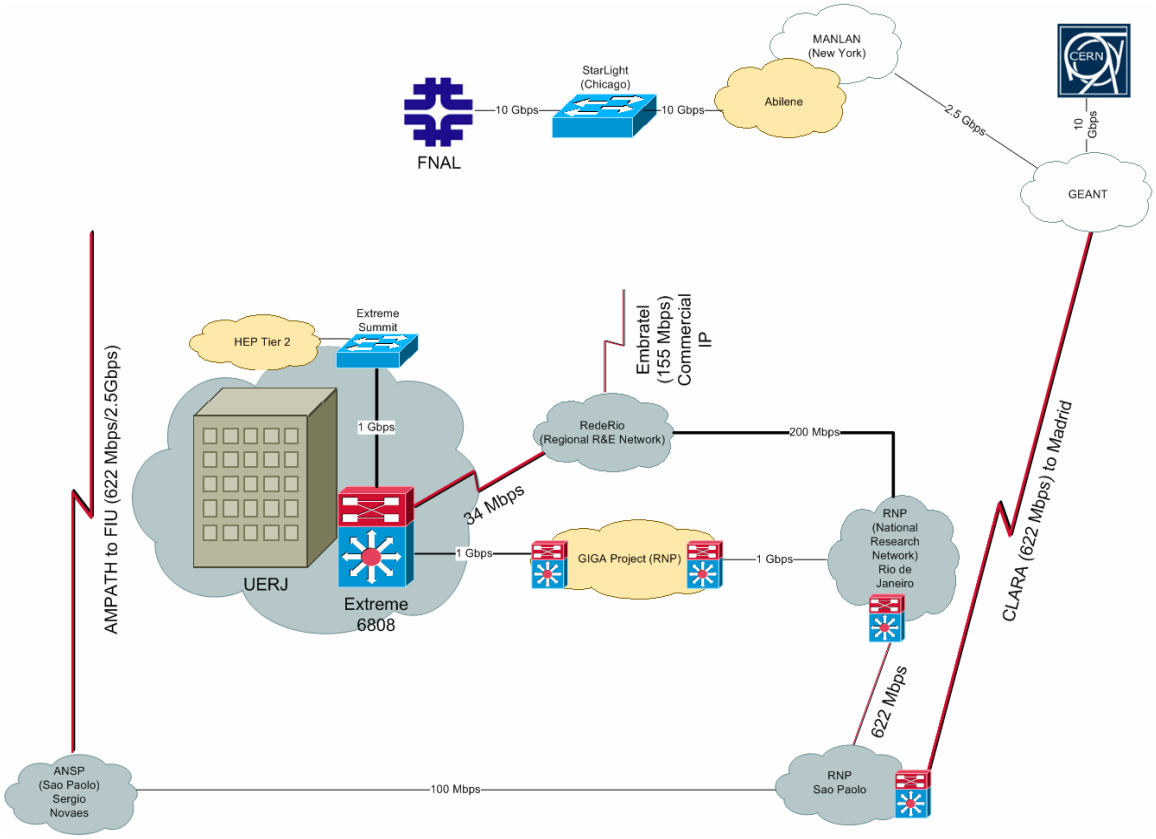

Figure 4 Brazilian network connections, in Brazil and to the US via the CHEPREO link and also (temporarily) via GEANT

The National Research Network of Brazil (RNP) connects several regional and metropolitan networks, such as RedeRio in Rio de Janeiro or ANSP in Sao Paolo. RNP's inter-regional GIGA project aims at connecting the research communities in Brazil with a high-speed optical network. Currently stretching 700 km and connecting Rio de Janeiro and Sao Paolo with a gigabit network, it is expected to reach over 4000 km in the upcoming years and to extend to the northeastern regions of the country. As illustrated in

the map, Brazil has a direct connection (622 Mbps) to the US and another link to Europe (Geant) via the Latin America project CLARA.

The connection to GEANT is for general purpose use at moderate bandwidth, although Caltech and the Brazilian groups were were permitted to send 500 Mbps for a short demonstration shortly before SC2004.

The CHEPREO project managed the connections to South America, including the provisioning, the deployment and the configuration of the then-new CHEPREO link, that was temporarily upgraded from its production capacity of 622 Mbps to 2.5 Gbps for the duration of SC2004. This link was specially deployed between Miami and Sao Paulo just for the Bandwidth Challenge, and allowed High-energy physicists at the University of Sao Paulo to participate. During the Bandwidth Challenge, a new record was set for a Latin American research network with throughputs up to 2.95 Gbps between the conference Center floor in Pittsburgh and the computational clusters in Brazil. The traffic generated on the link between Sao-Paulo and Pittsburgh during the demo is shown in Figure 6. This was a first big step in enabling physics groups in Sao Paulo, and also Rio de Janeiro (through the use of the GIGA project) to take part fully in the LHC and FNAL Run2 physics programs, through the efficient use of Gigabit/sec scale networking, together with their Tier-2 and Tier-3 regional centers. In many ways this step was made possible by the CHEPREO project.

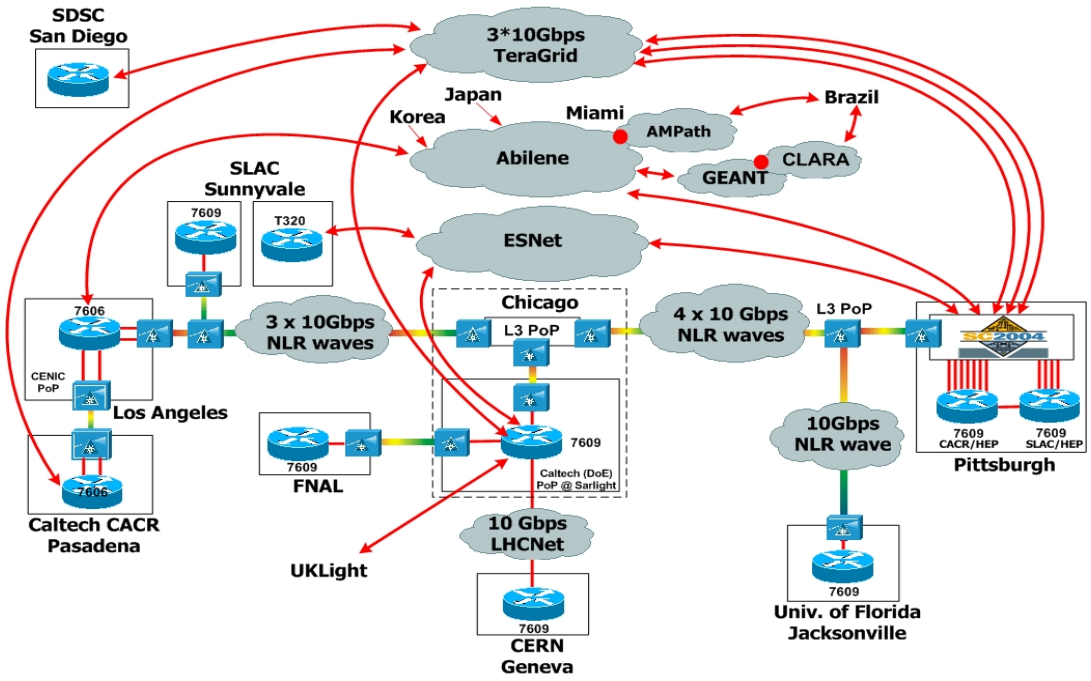

Figure 5 High Speed TeraByte Transfers for Physics Demonstration (SC2004)

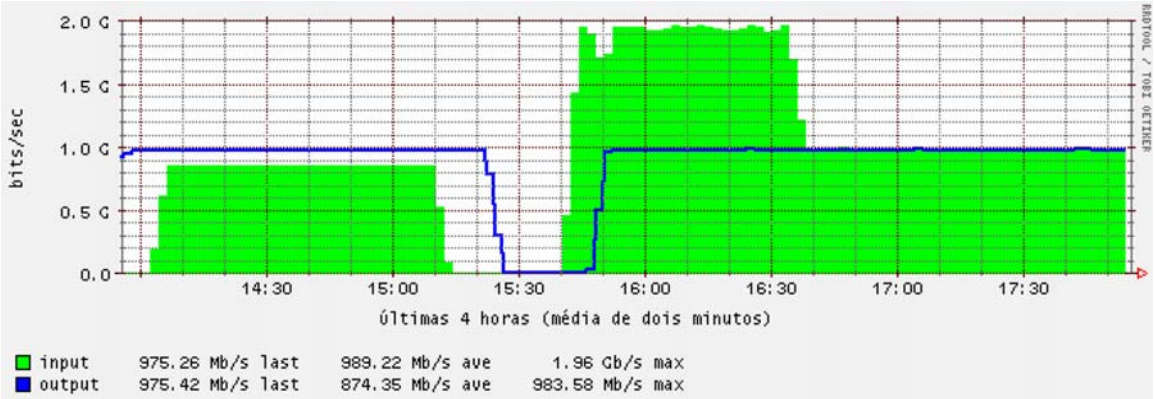

Figure 6 U.S. to Brazil traffic during the SC2004 HEP demonstration.