



The COOK Report on Internet Protocol Technology, Economics, and Policy



© Cook Network Consultants

Shaping a New Science Infrastructure Global Collaboration Mandates Inventive Solutions to Resource Constraint Problems

Editor's Introduction

This second and final part of our survey of Harvey Newman's work examines the LHC global Optical VPN. We find out how Harvey established a group that made alterations to the network's use of TCP/IP to better fill his wavelengths and solve other engineering problems found in using his global optical links. These achievements include the optical VPN network monitoring system known as MonALISA, and the FDT application. His buildout included the establishment of ICFA, a global organization monitoring the needs of the collider community. This group has met with success in building research community owned dark fiber into places like eastern Europe, Latin American and India. An off shoot has been programs in global network based e-science and the development. Another off shoot has been EVO a versatile global tele-

presence system running atop the optical infrastructure. We shall look at all these developments that one would hope will have commercial application but that are, as yet, not well known outside the research world.

Building the Global VPN

Newman: To make the LHC into a viable instrument for fundamental research, we had to build a global network collaboration of unprecedented scale. The global network diagram of the network of "tiered" networked facilities needed to handle the data. The concepts of the data flow patterns that will occur, and the bandwidth required to support those flows, have evolved over time. The present picture has not changed much in the last couple of years, but it may change again in the near future, when we have large quantities of real data from

Volume XVIII, No. 2
May 2009
ISSN 1071 - 6327

the LHC to contend with.

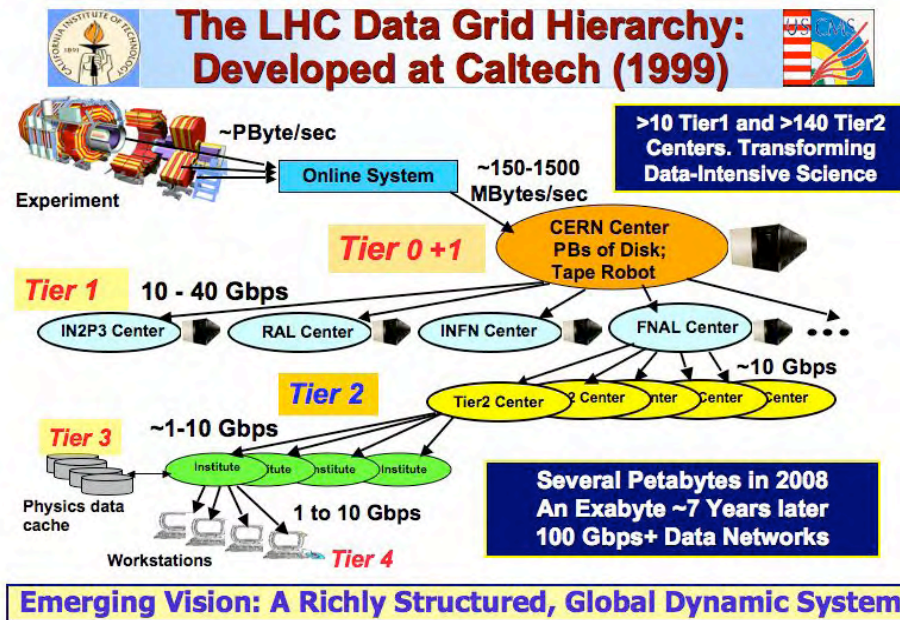
COOK Report: As it evolves will you have more Tier3 and Tier4 centers, or other centers rather than the "tier" ones?

Newman: There are now 140 Tier2 centers and this is already at, or close to the level where those organizations that can afford them already have them. There is basically one Tier3 computing cluster for each individual physics group, which they install and manage to satisfy their local needs. Tier4, at

On the Inside

**Research & Education
Innovation
Contents p.86**

**Please read Explanatory Note
page 85**



the bottom of the picture, is the ensemble of thousands of individual desktops and laptops belonging to the researchers. So we don't expect the concept of the lower tiers, and the number of facilities or systems in each layer, to change much over the next few years.

One issue with the **grid systems which have been constructed so far** (unlike those we originally conceived and have since continued to develop) is that they **are based on middleware without an end-to-end global view**. This makes operation of the system relatively difficult, and manpower intensive.

From the moment we first created the LHC Data Grid Hierarchy shown in the illustration above, and increasingly as we studied how it could operate with simula-

tions, we realized that we would need systems having a global view with the ability to follow thousands of actors and many tens of thousands of different tasks at once. It was also clear that we required the ability to understand what is happening in the end systems in order to isolate any problems, whether related to network throughput, reading or writing to storage, or the scheduling of jobs where the data and computing resources need to be co-located.

Dealing with these issues would require global systems, and ones with the capability of responding to problems in real-time. We then put considerable effort into learning how to build this new class of systems, starting with the simulation tools that were developed to study the global distributed system in the first place.

And we also went into testing the limits of the networks (and later developing tools and methods to extend these limits) because I knew that the capability of the networks would determine how dynamic the data flows could be, as well as the optimal balance between the use of large centralized facilities, and the use of a greater number of distributed smaller facilities. I knew that all of these factors depended crucially on our ability to get very high throughput among the sites. This started us down the path of learning how to use long-distance networks very well.

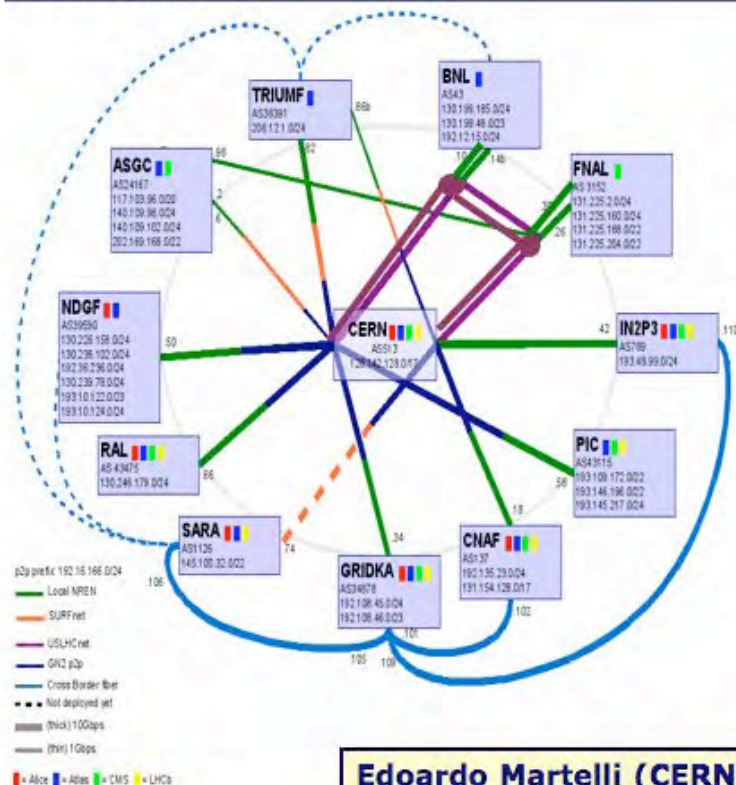
Another thing that really helped in 2001, which was about two years after completion of the data hierarchy diagram, is that one of my colleagues at Caltech, Stephen Low, came to my office and said he wanted to collaborate with me. Since he is in another field, computer science, this surprised me at the time. Steve knows a great deal about the theory and practice of TCP congestion algorithms. He showed us a new and very much improved congestion algorithm called FAST TCP that he had developed. Seeing FAST's capabilities helped to motivate us further to continue testing the limits of networks over the years, including our presence, and our records for



LHC Optical Private Network (OPN)



LHCOPN - current status



- ◆ Close collaboration with CERN & GEANT2
- ◆ Working with GEANT2 to achieve fiber diversity for European circuits to/from CERN
- ◆ Inter-Tier1 links for resilience, redundancy:
 - ◆ BNL, FNAL: US LHCNet
 - ◆ CNAF, SARA, GridKa, IN2P3: Cross-border dark fiber
 - ◆ NDGF, ASGC, TRIUMF: Reduced BW backup
 - ◆ RAL, PIC: No backup
- ◆ Evolution to N X 10 Gbps: Fermilab, BNL, US LHCNet; SURFnet

Edoardo Martelli (CERN)

throughput at the annual Supercomputing conference, every year since 2002. This also led later to a new application called Fast Data Transfer (or FDT) which was written by a member of my research group, Iosif Legrand.

The slide above shows the core of the present LHC network, called the LHC Optical Private Network (or LHCOPN). It shows the basic structure of the networking between the so-called Tier 0 at CERN and the Tier1 centers elsewhere.

The slide shows a concentric view - especially from the

point of view of the central management at CERN. In reality there is more to the overall system than that. The most important process is indeed to take the data generated at CERN, have it processed there for the first time, and then distribute that data to all the national centers.

But in the overall system, it is equally important to get the data needed to the Tier 2 centers, where most of the analysis as well as the production of simulated data is done. The Tier2s also have local control over some of their own resources, making them well-adapted to serve

the local needs of the physicists, as well as some of the needs of physicists at neighboring institutions.

The hierarchical pictures that you see in these two slides are useful to focus people's thinking, but they don't fully communicate the complexity of the data flows and kinds of workflows that you actually see. Namely people at a Tier 2 center would like to be able to get data from any place in the system. That is to say any other Tier 1 system or another Tier 2. People working on similar physics goals at different sites often work together and share results

and develop new algorithms jointly, using the same of different Tier2s. So the data needs to be distributed among the Tier2s so as to support this distributed analysis efficiently.

COOK Report: Is it reasonable then to say that all of your centers band together in a kind of private optical internet where each needs to be able to communicate with any other?

Newman: We'd like for them to be able to do that and, if you have limited network resources, you do have to manage how they are used. We also want to support transfers where Tier3 clusters serving individual physics groups get data sets from a

neighboring Tier2, and on a small scale from virtually anywhere. For us, "small scale" would be on the order of 1-10 Terabytes, so this can be done with throughputs of one to a few Gbps.

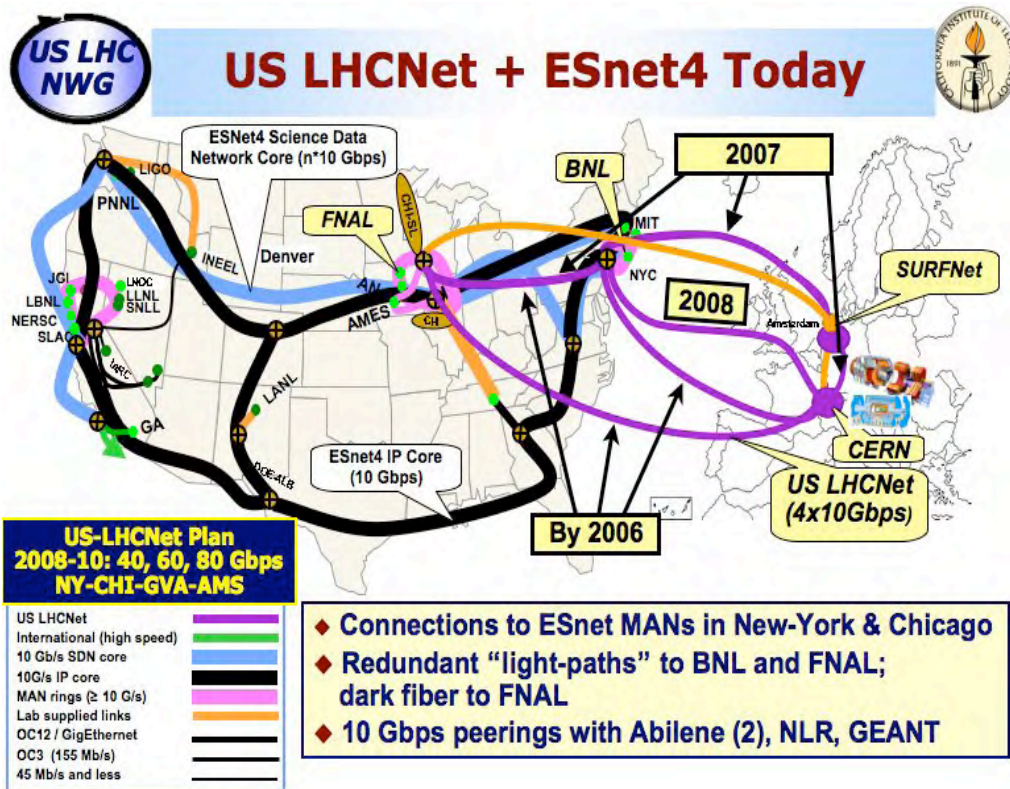
Your article about Surfnets [January 09 **COOK Report**] describes a relatively rich picture, where everyone has plenty of resources. We are not quite there. Even in our network roadmaps where the bandwidth to be provided is large, we think we are going to be resource constrained. We will therefore need artificially intelligent systems like MonALISA that allow us to view and manage what's going on, as well as deal with problems in real-time.

Capital Constraints Meet Network Management and Growth Needs


COOK Report: Are the resource constraints that you face more a matter of capital, or of human engineering, namely training people to use the resources?

Newman: Capital constraints are more important. As far as network management is concerned, we can view 100,000 things going on at once. All system managers can have at their disposal appropriate information to help them understand and mitigate problems. Many of the operations can be rather automated. I don't think human engineering is the main conceptual problem because we have learned how to build global systems like this - as exemplified, for example, by MonALISA. So the first issue is indeed capital, in particular the cost of networking and what the funding agency will provide.

The slide to the left illustrates the synergy in the United States between our network called **US LHCNet and ESnet, the Energy Sciences network**. ESnet serves all of the DOE funded programs, which are largely based at the national laboratories. These programs are not only in high energy physics but



ESnet Traffic has Increased by 10X Every 47 Months, on Average, Since 1990



J. Metzger
ICFA DDW07

also in many other fields such as climatology, geophysics, fusion energy, materials science, bioinformatics and other areas. But one of the main drivers of ESnet's present capabilities and forward planning is high energy physics.

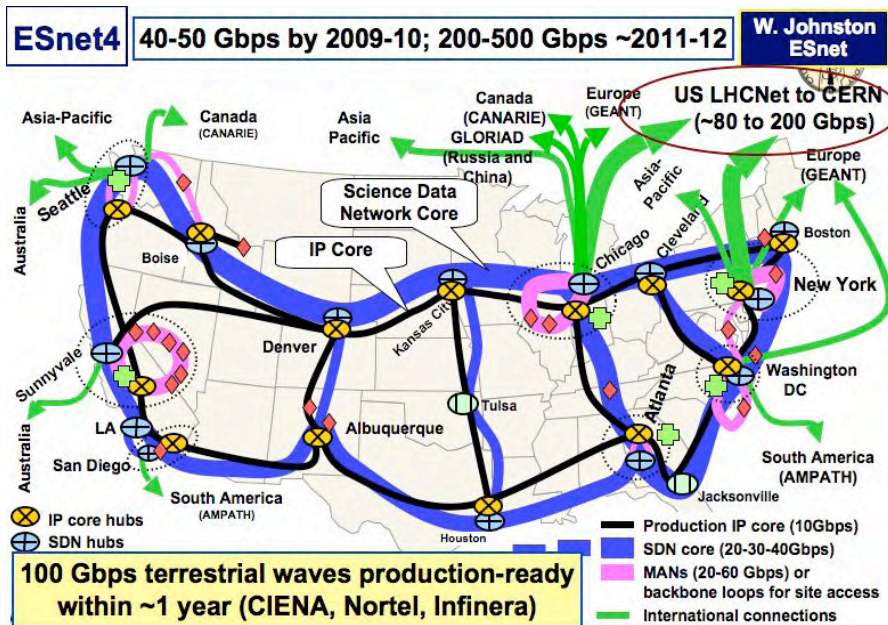
Recently the folks at ESnet have been doing projections of demand that show that climatology, fusion energy and bioinformatics will have similar demands to those of high energy physics within the next few years, and this is becoming an issue in their planning. They are concerned that despite their

forward-looking planning, they might find that the networks are not quite sufficient to satisfy the needs.

I already mentioned that historically, network utilization in high energy physics has grown at a rate between a few hundred and a thousand times per decade. ESnet has kept measurements of traffic that it has accepted into its network over time. Basically it has grown by a factor of 10 every 47 months on the average. You can see each factor of 10 as well as how many months it took historically on a semi-log plot. At the very end you see some of the points rising markedly higher. It is too soon to know whether this is a deviation from the norm.

So the scale of US LHCNet (shown in purple) and ESnet (in blue and black) have evolved together as the requirements of the LHC program have evolved. The black line shows the general purpose backbone of ESnet, which is currently 10 Gbps. The thick blue line shows ESnet's Science Data Network, which has been provisioned specifically to support the largest data flows, such as those generated by the LHC program.

Given the present state of technology, it is hard to make a road map that includes the technologies we will use that



will grow by quite this much. Nevertheless, historically this is what we have experienced. An increase of about 10 times every 47 months means about 300 times in a decade. You may remember that some of the other roadmaps that I showed were about 1000 times in a decade.

The slide ESnet4 (bottom page 5) is one of the first pictures that shows the transition around 2011-12 to 100 Gb wavelengths on the ESnet foot-print. Although this is just a drawing done in late 2008, it illustrates what will be realistically possible. The slide also shows a match with the US LHCNet outlook, shown in the arrows in green going off to CERN. The bandwidths shown for this period

are between 80 and 200 Gbps.

The long-term roadmap is shown in the HEP Bandwidth Roadmap illustration below. In 2001 I first created this private roadmap to help me answer "where am I going with all this?" When I first made it, the evolution in bandwidth was a bit faster but I've since moderated it in the out-years, to match what is likely to be affordable while meeting the needs as best we can. The first few steps have been close to what was originally foreseen, and for the last three years I haven't changed anything. In 2002 I gave the private roadmap to one of my colleagues who was then the head of computing at Fermilab, and he

showed it in a plenary talk at the big conference of the year in Amsterdam. I was in the audience and a bit shocked, but after a few seconds I realized he had done me a big favor by getting the roadmap out into the public view.

COOK Report: This must've allowed you useful feedback?

Newman: Indeed it did. What usually happens is that the first people who see my roadmaps say you know this is insane; clearly impossible. Then later they come to understand that these estimates are about right. And then when reality sets and what usually happens is they say these estimates are pretty good; they're the best available; but, given what we have seen occur, they were probably a bit too conservative.



HEP Bandwidth Roadmap for Major Links (in Gbps): US LHCNet Example



Year	Production	Experimental	Remarks
2001	0.155	0.622 -2.5	SONET/SDH
2002	0.622	2.5	SONET/SDH DWDM; GigE Integ.
2003	2.5	10-20	DWDM; 10 GigE Integration
2005-6	10-20	2-10 X 10	λ Switch; λ Provisioning
2007-8	3 to 4 X 10	~10 X 10; 100 Gbps	1 st Gen. λ Grids
2009-10	6 to 8 X 10	~20 X 10 or ~2 X 100	100 Gbps λ Switching
2011-12	3 to 5 X 40 or 2 X 100	~10 X 100	2 nd Gen λ Grids Terabit Networks
2013-5	To Terabit	To MultiT bps	~Fill One Fiber

Paralleled by ESnet Roadmap for Data Intensive Sciences

This reminds me of an anecdote about Oppenheimer. It is said that he attended a seminar at Los Alamos. At the end he went up to the speaker and said "you know I really enjoyed your talk: your ideas are crazy; buy they're not crazy enough to be correct."

The column on the left of the slide shows "production networks." What you see in yellow has already been implemented. Experimental networks mean that every year

or two you do exercises in order to prepare the network of the next generation. The point of these exercises can mean achieving throughput on a new scale, or working with new devices like we did recently with our Ciena multiplexers, before we put them into production service. As we demonstrated at this conference, we are indeed working with 10 x 10 Gb links and with 100 Gb links in experimental mode. These links are what we would like to put into production within a couple of years. On the right what you see is a kind of shorthand for some of the technical accomplishments of those periods for example the first use of a full gigabit Ethernet link in 2002, or the recent move to more dynamically provisioned networks.

Lambda Switching Advancements – VCAT and LCAS – FDT Gets More out of Each Lambda

The dynamism of lambda switching has been exciting. We are not actually switching full Lambdas at this point we are switching virtual circuits within them by means of emerging protocols like VCAT (virtual concatenation) and LCAS (a scheme for dynamically adjusting the size of the channels without breaking the connection). Here you are doing time division multiplex-

ing, and bonding sets of timeslots so they appear to users to be a single network channel. We use this in production in our Ciena multiplexers now.

In the future we foresee having dynamic networks of this sort but on a much larger scale, and somewhere in the middle of the next decade we foresee reaching terabit per second networks.

The slide Major Advances in data Transfer Applications below explains the process that we went through in order to understand how to get high throughput on our links. We learned how to do this progressively. The ramp up to the point where we could get

full use of a 10 Gb link was done by about 2005. For a while it was limited by the capability of the computer bus. PCI buses had a theoretical capacity of about 8.5 Gb per second and network interfaces would give you 7 to 7.5 Gb per second of sustained throughput. Then PCI express came along and we learned how to get nearly 10 Gbps flows with a single pair of servers and a single pair of interfaces - altogether cheap systems. And at this point the limitation became the reading and writing speeds of disks.

In 2005 my colleague Iosif Legrand wrote a Java application called Fast Data Transfer (<http://monalisa.cern.ch/FDT>)



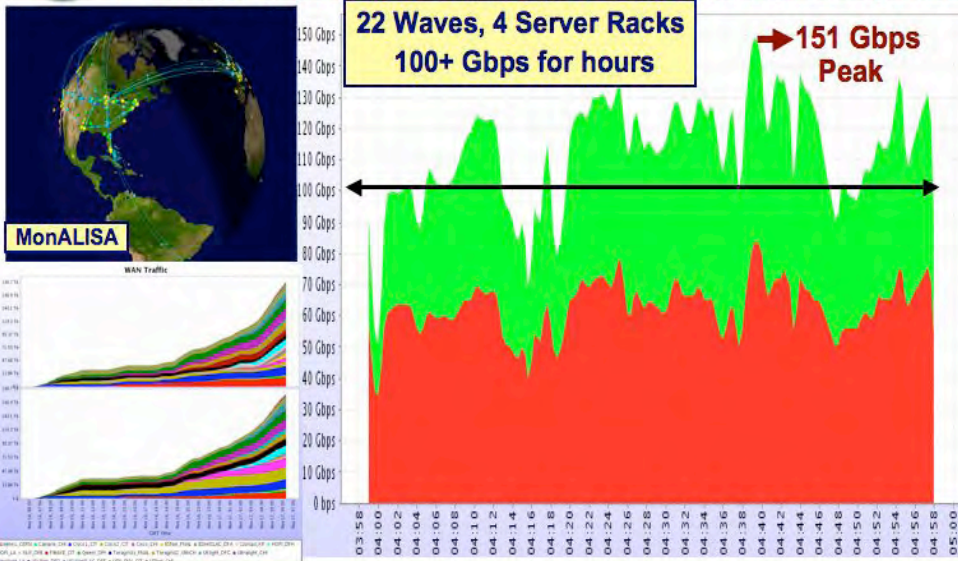
Major Advances in Data Transfer Applications Led by HEP with Computer Scientists and Network Engineers



- ◆ **2000-2007: HEP with computer scientists and network engineers developed the knowledge to use long distance networks efficiently, at high occupancy, for the first time**
 - **“Demystification” of large long range data flows with TCP:**
 - ➔ **2004-2005: Up to 10 Gbps per flow;**
 - ➔ **One to a few server-pairs matches a 10 Gbps link**
 - ➔ **Aggregate from 23 Gbps (SC03) to 151 Gbps (SC05)**
 - ➔ **Major advances in the TCP stack (e.g. FAST TCP), Linux Kernel (2.6.26), end system architecture, network interfaces and drivers**
- ◆ **2006-2008: Moved to mature storage-to-storage transfer applications; transfers among storage-systems**



Bandwidth Challenge at SC2005 (Seattle): Caltech, CERN, FNAL, BNL, SLAC, UM, UF, ESNet, I2...



475 TB Total in < 24h; Sustained Rate of 1.1 Petabyte Per Day

↳ that runs on any major platform, and achieves new levels of stable throughput. FDT does much more than just deal with the protocol. It can measure the capacity of an end-to-end network path in real time. It takes a large set of up to thousands of files, and fills large buffers with them, while sizing the buffers to match the capacity of the network path. It looks at the end systems in order to ascertain their capability. It sees how many disks there are and assigns one thread per disk. And it sends the buffers to the network at a rate which is compatible with the measured capacity of the link along the whole path.

To summarize -- in 2005 we started to move away from memory to storage systems,

since high speed data transfers are the what is crucial for us to do our science.

The slide above, Bandwidth Challenge at SC2005, shows the first bandwidth challenge with a goal of reaching 100 Gb per second. We were in Seattle, which is a great place to get access to multiple networks. There were some thirty-six 10 Gbps waves coming to the show floor and by contacting many of our friends and colleagues who run the networks we arranged to use 22 of them. There were actually two high energy physics booths that worked together: one from Caltech where we had 15 and one led by Fermilab and the Stanford Linear Accelerator Center (SLAC) that had seven.

This is the MonALISA plot showing the accumulated data flowing over each of the waves on the left, and the aggregate flow on all of them into and out of the exhibit floor on the right. We hit another benchmark of being able to reach a sustained data flow that was the equivalent to about a petabyte of data a day.

COOK Report: And the data flow involved moving data from one set of storage devices to another?

Newman: We had a lot of sites involved – Florida, CERN Caltech, Brookhaven and others, and at this stage the data transfers were mainly memory to memory. We also had four full racks of servers to reach our benchmarks. Some of the equipment worked very well; and some not so well.

Now Slide SC06 BWC: Fast Data Transfer at the top of the next page shows the results we obtained with the FDT application at Supercomputing 06 in Tampa, Florida. Iosif Legrande written the FDT application specifically for this conference, and its first version was ready for testing just about six weeks before the conference.

We tried it out and the results were pretty amazing. Each one-rack unit server (which is




SC06 BWC: Fast Data Transfer

<http://monalisa.cern.ch/FDT>



- ◆ An easy to use open source Java application that runs on all major platforms
- ◆ Uses asynch. multithreaded system to achieve smooth, linear data flow:
 - Streams a dataset (list of files) continuously, through an open TCP socket
 - ➔ No protocol Start/stops between files
 - Sends buffers at rate matched to the monitored capability of end to end path
 - Use independent threads to read & write on each physical device
- ◆ Secure: Can "plug-in" external AAA APIs from major projects

- ◆ **SC06 BWC: Stable disk-to-disk flows Tampa-Caltech: 10-to-10 and 8-to-8 1U Server-pairs for 9 + 7 = 16 Gbps; then Solid overnight. Using One 10G link**
- ◆ **17.77 Gbps BWC peak; + 8.6 Gbps to and from Korea**



Outlook for 2007: ~70-100 Gbps per rack of low cost 1U servers

one and three-quarter inches high in a rack) with 4 disks inside could read and write at 200 Mbytes/sec, which matches the speed of the disks, even when reading and writing across long range networks. As you increase the number of servers, and thus the number of flows, you can see in the slide that each flow remains pretty stable. The flows don't fight each other.

On the right-hand side the top bar chart (labeled "I Legrand") shows the startup of the flows, while the chart underneath it shows the results we had when we left all the flows running during the overnight period. This is 15-16 Gb per second, by the way, on one 10 Gb link. For

some reason other groups at SC06 had not thought of using the link in both directions, or perhaps (without FDT) they were unable to get an increase in the overall throughput.

There is a problem with TCP where you get a backpressure from the acknowledgments. However with the FDT application, where we are not trying to send data as quickly as possible but are instead sending the buffers at a rate that we know the link can take, the back pressure doesn't disrupt anything. We could reach nearly 18 Gb per second on a single 10 Gb per second link by having 10 pairs of servers transferring data in one direction, and another 8 pairs transferring

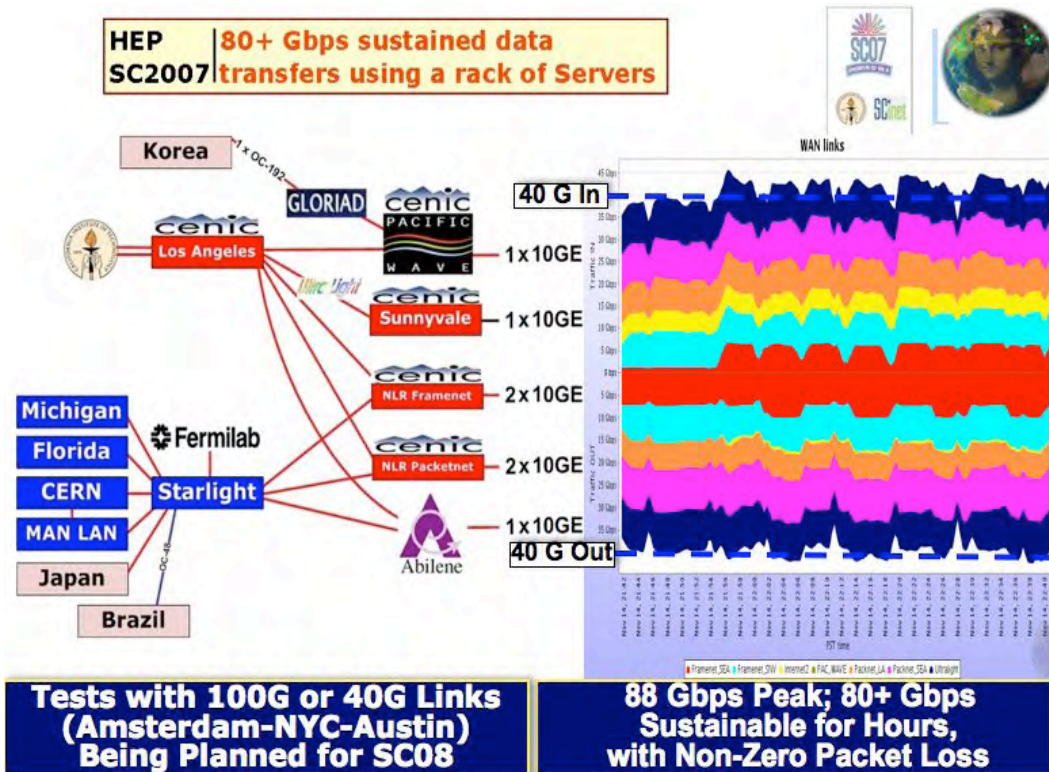
data in the other direction.

COOK Report: So to the people whom you're asking for funding you are demonstrating that you'll be improving the technology in very interesting ways?

Newman: Correct. We are doing two things. First **we have designed networks where we intend to use them at a very high occupancy rate, such as**

80% or even 85%. We don't know exactly at which level of use in a production environment we will need to back off. This is nothing like standard networks. But with our virtual circuit channels we can do this by using a set of channels of nearly any dimension inside the 10 Gb link to support a variety of medium and large flows. **We can use the remainder of the link for general-purpose networking.** We could, for example provision dedicated channels adding up to a capacity of eight Gbps, and leave the remainder for general-purpose networking.

COOK Report: And by general-purpose networking



Tests with 100G or 40G Links (Amsterdam-NYC-Austin) Being Planned for SC08 **88 Gbps Peak; 80+ Gbps Sustainable for Hours, with Non-Zero Packet Loss**

servers we could simultaneously sustain 40 Gbps in and 40 Gbps out. We therefore had shown that we could fully match a next-generation link, and were ready to use them in production.

Then we started talking about how to do 40 and 100 Gb per second networking at SC08. You can see that using FDT and several waves we could sustain 80 gigabits per second in both

you mean what? Web and e-mail?

Newman: Yes. Web access and e-mail and all the usual traffic. **But when it comes to this, like Surfnets, we do not do much at layer 3. We avoid routing wherever possible and as explained earlier don't have what you would call a "carrier class" router. But we do routing in and out of our domain, and have peerings with all the major research and education networks at our points of presence. Our own network is nearly all done at layer 2, and at the core it is layer 1.** We have a software framework with which we build optical virtual circuits, which has some simi-

larities with Surfnets; except that Surfnets does not do any dynamic configuration at the optical layer (layer 1) using VCAT and LCAS channels.

Overall, our results at SC06 showed that it took very little equipment to fully match a 10 Gbps link. From these results we concluded that if I had one rack of servers running FDT, I could reach 70-100 Gbps per second. So this became our goal for SC07.

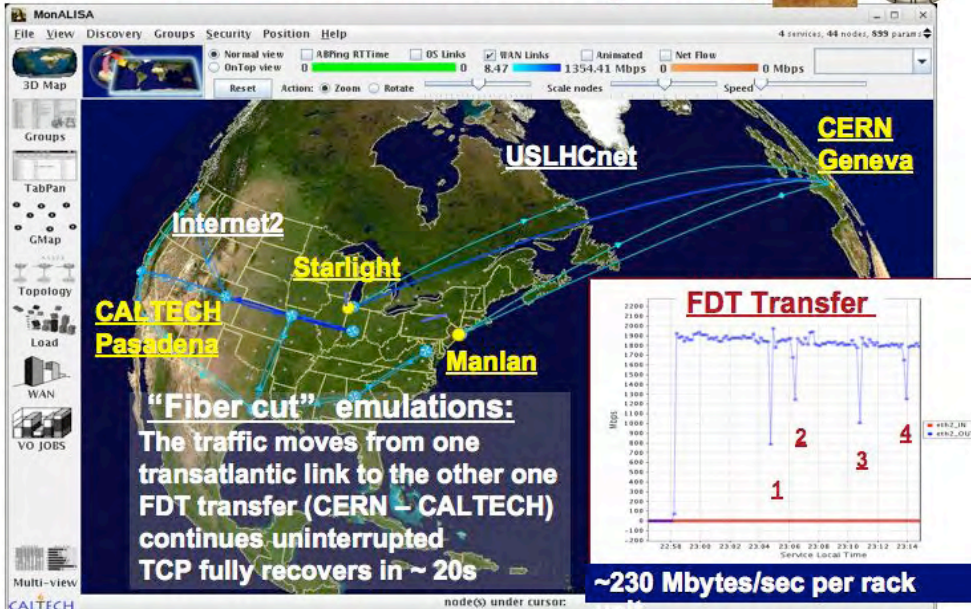
The HEP SC2007 slide above shows us essentially reaching our milestone at SC07. Reno Nevada was not as good a place as Tampa to do networking. At Reno we essentially had access to all the waves that existed and we showed that with one rack of

directions for hours. **And we also found that by putting our application FDT together with Steven Low's FAST TCP, that we could sustain these rates even in the presence of some packet loss. Until that time, any significant packet loss with TCP over long-distance had meant death to large flows. But with this TCP stack, and this application, even with up to a few percent packet loss we still got extremely high throughput.**

COOK Report: With packet loss, the problem is you have to wait for the acknowledgment to come back and it doesn't come?



FDT Automatic Path Recovery: Fiber Cut Emulations (4/2007)



matches the speed its four disks. You can actually see that as the disk head moves in from its outer edge, the data rate passing the disk head decreases slightly and so the transfer rate also goes down. Overall, you now have wide area transfers going on at speeds equal to what you could do over the distance of a meter.

The next part of this exercise was, since we have multiple links across the Atlantic, to find out how the system would react if we cut one of the links. When we did emulate a link cut, the MonALISA service automatically restored the connection over a different physical path within a second or so, namely in a time short enough that the TCP session was not interrupted. All that happened is that it took just 20 seconds

Newman: Right. What we found is that, with the way this TCP stack was tuned, the recovery was sufficiently fast to allow it to just keep going. It had a number of threads and packet loss would affect only one at a time.

by checking the end systems' and the network path's capability, before sending the data. When the transfer starts, the trace in the lower right of the slide shows that the single 1U server used produces a smooth data flow of 230 MB per second, which

As shown in the FDT Automatic Path Recovery slide above, we also had developed services in our MonALISA system that could dynamically set up a circuit across the network between two sites, send an entire dataset, then tear down the circuit when the transfer was done.

The slide illustrates a transfer test with a single server pair, where we emulate cuts on our transatlantic circuits to see how the network and the transfer application behave. The MonALISA service starts



MonALISA: Monitoring Grids, Networks, Compute Nodes, Running Jobs, Processes

- Built for speed and global scale: 5k messages/sec/server; multi-threaded engine schedules ML services
- Autonomous agents auto-discover and collaborate in real-time for a variety of tasks

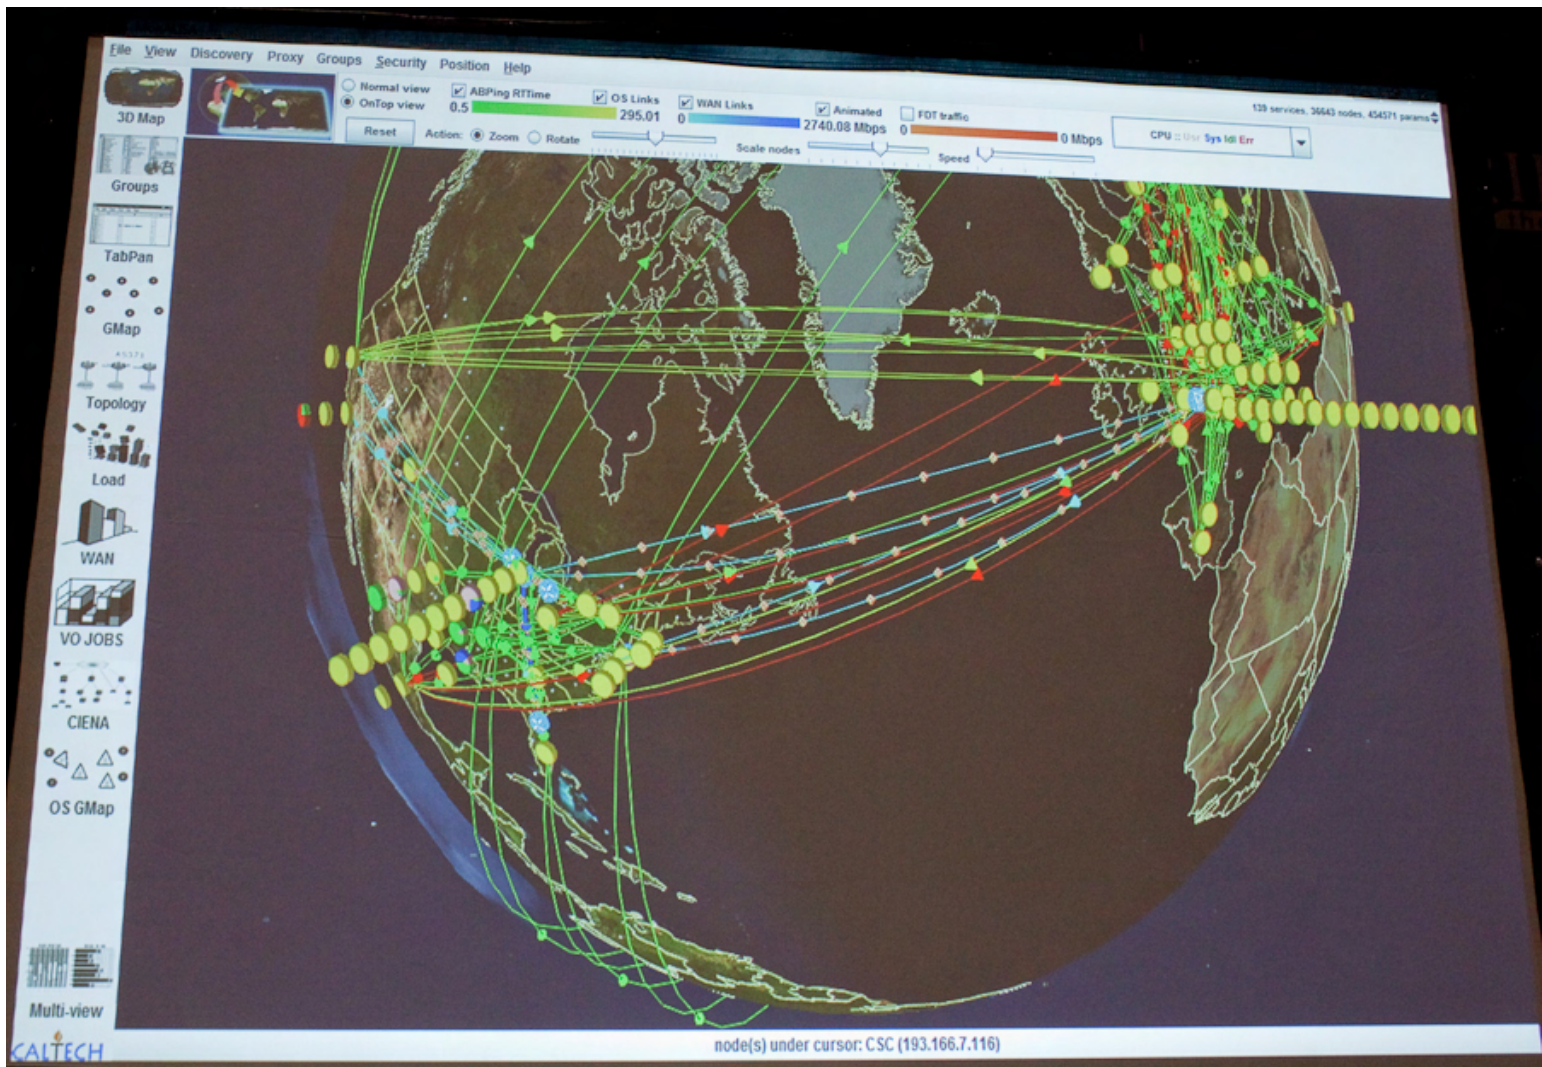
CENIC Innovation Awards '06 and '08
Internet2 IDEA Award '07



- MonALISA Today**
Running 24 X 7 (5 Years)
Now at 360 Sites
- Collecting > 1,000,000 parameters in near real-time
 - Update rate of >20,000 parameter-updates per sec
 - Monitoring
 - * > 40,000 CPUs
 - * > 100 WAN Links
 - Many Thousands of Grid jobs running concurrently

- Major Communities**
- ALICE
 - EVO
 - OSG
 - CMS
 - DO
 - STAR
 - LCG RUSSIA
 - SE Europe GRID
 - APAC Grid
 - UNAM Grid (Mx)
 - ITU
 - Internet2
 - ULTRALIGHT
 - GLORIAD
 - US LHCNet
 - RoEduNET
 - Enlightened

A Sampling of End To End Monitoring and Global views



for the TCP protocol to fully recover, and then the transfer continued as before.

The slide below describes MonALISA. This slide is covered up by the words but if you have the animated version of it you could see underneath the word boxes a whole bunch of pictures illustrating the diverse of services and applications provided by the MonALISA system. These include the view of networks on the global display you saw at the Caltech booth at SC08; the net throughput between

any two sites shown in a histogram; the network topology that is derived in real time among all the routers and switches; the job-lifelines that are the steps in any given computer showing how long it takes to get from one step in a running program to the next; and job statistics for each of several different organizations, accompanied by pie charts showing the division of resource usage.

Editor's Note: The picture above is one that I took of Harvey's laptop in Austin. It

shows what he has described in the preceding paragraph.

COOK Report: Could MonALISA be thought of as an operating system for this kind of private optical network?

Newman: Yes. It is a global distributed system meant to support network operations and grids. It can also profile the end systems and it even has been used in running applications as a kind of super spy to show everything that goes on as the application progresses.

COOK Report: If you look at existing network tools, is there anything comparable to MonALISA in an earlier generation or in another context?

Newman: I don't think so. I think there are large systems that have a huge amount of code and are run by many many people. But this is an artificially intelligent system that has a near-zero operational burden because everything is automated.

COOK Report: If the climate people get their own very large networks is it likely that they will be able to borrow what you have developed to help them run their own systems?

Newman: They could. It means gathering data on a very large scale from a large number of sources. This is very simple in the sense that there are no rapid strategic operations. Namely you don't have to change configuration or workflow in any fundamental way to respond to a dynamic load that you did not really predict or determine beforehand.

There are 340 implementations of MonALISA running, and the operational burden does not grow significantly with each new installation.

COOK Report: Do you have an implementation of

MonALISA then the running in virtually every node in your LHC network?

Newman: Yes absolutely. We have agent-based services that gather information from all the switches and routers and the loading on all the links. They also can follow the source and destination, as well as the progress of every data transfer in the network. In the case of FDT, MonALISA agents profile what is happening within the end-systems and react in case a system gets too loaded. All of that. It is a large intelligent monitoring, command and control system.

These Networks Need Dynamically Provisioned Circuit Overlays

In the following slide Bill Johnston, the former ESnet manager, describes his views on the need to have dynamically provisioned virtual circuits across these kinds of networks, in order to provide guaranteed quality of service for a designated set of data flows. The ability to provide their users with useful and relevant network monitoring operation on an end-to-end basis is another important requirement.

Now people have adopted these guidelines. I would



W. Johnston, ESnet Manager On Circuit-Oriented Network Services



- In order to
 - facilitate operation of multi-domain distributed systems
 - accommodate the projected growth in the use of the network
 - facilitate the changes in the types of traffic

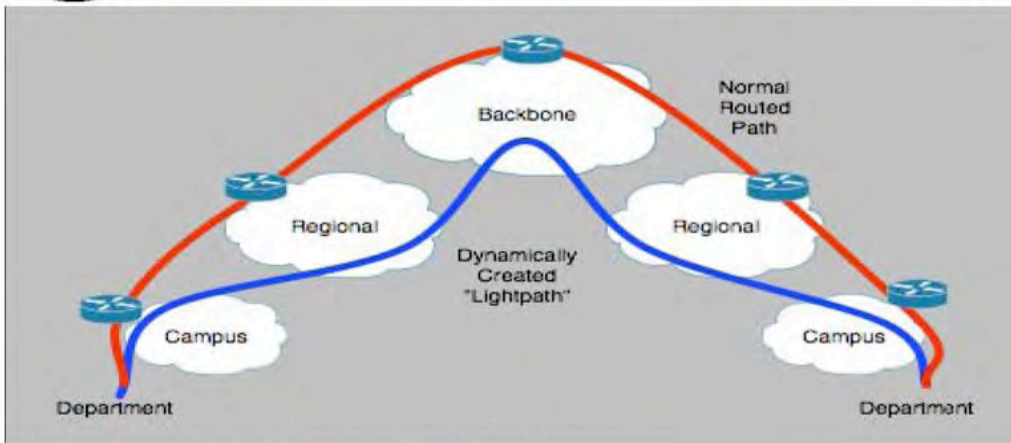
the architecture and services of the network must change

- ➔ **The general requirements for the new architecture are that it provide:**
- 1) Support the high bandwidth data flows of large-scale science including scalable, reliable, and very high-speed network connectivity to end sites
 - 2) Dynamically provision virtual circuits with guaranteed quality of service (e.g. for dedicated bandwidth and for traffic isolation)
 - 3) provide users and applications with meaningful monitoring end-to-end (across multiple domains)

Bandwidth Guarantees; Traffic Isolation; E2E Circuits With Well-defined Characteristics; Bandwidth Scaling/Rescaling
These features cannot (adequately) be addressed by standard IP networks



Internet2's New Backbone with Dynamic Circuit "Light Paths"



- ◆ Transition essentially completed in 2007
- ◆ Initial deployment – 10 x 10 Gbps wavelengths over the footprint
- ◆ First round maximum capacity – 80 x 10 Gbps wavelengths; expandable
- ◆ Scalability – potential migration to 40 Gbps or 100 Gbps waves
- ◆ Dynamic circuit network (more later) deployed last Fall; available to US Tier2s and Tier3s with backhaul to protect the campus nets

for the overlay of these dynamic circuit Lightpaths.

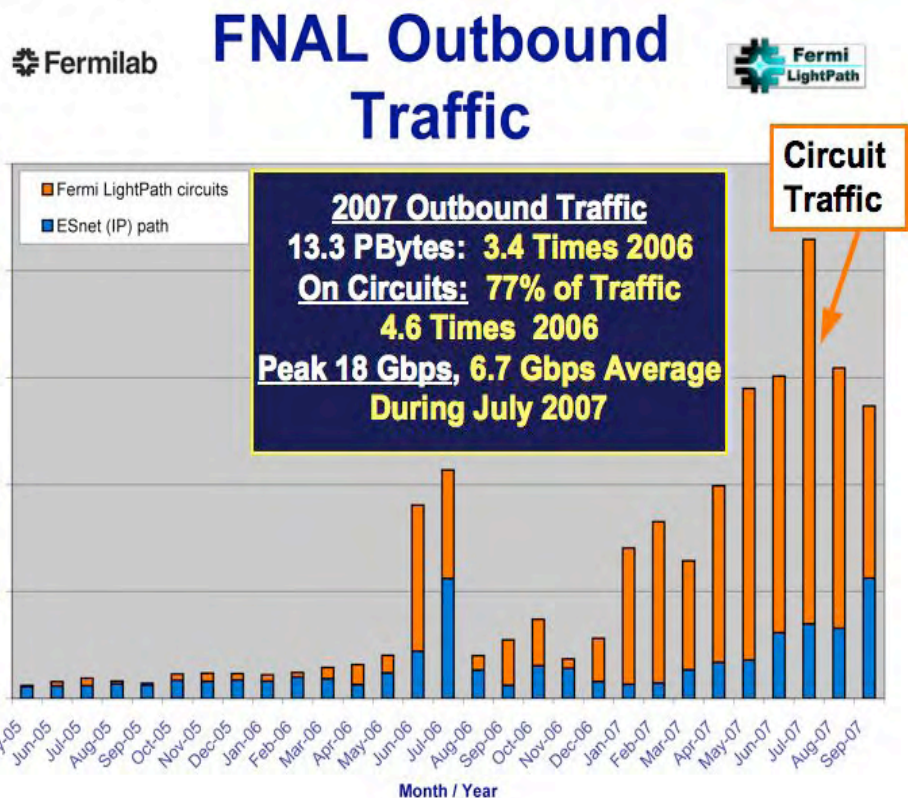
As in the case of US LHCNet, the Internet2 instantiation of a dynamically allocated overlay lightpath network is different from that used by Surfnet in that they also use VCAT and LCAS to channelize the bandwidth. Within these 10 Gb links we can provision channels as small as 50 Mbs per second, or as large as the entire link.

claim in addition that the last point (on monitoring) requires these globally distributed systems. Only systems like MonALISA are able to do the monitoring on an end-to-end basis, react to problems and deal with them automatically, and avoid having to grow the staff required to look after the system as the level of network usage increases.

will meet the users' needs. Internet2's new backbone has been given the capability

COOK Report: This permits a more effective use of your constrained resource? Correct?

Johnston goes on to say that without such a framework of dynamically provisioned virtual circuits, you cannot run the networks in a way that





ICFA Standing Committee on Interregional Connectivity (SCIC)

- ◆ **1996 Visionary Statement of ICFA: Our Major Collaborations Should Gear Themselves for Remote Participation**
- ◆ **1997 Network Task Force [Dedicated to David O. Williams]**
- ◆ **SCIC Created in July 1998 in Vancouver. Charge:**
- ◆ **Make recommendations to ICFA concerning the connectivity between the Americas, Asia and Europe**
- ◆ **As part of the process of developing these recommendations, the committee should**
 - **Monitor traffic on the world's networks**
 - **Keep track of technology developments**
 - **Periodically review forecasts of future bandwidth needs, and**
 - **Provide warning of problems**
- **2008 Reports and Presentations:** <http://cern.ch/icfa-scic>

Newman: Yes. I never liked a situation where all you could do was slice up the 10 Gb link into fixed 1 Gb links, since that is rarely, if ever, a way to use the available network bandwidth efficiently.

The slide at the bottom of the previous page shows that most of the large flows coming out of Fermilab run over dedicated circuits of the type described above. **Circuits are now taking hold as the main mechanism supporting large flows.**

Digital Divide Issues

Newman: In 2002 I became chair of the ICFA technical panel that oversees international networking for the field of high energy physics, as described in the slide above.

ICFA is the **International Committee on Future Ac-**

celerators. It was put together as a forum where the leaders in the field – in particular the directors of the accelerator laboratories – could come together and discuss both the present status and future of the field and any other related major issues. Apart from accelerators and reviews of the physics program, they sometimes took on major issues which

they dealt with through technical panels.

For example one program deals with detector instrumentation schools. It is like an outreach activity. They have workshops that train students on how to use detectors and educate them on the capabilities of different detector instrumentation. There is another one on advanced developments in accelerator physics.

In 1996 their discussions led them to conclude that the field would be dominated by larger and larger collaborations. Consequently they issued a visionary statement that said **the major collaborations in high energy physics should organize themselves to fully support remote participation.** They then commissioned a task force on networking, to review the status and outlook



SCIC in 2007-8

<http://cern.ch/icfa-scic>

Three 2008 Reports: An Intensive Year Rapid Progress, Deepening Digital Divide

- ◆ **Main Report: "Networking for HENP"** [HN, D. Nae et al.]
 - ➔ **Includes Updates on the Digital Divide, World Network Status; Brief updates on Monitoring and Advanced Technologies**
- ◆ **34 Appendices: A World Network Overview**
Status and Plans for the Next Few Years of Nat'l & Regional Networks, HEP Labs, & Optical Net Initiatives
- ◆ **Monitoring Working Group Report** [R. Cottrell]

Also See:

- ◆ **TERENA (www.terena.nl) 2007 Compendium:**
In-depth Annual Survey on R&E Networks in Europe
- ◆ **<http://internetworldstats.com>:** Worldwide Internet Use
- ◆ **<http://websiteoptimization.com/bw>:** Broadband (the new Digital Divide)
- * **SCIC 2003 Digital Divide Report** [A. Santoro et al.]

for usage of networks by our field. After the task force had reported, it was clear that this would be an ongoing issue. Consequently they created the Standing Committee on Inter-regional connectivity (SCIC) to oversee the status and requirements for networks, in and among the different world regions. The charge to the committee, which presents its results annually to ICFA in a series of extensive reports, is summarized in the slide.

From the point of view of the US and Europe they could see that one of the prime points of consideration was that these collaborations had become truly global in nature. I served on the SCIC committee from its inception, and when I took over as the chair in 2002, we began to try to divine what the most significant issue was that we as a committee could tackle. We naturally came to the conclusion that the growing digital divide among regions within the scientific community was just such an issue. Many regions of the world, from central and southeast Europe to south and southeast Asia to Latin America, had much less connectivity, and a correspondingly lesser ability to participate effectively in these collaborations.

On the slide "SCIC in 2007-8" on the preceding page, you see my introduction of our

2008 reports, presented to ICFA in February 2008.

Every year we do a complete set of reports the locations of which are mentioned in the slide above. In our reports we deal with all the major research and education networks and some of the major R&D projects in networking. For each of them we ask for an update which we include in a short appendix which can be anywhere from 2-10 pages long. [**Editor's Note:** Harvey released the 2009 SCIC Reports, including 39 Appendices and his presentation, on Feb 10. Fascinating – You can find the full set of documents and the slides at:

<http://monalisa.caltech.edu:8080/Slides/Public/SCICReports2009Final>]

The monitoring part is put out by Roger ('Les') Cottrell. He leads the Internet End-to-end Performance Monitoring (IEPM) activity at SLAC. Using a worldwide ensemble of monitoring servers and

nearly 1,000 monitored computers, he checks the packet loss rates and the variation in round-trip times, and derives the throughput capabilities along different network paths between the US or Switzerland, and more than 150 other countries.

COOK Report: Then this is somewhat like an encyclopedia that allows you periodically check the network capability throughout this very huge and complex global system?

Newman: Indeed. It also allows us to formulate high-level statements about the field, and it provides a quantitative roadmap, showing the state and the evolution of the digital divide. Especially which regions are, or are not making progress relative to other regions.

The side below is our high-level mission statement, which has not changed since 2002. It says that we should



SCIC Main Conclusion for 2008

- ◆ ***As we progress we are in danger of leaving the communities in the less-favored regions of the world behind***
- ◆ ***We must Work to Close the Digital Divide***
 - ➔ ***To make physicists from all world regions full partners in the scientific discoveries***
 - ➔ ***This is essential for the health of our global collaborations, for our field, and for the world community***



Work on the Digital Divide from Several Perspectives

- ◆ **Raise Awareness: Locally, Regionally & Globally**
 - ◆ **Digital Divide Workshops: Rio, Daegu, Cracow/Sinaia, Mexico City**
 - ◆ **Diplomatic Events: WSIS, RSIS, Bilateral: e.g. US-India**
- ◆ **Technical Help with Modernizing the Infrastructure:**
 - **Provide Tools for Effective Use: Data Transport, Monitoring, Collaboration**
 - **Design, Commissioning, Development**
 - **India "Knowledge Network", and Link to US, Japan and Europe**
- ◆ **Share Information: Monitoring, Tracking BW Progress; Dark Fiber Projects & Pricing**
 - **Track Planning (focus on LHC) and Leading Edge Progress**
 - **Model Cases: Brazil, Poland, Slovakia, Czech Rep., China ...**
 - **Encourage Access to Dark Fiber; Modern technology choices**
- ◆ **Encourage, and Work on Inter-Regional Projects**
 - **Latin America: CHEPREO/WHREN (US-Brazil); RedCLARA, EELA**
 - **Mediterranean: EUMEDConnect; Asia-Pacific: TEIN2**
 - **GLORIAD Global Optical Ring**

go after closing the digital divide.

There are several ways in which we can help close the divide, as summarized in the slide above. We communicate with the governments in each of the countries to tell them about the enablement for their societies that will occur if they support these activities - not to mention modern network infrastructure for their own countries. We show them where they stand in terms of their connectivity relative to other countries and other regions. This is a program of education through which we work by means of workshops and other activities.

We help them to understand how to get access to dark fiber, which in some cases such as Brazil and Romania has led to great improvements over the last few

years.

The latest workshop in October 2008 in Mexico City was very valuable. I think it helped bring about some progress in Mexico despite a terrible monopolistic situation there with the national carrier. We were thinking of having one in Russia in 2008 but we usually try to do this only when there is a hope of positive change. Things were looking favorable there for a while, but in 2008 looked rather less so. Now the outlook is looking more favorable again so we will consider a workshop there in 2009. Another site we are considering is Argentina.

COOK Report: What happened to Gloriad? I think it got chopped in two?

Newman: Well there remains a relatively low bandwidth link (155 Mbps) at the

border between China and Russia, which is in a very isolated area. Relatively good network links exist in the Moscow area, as well as between Moscow and Europe, and these links are expected to progress to 10 Gbps in 2009. But across the country progress has really slowed down, and I have not heard much lately.

We also participated in the World Summit on Information Society and in various bilateral meetings. I went to a whole series of bilateral meetings between the US and India and then participated in the Indian Knowledge Network initiative when we were called upon to show high throughput into and out of India.

About every 18 months we have a conference on Computing in High Energy Physics (CHEP) and in February of 2006 we helped arrange to have the conference in Mumbai at the time of our network demonstrations. The President of India gave the introductory speech and we did a demonstration of high-speed networking. We managed to arrange a donation from VNSL - they didn't really consider it a donation we later found out - of four STM-1 links which at the time was the most total bandwidth that you could get on the SeWeMe3 cable. We used

those pretty efficiently and achieved by far the most dataflow ever going in and out of a site in India at one time: more than 15 Terabytes received over a period of a couple of days.

COOK Report: Who in India conceptualized the Knowledge Network?

Newman: Sam Pitroda. I also have worked on occasion during the last several years with Anil Srivastava. We had a very active meeting right after the CHEP 2006 conference and at the level of the president the policy push for investment in a much higher speed (ideally 10 Gbps) in-

ternational link to India was quite high. But the bureaucracy became almost impossible to overcome.

Sam Pitroda is in the midst of things and is doing all he can. Over time it seems that the initial goals have become more modest, but progress is being made. The Prime Minister approved the Knowledge Network in January of 2007, but it was only last Fall that I had a visit from the two people, one from an academic institution and another from a network organization, who were suddenly assigned to get the Knowledge Network launched and the first 50 institutions connected

within the next two months. I recently heard that they had indeed successfully launched the Knowledge Network, and I hope to hear from them again soon with some details of their progress.

So in general in the SCIC we share information, monitor progress, track bandwidth and throughput improvements in each region. When presenting our results in any country, we try to remove any illusions they may have about how advanced or not-advanced they are relative to their peers in their region, and in different regions. We present model case studies, from countries that have

TERENA Compendium of National Research and Education Networks In Europe/Network

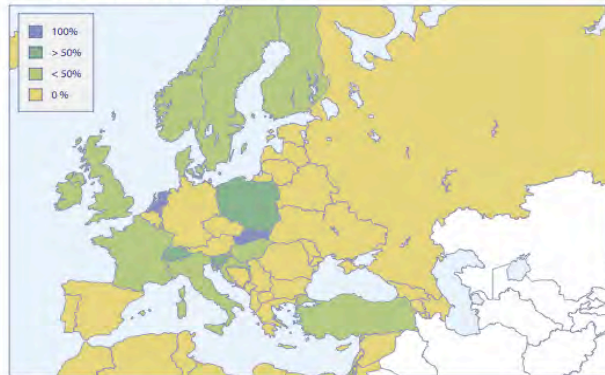
3A.7 Dark Fibre

Some NRENs own dark fibre or have IRUs² or lease dark fibre and can decide themselves what technology and what capacity to use on their fibre. NRENs were asked if they currently have IRUs or own dark fibre, or if they plan to get it during the coming two years. NRENs were also asked approximately what percentage of their backbone is dark fibre, in km, in point-to-point distances.

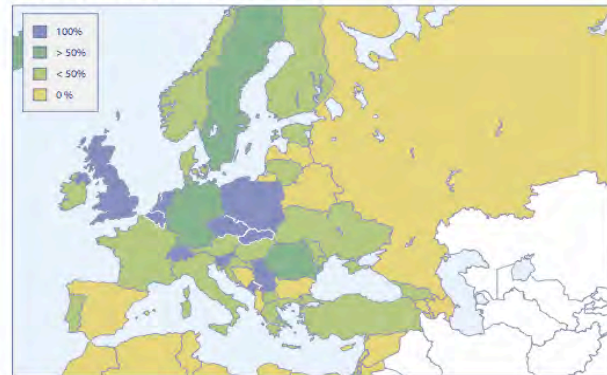
The maps below³ illustrate the rapid developments in this area in recent years. Many (although not all) NRENs predict a further increase in the percentage of their network that is dark fibre by 2010.

Legend: dark blue is 100% dark fibre; yellow is either no dark fibre or no information from that country for that year. Note that for 2008, 2007 data has been used for some countries that did not answer this year.

Map 3A.7.1 Dark Fibre on NREN backbones, 2005



Map 3A.7.2 Dark Fibre on NREN backbones, 2008



² IRU stands for 'Indefeasible Right of Use'. This is the effective long-term lease (temporary ownership) of a portion of the capacity of a cable. See, for example, <http://whatis.techtarget.com> for more information. The distinction between an IRU and a lease is becoming less clear; therefore, these two categories have been combined.

³ Idea developed by RedIRIS, Spain.

really forged ahead. Brazil, Poland, the Czech Republic, Slovakia and China are now among the leading examples.

We encourage access to dark fiber wherever it makes sense. This encouragement was especially transformative for Brazil where the leaders of the national research and education network RNP took up the charge, and deployed a new network covering the eastern half of the country, by bringing together many dark fiber footprints acquired from the utility companies and others.

And we also have helped them understand how to make modern technology choices, so that they could build these things in an affordable way.

Trans-European Research and Education Network Association) has picked up in its annual Compendium that for national research and education networks to succeed, an important element in the formula for success is access to dark fiber, so they can build their own infrastructures. In the past three years in Europe use of dark fiber has spread significantly. The page from the 2008 Terena Compendium is shown at the bottom of the preceding page.

But remember as to the possibility of doing this in United States that you're not allowed to give away what you can sell.

COOK Report: I think that is absurd.

Newman: I am inclined to agree with you.

Then there are also bilateral projects that help a lot. The US - Brazil projects were very valuable in transforming Brazil.

In Europe, there are several others, listed in the Terena Compendium. GLORIAD you know about. EUMedConnect is the only initiative providing connectivity at the level of 34 to 155 Mbps to countries in North Africa. TEIN2 (and now TEIN3) provides improved European - Asian connectivity.

COOK Report: Has anyone been able to go through your experiences with these cases and look for patterns of commonality that would allow the insights gained to be transported elsewhere? As for example to the United States?

Newman: Terena (the



SCIC Monitoring WG PingER (Also IEPM-BW)



R. Cottrell

Monitoring & Remote Nodes (2/08)

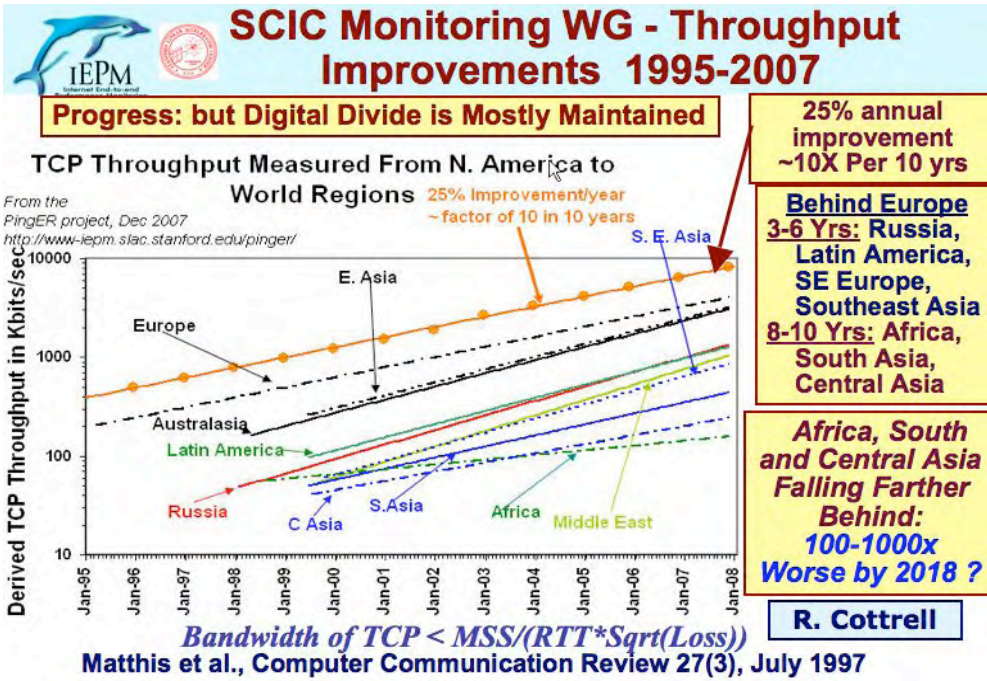
- ◆ **Measurements from 1995 On**
Reports link reliability & quality
- ◆ **Countries monitored**
 - ➔ **Contain 96% of world pop.**
 - ➔ **99% of world's Internet users**
- ◆ **800+ remote nodes at 675 sites in 155 nations; 40 monitoring nodes 87 Sites in 45 African countries**
- ◆ **Strong Collaboration with ICTP Trieste and NIIT (Pakistan)**
- ◆ **Excellent, Vital Work**



PingER Regions



Countries: N. America (2), Latin America (20), Europe (40), Africa (45), Middle East (12), Central Asia (9), South Asia (7), East Asia (4), Southeast Asia (9), Russia (1), China (1) and Oceania (4)



that are remarkably parallel. There are some cases in which one region is catching up with another. Toward the bottom of the slide you see that central Asia, south Asia including India and Pakistan, and Africa are really very far behind. It is also disturbing that the trend is towards them falling even farther behind.

There is a rather unpleasant conjecture that if things stay as they are, by 2018, some of the regions at the bottom of the graph will be two or three orders of magnitude behind the others. When you consider the present situation indicated in the slide, where India, Pakistan and central Asia are 8-10 years behind, and think back to the connectivity in the US 8-10 years ago, you can appreciate what this means. There is also a lot more information about

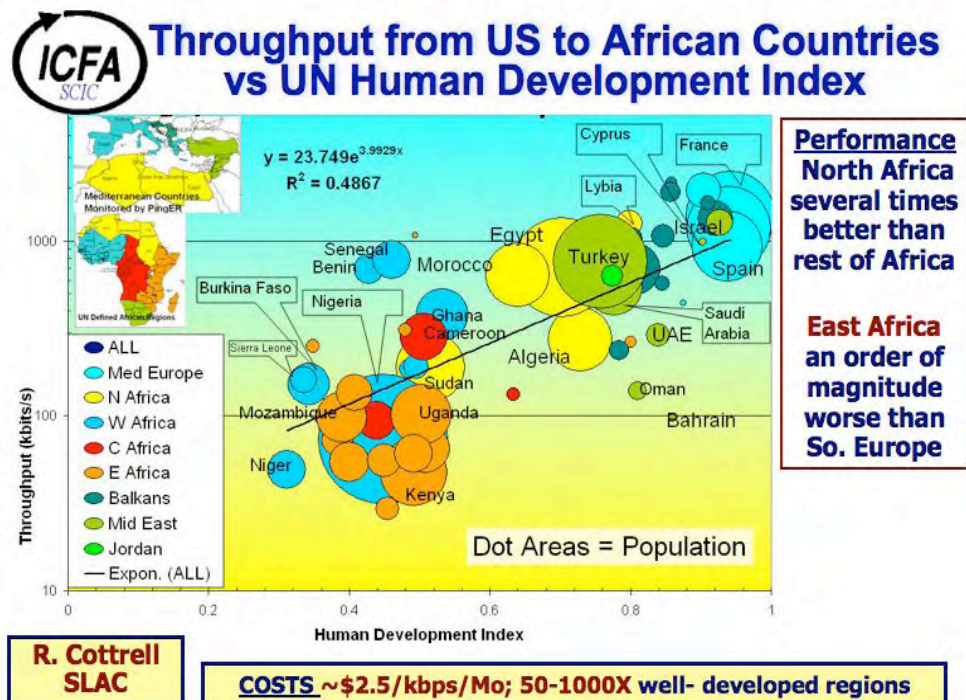
Monitoring Regional Connectivity

The slide at the bottom of the preceding page focuses on our SCIC monitoring project. It tells you how we have been monitoring a sparse sampling of connectivity in a huge number of countries.

To make this more understandable Les Cotrell has organized these countries into groups shown at the bottom. These are organized into so-called regions so that you can see what is happening in each.

COOK Report: The connectivity he's measuring is not connectivity on your high energy physics network but rather conductivity on the commodity global Internet yes?

Newman: Yes. On the slide Throughput Improvements immediately above you can see how these different regions are progressing. The straight lines are an idealization. If you look at the actual data, you see that they have more wiggles. Then if you apply some smoothing to the data you have straight lines



Africa among the slides, including the fact that the average cost per kilobit per second there is 50-1000 times higher than what we are used to. Also the computer density is much less which often means that many people are sharing a single computer.

As you can see the slide at the bottom of the preceding page has more information about the African situation. This one shows throughput from the US to some of these African countries, compared to some well-known international economic indicators. This one also shows the United Nations Human Development Index (HDI) which is on a semi-log scale where vertical differences are huge. You can see where different countries lie. The area of each dot is proportional to

the population of the country it identifies.

The different colors are different world regions. You can also see, as you would expect, that people who are low in the Human Development Index, which means access to economic welfare and education, are also very much lower in terms of throughput. North Africa and Turkey are somewhere in the middle.

Digital Divide Workshops – Putting Brazil on the GLIF

We had the first of these in Rio in February 2004. We discussed digital divide issues of course but also before the conference we had tutorials on advanced technology so that students could benefit from learning about our

methods and how to improve things. Before we started the first workshop we had a meeting in 2002 to discuss these issues, and to launch a Grid initiative in Brazil centered around the creation of Tier2 and Tier3 clusters in Rio de Janeiro and São Paulo respectively. This initiative, and the Digital Divide Workshop, have since served to bring Brazil into the fold, by making the physics groups there effective partners in LHC computing.

In Brazil we also had found out from Don Reilly about the existence of this dark fiber that Electrobras, the electric company had gone bankrupt and left behind some 15,000 km of unlit fiber. Our first thought of course was to try to figure out how to use the dark fiber.

Now the research and education network of Brazil was at first very suspicious and we had the impression that they were wondering what were we (North Americans) doing there? The first discussion was not terribly comfortable, but when we returned sometime in 2005 they showed us a reverse auction that they had carried out to acquire links on a fiber footprint throughout the country. Progress was rapid and by 2006 there was a 10 Gbps core forming a triangle between Rio de Janeiro, Brasillia, and São Paulo.



First ICFA Digital Divide Workshop UERJ, Rio de Janeiro, Feb. 16-20 2004



HEPGRID AND DIGITAL DIVIDE WORKSHOP

NEWS:


- Bulletin: [ONE](#) [TWO](#)
- WELCOME BULLETIN
- General Information
- Registration
- Travel Information
- Hotel Registration

Tutorials

- ◆ C++
- ◆ Grid Technologies
- ◆ Grid-Enabled Analysis Networks
- ◆ Collaborative Systems

Theme: Global Collaborations, Grids and Their Relationship to the Digital Divide
For the past three years the SCIC has focused on understanding and seeking the means of reducing or eliminating the Digital Divide, and proposed to ICFA that these issues, as they affect our field of High Energy Physics, be brought to our community for discussion. This led to ICFA's approval of the first Digital Divide and HEP Grid Workshop.
 More Information: <http://www.lishep.uerj.br>

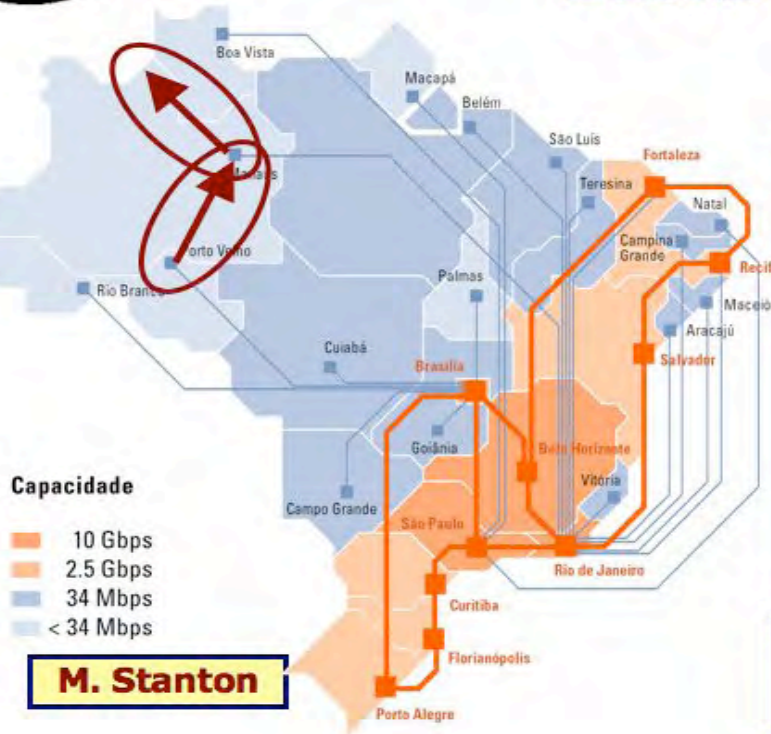
SPONSORS



Sessions & Tutorials Available (w/Video) on the Web



Brazil: RNP2 Next-Generation Backbone



New vs. Old
A factor of 70 to 300 in Bandwidth

2006:

- ➔ **Buildout of dark fiber nets in 27 cities with RNP PoPs underway**
- ➔ **200 Institutions Connected at 1 GbE by End-2008 (Well-advanced)**
- ➔ **WHREN (NSF/RNP) Link to US Now to 10G; 622M Link to GEANT**

Now extending to the Northwest; Dark fiber across the Amazon to Manaus

The loops to the north and south of the core triangle run at 2.5 Gb per second. That represented a factor of 70 to 300 times more bandwidth than anything in their previous network. Also once they started asking around they found much dark fiber throughout the country, owned by the railway or the electric company or a gas consortium.

COOK Report: So they were doing all of this by private investment, isolated from each other and from the national government, and no one had much of an idea where anything was?

Newman: A reasonable conclusion. It was the national

research and education network organization RNP that investigated, pieced it all together, and then acquired access to much of the fiber. In addition, as mentioned above, they proposed a project, approved by the government, to install a dark fiber network in 27 of Brazil's state capitals. **RNP found that they could set up a metropolitan network connecting schools and government institutions and research institutions for about 30% less than what they had been paying the incumbent, and get 1000 times the capacity.** The project, which will connect 200 institutions at 1 Gbps, is now well advanced. [Some of the specific data on

this is not among the slides used here, but is in the 2009 SCIC slide set.]

Do you see the dark red arrows at the upper left? There is Manaus, an amazing city of 2 million people in the Amazon jungle with nearly no connectivity. They made a plan to build a 622 Mbps link from Puerto Velho to Manaus across the Amazon jungle. It took them a bit more than a year to complete it. A future plan is to extend this to the sea and link up internationally.

The question became one of who would use this new network? One of our key roles was to assure the research and education network de-

velopment people, Mike Stanton in particular, that if he built this, it would not go unused. We promised that we would show him how to use it effectively.

COOK Report: When you say to a man like Mike Stanton -you'll show him how to use it - I gather that with all these other reports of high energy physics optical network's in many other country's and the ICFA a annual documents you've been telling me about that you've accumulated a large body of knowledge. You then can show Stanton what has been happening in other countries, and extrapolating from their experience that he can be certain that it can then be transferred to his country?

Newman: Yes. He'd been reading our reports and appreciated them. Along with this there were two groups of physicists, led by Professor Alberto Santoro in Rio de Janeiro and Professor Sergio Novaes in Sao Paulo, who were building two Tier2 centers. This was the focus of scientific activity that would use this network. In addition they also participated in several of our demonstrations at the annual Supercomputing Conferences.

Building on this experience, about two years ago, Novae wrote a proposal together with Luis Lopez the head of

the São Paulo regional network (ANSP) to establish a grid computing network throughout the state of São Paulo called GridUNESP, and the regional government approved it. The focus of this grid is a new Tier1 facility, the first Tier1 in Latin America, which is now being built at Novaes' university. 16 Tier2 centers connecting to the Tier1 are also being built, at each of the 16 university campuses in the state. In addition to high energy physics, this new grid facility will serve many other fields of scientific research. This infrastructure work opened up a lot of new possibilities. The richest area in the country became the most enthusiastic which is São Paulo in the lower middle of the Brazil RNP2 slide above.

The next step related to this happened at the ICFA Digital Divide Workshop that we held in Mexico City in October

2007. We sat down with Mike Stanton (representing RNP), Luis Lopez (on behalf of ANSP), and my colleagues who run the AMPATH network in Miami, and we worked out a scheme by which they could work together and upgrade the link that they have now that goes from São Paulo to Miami which is funded by the state of São Paulo mostly and partly by the NSF. The first 10 Gbps link paid by São Paulo is expected to be in service in February 2009, and the second one paid by the national network will be in service a couple of months after that. As a result the connectivity of the grid computing facilities in Brazil to the US, and onward to Fermilab, many US universities or CERN will be vastly improved. The new links will also be used in part to support better connections to the internet for the whole country.

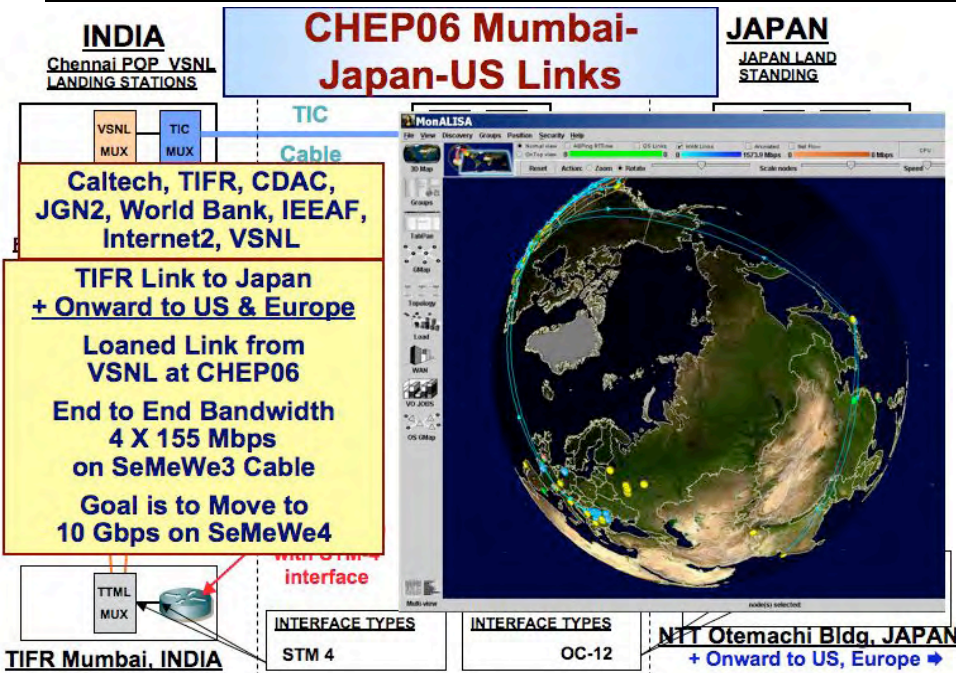


CHEP06: Pres. of India Collaborating with US, CERN, Slovakia via EVO



Data Transfer of 15 TBytes in 2 Days

Helped Spark "Knowledge Network" 500-1000 Institutions at 100 - 1000 Mbps



could actually get high throughput across the continent and across the world.

The CHEP-Mumbai-Japan-US slide to the left is a bigger picture with a few details about how we got 155 Mb per second on each of the four links. The goal is a 10 Gbps international link on Se-We-Me4 (or another cable). This hasn't happened yet, partly because of a lot of counter pressure from VSNL which wants to keep the kind of charging that they have had in place.

The next two slides show the ICFA Digital Divide Workshop

Other Workshops

The slide on page 23 above is taken from the CHEP06 high-energy physics computing conference in Mumbai. You see there on the left the President of India and me looking at a display where we running our collaborative system, and he is communicating with our colleagues at Caltech, CERN, Florida and Kosice (Slovakia). We also had a live stream from the Winter Olympics to illustrate some of the system's additional capabilities. On the right you see the MonALISA global display and network monitoring results, showing the traffic over the links from Mumbai to Chennai and from there to Singapore and onto Japan and across the Pacific to the West Coast of the US and on to Europe. The achieved data transfer, of 15 TB in two days, has helped

encourage the creation of the Knowledge Network I mentioned earlier in our conversation. It showed that you



International ICFA Workshop on HEP Networking, Grid and Digital Divide Issues for Global E-Science



National Academy of Arts and Sciences Cracow, October 9-11, 2006 + Sinaia Romania
<http://icfaddw06.ifj.edu.pl/index.html>



Sinaia, Romania October 13-18, 2006 Mihai Petrovici
<http://niham.nipne.ro/events2006/>



**SLOVAK Academic Network January 2008:
1 GbE to 10 GbE Switched Ethernet**

SANET - Slovenská akademická a dátová sieť
(február 2008)



**2500x Increase
Since 2002**

**Weis
Horvath**

- **1660 km of Dark Fiber CWDM Links**
- **2002 - 2004: Dark Fiber Links to Austria, Czech Republic, Poland**
- **2005-6: Complete 1 GbE links to all main sites**
- **2006: 10 GbE Cross-Border Dark Fiber to Austria & Czech Republic (11/2006); 8 X 10G over 224 km with Nothing In-Line**
- **2007-8: CBDF links to 10 Gbps; Transition backbone to 10G has begun**

that we held later that year, dealing with digital divide and grid computing issues with a focus on central Europe. The week-long workshop took place in Kraków Poland and also in Sinaia, Romania, in two parts during one week.

the universities and research centers from Bratislava in the West to Kosice in the East were connected at 10 Gbps.]

Notice that the international links with their neighbors are also shown. They were

among the first to support cross-border dark fiber to the neighboring countries. By setting up a dark fiber ring among Austria, Czech Republic, Slovakia and Poland, all four countries are able to share a single link to the GEANT2 pan-European network, which has a high subscription cost.

Another thing that they did was show the ability to send data a long way with nothing in between. Namely eight 10 Gbps links over 224 km without any intervening amplification.

The slide below contains a little bit about the Mexico City workshop in October 2007. We had tutorials on advanced technologies for the students before the work-

And on the Slovak Academic Network slide you see a successful example. The Slovak Academic Network in 2002 had as its fastest link 4 Mb per second. Since then they have connected up all the universities with a network which is represented by a Gigabit backbone in white on the map and then starting a couple of years ago migrated to 10 Gb per second which is shown in gold. [By January 2009, all

International ICFA Workshop on Digital Divide Issues

<http://fismat.uia.mx/HEP/ICFADDW2007>

ICFA
Mexico City
October 24-27, 2007

Local Organizers
S. Carrillo, UIA (Chair)
A. Zepeda (Cinvestav)
C. Casasus (CUDI I2)
L. Nellen, UNAM
G. Contreras, Cinvestav
J. Martinez (Cinvestav)

> 100 Participants
70% from Latin America
Presentations: Latin Am. NRENs; Univ., Lab Directors, LHC Experiments
Tutorials on Advanced Tech.: EVO, Networks, Dynamic Circuits

Latin Am. Country	NREN Organization	Nat'l Connections [Access BW often Less than this]	External Capacity	A.Santoro M. Stanton
Argentina	INNOVA-RED	256 Kbps – 34 Mbps	90 Mbps RedCLARA (temp. disconnected 12/2006-1/2007)	
Bolivia	ADSIB	64 – 128 Kbps	1.5 Mbps (commodity)	
Brazil	RNP	up to 10 Gbps	2.5 Gbps (incl. RedCLARA+ANSP)	
Chile	REUNA	up to 310 Mbps [Access 10 Mbps]	90 Mbps RedCLARA	
Colombia	RENATA	10 Mbps	10 Mbps RedCLARA	
Costa Rica	CR2Net	32 – 512 Kbps	10 Mbps RedCLARA (disconnected)	
Cuba	RedUniv	19.2 Kbps– 2Mbps	Not known (commodity)	
Ecuador	CEDIA		10 Mbps RedCLARA	
El Salvador	RAICES		10 Mbps RedCLARA	
Guatemala	RAGIE		10 Mbps RedCLARA	
Honduras	HONDUnet		Not known (commodity)	
Mexico	CUDI	155 Mbps [Access 2-34 Mbps]	1 Gbps-RedCLARA, 1 Gbps CENIC (PacWave)	
Nicaragua	RENIA	100 Mbps	10 Mbps RedCLARA (disconnected)	
Panama	RedCYT		45 Mbps RedCLARA	
Peru	RAAP	10 Mbps	45 Mbps RedCLARA	
Uruguay	RAU	64 Kbps to 1 Mbps	34 Mbps RedCLARA	
Venezuela	REACCIUN	155 Mbps & 34 Mbps	90 Mbps RedCLARA and Ampath	

shop, and talks via EVO from the Fermilab and CERN lab directors as well as my colleague Barry Barish who heads the International Linear Collider project.

A year after this workshop there is now movement toward a more modern network in Mexico. As I also noted the key people from Brazil who attended and the conversations with them that continued to push forward progress in Brazil.

Part of the motivation of the Mexico conference was the poor collectivities of the various countries in Latin America. Pink are the connections that I consider really bad. Yellow is on a scale that is so-so, and Brazil and Mexico and white are doing relatively well compared to the others.

What is highlighted in the slide however, is that the bandwidth on access links from the universities in Mexico to the backbone is still

very low. In fact we were only able to have remote participation at the conference through a short term donation from TelMex of a 34 Mbps link. Use of the link was limited to the conference. The rest of the university had to share their usual 2 Mbps connection.

Enabling Virtual Organizations

A Low Cost but Highly Efficient Global Telepresence System

The first problem I had tackled, as we discussed, was networking. Then **starting in 1994 when we knew we were going to begin working on the LHC** – and at that time I'd only been working on networking for 12 years -- **we realized that the next problem to be tackled was one of remote collaboration.** We started in 1994 with funding from the Department of Energy of \$45,000, which was just enough to purchase three UNIX-based workstations each of which had a \$ 6,000 graphics card. We found that Silicon Graphics had a kind of workstation you could carry on your back using an over the shoulder sack. In addition to a slot for the video card it had a port for an overhead projector that would enable the workstation to display via the projector the video that the computer was processing.

With this equipment in mind I made some configurations aimed at doing videoconferencing over IP. This was in parallel to the usual videoconferencing systems that ran over an ISDN line that we already had and that cost a lot of money. In 1994-95 we began to think a lot about

how to establish videoconferencing using a web browser. We did some tests and then in 1996-1997 our team wrote the first such system called VRVS – the Virtual Room Videoconference System.

We had a plan to put it into production by March of 1997, but when we exposed it to our colleagues in the fall of 1996 they said "No we cannot wait until March. We need to use it starting in January." So we pushed the release date up to January 97.

COOK Report: How would you characterize the group of people who wanted it?

Newman: They were researchers mostly in Europe and the US. We started with about 100 sites and since that time it is grown by a factor of about two every year. VRVS reached about 30,000 users and has been rewritten about four times. It was originally written in C and ended up written in C++ with some Java, with some legacy pieces left written in C. In 2006 we just froze the code because we couldn't see developing it anymore, because of its ten year old roots.

Philippe Galvez (VRVS and EVO's chief architect) and I then decided that it would be better to simply replace the whole thing by a new system written entirely in Java. Consequently, in 2006 we started developing the new system.

COOK Report: C and C++ are not used that much anymore correct?

Newman: C++ is still used widely in scientific computing and is the main language used in many fields of scientific research including high energy physics. But for remote applications, C++ compared to Java is hard to work with. Java also has the advantage of being largely platform independent, and it has better support for multi-threaded programs. You can write C++ threads but doing so is painful.

We called our new system EVO (Enabling Virtual Organizations). It went into production in June of 2007 and our user base started over with that event. We shut off VRVS a few months after that. Since that time, our EVO user base has continued to grow exponentially.



EVO Objectives

- ◆ Provide a global-scale robust real-time collaboration service to the LHC experiments and other major research and education programs
- ◆ Support and foster much more effective collaboration, across the full range of working environments from laptop to conference room to control room and auditorium
- ◆ Build and operate a system that can support the entire research and education community, on a global scale
- ◆ Design, build and evolve a system that can do all this, Operated by a Small team, with the required high level of functionality, and Low Cost

EVO.CALTECH.EDU

EVO is a real-time collaboration service built upon our global system MonALISA.

COOK Report: What do you mean by built upon?

Newman: It uses MonALISA as an underlying monitoring and communications infrastructure to help in setting up the video streams.

COOK Report: So in a sense EVO is a component of MonALISA?

Newman: No I would describe it as an application that uses MonALISA. EVO is written as a separate application and uses MonALISA to help with the interconnections among the servers that manage and transmit the video streams. It automatically configures and optimizes the interconnections among the EVO servers as well as between each EVO

end-client and one of the nearby servers, for example.

The Economics of EVO

COOK Report: And if we ever somehow got a dark fiber-based educational network in United States could EVO and MonALISA run on that network?

Newman: Yes they are both written in Java and completely transportable. The point is they have been designed to be highly scalable and pervasive, and to be used by a very large community.

EVO supports collaborative meetings that span the full range of working environments, from a desktop/laptop to a small workgroup to an auditorium to a control room. It also interfaces to all the usual H.323 MCU-based


videoconferencing systems and also to Access Grid. Using EVO, people in all these different working environments are able to successfully collaborate together.

Having a system like MonALISA that monitors the system from end to end in real time, and handles some of the monitoring and communications operations as well as problem mitigation and recovery automatically, enables our small team to keep the EVO operation running smoothly.

I would say that, including handholding of end-users which is not a normal commercial function, we only dedicate two to three full-time engineers to operations and support. The rest goes to ongoing system integration and development, testing new video and audio devices and meeting user's requests for additions or changes to the system.

COOK Report: Two to three people globally?

Newman: It does have some local support but, yes, at a global level we try to provide only what the industry would call third level support. For CERN we provide also first level support but for other organizations, in general, we are supposed to provide only third level support. For example if a guy cannot find



What is EVO?

- ◆ **Unique leading edge Videoconference Technology**
- ◆ **A World Wide Videoconference and Collaboration service for the High Energy Physics community.**
- ◆ **A scalable and self-healing collaboration infrastructure used with powerful agnostics videoconference client.**
- ◆ **A unique way to bridge together all the videoconference standards (H.323, SIP, etc...).**

1891

EVO.CALTECH.EDU

CALIFORNIA INSTITUTE OF TECHNOLOGY

out where to turn on his microphone or doesn't know which audio device to choose, apart from the CERN community and some people closely connected with the LHC, we don't support that. But if someone says oh all my meetings crashed and it sounds like there is a network problem, we would investigate.

EVO is designed to be self-healing. If one of the links or services goes down, it automatically reconfigures itself to work around the problem. It bridges across multiple videoconference standards from H. 323 to Session Initiation Protocol, as well as its own set of tools that are based on an academic set of protocols developed in the VRVS days.

COOK Report: Is all of this essentially open source?

Newman: No. What happened was that in 2001 we discussed this with the chair of our department when we were contacted by angels who provided us with some funding and we started a company. Caltech owns the intellectual property and we license it from them following the standard Institute Policy. This is the usual arrangement where the university owns the intellectual property and we have a right to use the technology that we invented. For some four years our company went through a venture-capital exercise. Now we are starting over again with a new company, and without any external funding.

COOK Report: My only reason for asking this is to ascertain some of the basic economics. If it is open source obviously you don't have to pay the fee for using it. But as it apparently is not open source and you are spreading its use now, what is the economic model? Are there usage fees?

Newman: In a business model we are looking at, considering the number of people using the EVO system and the funding required to operate it, the conclusion is that between \$100 and 200 a year per person would be enough to provide the necessary support to our community.

COOK Report: Speaking hypothetically, supposing someone in the new administration wanted to invest in this system for wide scale use in science education at the high school or community college level or in some other ways, what would be the issues that they would be facing?

Newman: Probably the fact that it is not open source.

COOK Report: There would have to be a license executed? With Caltech for example?

Newman: Well it could be with our company EVOGH,

since we've already licensed the technology from Caltech.

We have now been talking with the NSF, and they use EVO internally at the Office of Cyberinfrastructure. Some of the folks we work with see the advantage of having as many people use EVO as possible. But if we propose a nationwide deployment for research and education, then we might once again have to deal with the issue of it not being open source.

COOK Report: If someone looked at this and understood the potential and wanted to effectively reverse engineer it in open source, what would happen?

Newman: I don't know what you mean, really. Wanting to and actually doing it are two different things. When in an earlier context I talked to KDDI about VRVS, and we presented a very high valuation of our company, they were initially shocked. But after asking them to seriously consider how much it would cost them to re-create something like it, and discussing what it does for an hour, they became very friendly.

COOK Report: But if enough people decided that it has great potential. At least there should

be no reason that they couldn't sit down with each other and work out some kind of a solution that could open some interesting avenues?

Newman: That's true. But we've seen other companies attempt to do similar things, and what they turn out has been very primitive by comparison.

COOK Report: Are there any publicly available figures on the cost of licensing this? Did you just say in fact that it's about \$100 per person per year?

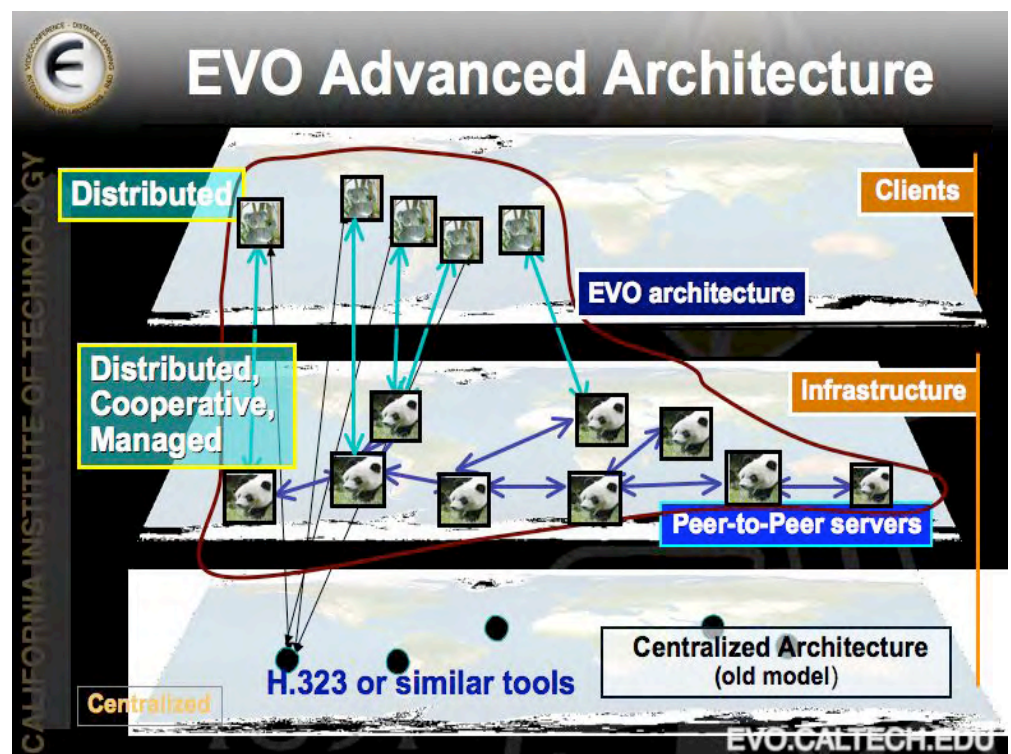
Newman: Well it depends on the scale of use. We have a business model and we do have a price list. We also have a standard agreement

for resellers who can resell it for a certain percent of the total costs and give us the rest. There is a price list that they get if they wish to become resellers.

More about the Architecture and Use

The Advanced Architecture slide below contrasts EVO with traditional systems that use H.323 and so-called multipoint control units (MCUs).

H.323 systems are centralized around the MCUs, so that end users on different MCUs are unable to communicate with each other unless they set-up manually a complicated and static environment called cascading. That architecture is contrasted with our own, where we have



a whole layer of servers called Pandas where all the inter-connections are done automatically in function of the best quality of the possible network paths. (Our team members who wrote the software like to use the name of animals for the different components. The EVO client is called a Koala, for example).

The Pandas make up a fabric of servers, and MonALISA services help with their inter-connection. When in use the system tests the connectivity between each Koala client and a short list of possible Panda servers, to figure out the best one to connect to. Building a fabric in this way, with optimizations at the edges and among the servers, makes the system highly expandable.

COOK Report: When on my own I asked my Economics of IP Networks mail list about their knowledge of or experience with EVO, Don Clark who runs the New Zealand Research and Education network said in that the H. 323 support had not worked well for them and added that their members have not gotten a great deal of use from experimenting with EVO. Did you see that comment?

Newman: No. All I can say at the moment is that Australia is right next to

New Zealand and they are very happy with it and are starting some joint development with us of some of the interfaces.

COOK Report: Ok fair enough, please continue.

Newman: Slide Key Architectural Functions shows some of the features. As I have pointed out, the connectivity and the routing of video and audio streams are automatically supported by the MonALISA-based infrastructure. There are alarm notifications. The end clients are intelligent in that they look for limitations of the end systems. So if you are going to receive ten high resolution videos from other sites and you are on wireless, for example, then the system will

automatically reduce the number of videos, or just send audio to you. There are automated adjustments going on all the time for all the participants. There are also important features such as being able to get through multi-layer Network Address Translation barriers, and adapting to and keeping up with a lot of different operating systems. EVO also includes a set of phone gateways which means you have a distributed phone system attached to this. If you are on the road and don't want to have to deal with video you can simply phone in to a nearby bridge.

COOK Report: Are there some similarities to Skype?

Key Architectural Functions

- ◆ **Dynamic auto-registration to high level directory services to provide a global infrastructure view.**
- ◆ **Automatic Alarm notifications when monitored parameters (system or network) go outside preset bounds.**
- ◆ **Automatic Re-activation of components and services.**
- ◆ **Continuous Monitoring of network quality (packet loss, jitter, latency) between its peers and its possible peers (short list).**
- ◆ **Automatic Rerouting to obtain the best performance/quality.**
- ◆ **Dynamics Services (video, audio, screen content, data, ...) that matches current resources/capabilities of the end-client machines and applications.**
- ◆ **Encryption among the core nodes (Pandas) and between each node and its clients (Panda ↔ Koalas)**
- ◆ **Access to real-time and historical monitoring data; extensive drill-down capability**

CALIFORNIA INSTITUTE OF TECHNOLOGY

EVO.CALTECH.EDU

Overall System Capabilities

- ◆ Works on MS Windows, Mac OS X, Linux !
- ◆ Connects to H.323 clients, H.323 MCUs, SIP clients, SIP MCUs, IP telephones, standard telephones !!
- ◆ Accepts qCIF, CIF, VGA, 4CIF, 720p, 1080i, 1080p resolutions!!
- ◆ UDP and TCP tunneling through one Port. Network Address Translation and Firewall transversal.
- ◆ Encrypt Video and/or Audio and/or Whiteboard and/or IM/Chat and/or data.
- ◆ Multilanguage support: 11 languages supported so far; template makes add'l languages easier.
- ◆ Administration of meetings: Become a meeting moderator; Mute/Unmute video/audio of a participant; kick out a participant; add a new participant as a moderator.
- ◆ Sharing of computer screens (any part) among participants of the same meeting.
- ◆ IPv6 is fully supported.

EVO.CALTECH.EDU

encrypt any of the streams. The interface essentially is a template and you can change some of the items in the template. This is used to provide versions in multiple languages, for example.

The slide "EVO: World Wide Topology" below shows a picture of typical day with the different servers and groups of people clustered around the servers with the connections being maintained by MonALISA. The grey tags represent users with a IM presence status set to not "active" (busy, away,...) while the yellow ones shows the "active" users (i.e. interacting with their computer).

Newman: Well with Skype, the number of people on a single conference is limited and the entire architecture closed. Skype is based on a pure peer-to-peer architecture at the end-client level while EVO is based on a pure peer-to-peer architecture at the server level, close to the core network backbone which give more control from a networking perspective. There are all kinds of things in other systems that are not extensible or scalable as they are in EVO.

all system Capabilities" above points out more of these capabilities. EVO can get through robust firewalls by opening a single port. We can

The slide "Over-

EVO : World Wide Topology

50 servers spread across 25 countries

MonALISA

EVO.CALTECH.EDU

The numbers

- ✓ **In production since June 2007**
- ✓ **18,900 registered users**
- ✓ **600 users (unique) per day**
- ✓ **Around 160 meetings per day**
- ✓ **Around 3,000 meetings per month**
- ✓ **Up to 130 users in the same meeting**

Building on the experience with the VRVS system, since 1996

EVO.CALTECH.EDU

25,000. Meetings can range from small conferences of the very few people up to +200 people in a single meeting that functions as a plenary session for one of these big experiments.

Because I generated these slides for a talk in Slovakia we have some information here on the use of the EVO in Slovakia. Slovakia has its own telephone gateway into the system and some of the people involved on our sister team in Slovakia are shown on the next slides

Our Slovakia team includes experts on Codecs. They have developed an H.263 codec and now also our first H.264 codec, which works very well. Both the H.263 and H.264 codecs support 1080i high definition video, which is the most that a mass market

EVO: Slovakia usage

- 662 Registered Users**
- 1300 Meetings with Slovakian Users**
- SANET Community: 490 Meetings with 3500 Connections**
- 1 Telephone Gateway +421 55 234 2420**

EVO.CALTECH.EDU

COOK Report: And those 50 servers support about how many clients?

Newman: In a typical day we get about 400 to 500 concurrent connections while on the LHC start up day we reached 1680 connections -- traffic that the system handled fairly easily.

The first slide above shows some high-level numbers on the usage. The number of registered users is just under

Slovakia – Caltech Collaboration Development of EVO system

- Dr. Pavel Murin**
High Energy Physicist; Leader of EVO project in Slovakia; Distributed archive system
- Dr. Viktor Michalcin**
EVO developer Video codecs; Video transcoding; New OpenGL ViEVO
- Dr. Pavel Farkas**
EVO developer Video + audio codecs; New OpenGL ViEVO; Pocket PC

EVO.CALTECH.EDU

Slovakia – Caltech Collaboration
Development of EVO system

Michal Pauliny
Video codecs
Whiteboard
EVO Booking
Video camera detection in ViEVO

Stefan Zavoda
Distributed archive repository system
Player for EVO recorded streams

Rasto Adamek
Video codecs;
Desktop sharing;
New OpenGL ViEVO for MAC OS X

Martin Harcar
Video codecs
Desktop sharing
New OpenGL ViEVO for MAC OS X

EVO.CALTECH.EDU

HD camcorder can do. 1080i is what you see whenever you watch high-definition broadcast TV. In principle we can also support 1080p but to the capture 1080p takes a

several thousand dollar camera era which we have not yet purchased. The next slide below shows a prototypical high definition setup for EVO using a PC, HD

camcorder, HD capture card and large plasma display for \$ 3000.

This setup is better in several ways than Cisco's Telepresence that costs a hundred times more per site, although Telepresence's audio is excellent.

COOK Report: The implications of being able to use a system like this for everything under the sun are rather mind-boggling. I can understand why the IP attorney said Caltech would be watching it pretty closely. Is this still under license from Caltech?

Newman: Yes. We have a worldwide exclusive right to use and develop the technology, under our license agreement with Caltech.

EVO Full HD Prototype

Example setup:

- ◆ Quad Core Dell PC
- ◆ Blackmagic Intensity HD video grabber card
- ◆ Sony HD camcorder
- ◆ Panasonic HD Plasma Monitor
- ◆ Phoenix Duet echo canceling mic
- ◆ Total: ~\$3000

EVO.CALTECH.EDU

OpenGL's Use in EVO

The EVO user interface uses OpenGL. You may ask what OpenGL does for you? Well it does a lot of things. It is supported on all the leading graphics cards in PCs and it has a lot of built in acceleration functions. It also supports the use of different libraries including all the re-

cent ones from Intel as well as manufacturers like NVIDIA. As a result of this a lot of the operations involving the handling of the video stream can be done with hardware support in a mass-market PC. In fact I would say many \$ 100 or even \$50 graphics cards will do this now. For \$ 100 - 150 we now have cards that support two full HD displays (which can also be large LCD or plasma TVs) providing one 4 megapixel screen space. We've come a very long way from that special board that was able to do video (at a much lower resolution and frame rate) for \$6000 each in those 1994 UNIX workstations.

U ViEVO features

- **All videos are shown in a single window**
- **Use GPU (Graphical Processor Unit) instead of CPU**
- **Real time:**
 - 2D and 3D special effects: Reflections, shadows, surface textures, etc.
 - Multilayer: transparency, text over video ...
- **All resolutions from QCIF to HD 1080p**

EVO.CALTECH.EDU

Using OpenGL also gives you a lot of very nice three-dimensional functions. The live video streams are painted on surfaces in a three

dimensional space. These are usually planes, but in fact we can paint the streams on any surface.

You can see Philippe and some other members of the team in the next two slides which illustrate some of the features and possibilities of the OpenGL-based EVO interface.

You can make the speaker big and arrange the others at the bottom. Or stack them up as they are shown above. We can easily arrange other things like walls of speakers and walls of listeners, and put other objects in the

ViEVO Application

Count_Dracula

Team Meeting

Selected Speaker

Physics plots

Mixing Content and Videos as Needed

New features coming soon. (Already working in our test lab)

- Possibility to connect a VR headset to an H.263 meeting coming on conventional H.263 hardware.
- An highly enhanced mixing functionality will allow a provide better audio reception for and deliver that can records only one audio stream in a H.263 stream.
- A new transcoding functionality will provide some connectivity to Accented meeting for H.263 and so.
- The transcoding function will be based on a new video audio codec supported by H.263. A key point will be a video codec supported by the H.263 codec.
- A new version of Proton (H.263) for Proton's video network Address Translation (NAT) environment.

7 participants

EVO.CALTECH.EDU

foreground.

You can put a lot of information in a three-dimensional space. And video is just one source of information. You can put things like some of our data distributions so you would have a combination of content and video for collaborative work. And you can also see that there are some things that while not quite so important, you do get for free - like ray tracing, reflections and transparency so that you can have surfaces with different qualities. In the video features slide above with Philippe in the middle, for example, you can see how his shirt is being reflected off the shiny surface on which the window stands. It's also possible to animate the objects in the 3-D space, although that is not a function users have asked for (at least not yet.)

You can fold up a lot of information into a small screen space. As a result we are thinking of collaborating with people like the University of Illinois at Chicago who like to have video walls. We will extend that sort of environment to other working situations in lower cost setups using some of the 3-D capabilities shown above. We will fold up some of the video-wall material into a space compatible with a given screen size, and do this in different combinations to accommodate anything

from a small screen to a few medium-size or larger screens in your office. In this way we can foresee multi-point sessions including sites with large video walls, collaborating with many other sites with screens of varying sizes in offices and/or small conference rooms.

COOK Report: Does it make sense to ask you at this point how all of EVO fits into the context of the Optiputer and Opti-portal? All that stuff is open source. It is intended to enhance collaborations but it doesn't look nearly as powerful.

Newman: Well relative to EVO those are meant for a very few sites and the cost per site is relatively high. Another difference is that the Optiputer sites are generally not bandwidth constrained, and they normally do not bother with the issue of com-

pression to conserve bandwidth. For them high quality means raw video, which has the advantage of low latency. But raw full HD video, using of order of one Gbps per site, is not something you can deploy widely at many sites to serve a large community.

Apart from the latency, EVO can get similar quality with a couple of Mbps, especially with the H.264 codec, and this makes it much more scalable. And to help improve on the latency, we are going to work on multi-threaded codecs using multiple cores in the new processors, and varying degrees of compression adapted to different bandwidth levels.

A Multiplicity of Capabilities

COOK Report: So in this sense, the license fee aside, the technology that you are

Self Managed Infrastructure

EVO Uses MonALISA

- ◆ Continuous monitoring of network quality between each server and its possible peers.
- ◆ Detects any network cut or peer down : Automatically reconnects the topology of servers to fix the problem.

CALIFORNIA INSTITUTE OF TECHNOLOGY

<http://monalisa.caltech.edu> EVO.CALTECH.EDU

explaining is much more extensible and more broadly usable over existing infrastructure?

Newman: Yes. It has many built in aspects that make it adaptable to a wide range of network conditions. The software agents it uses, for example, measure the bandwidth available and adapt to it.

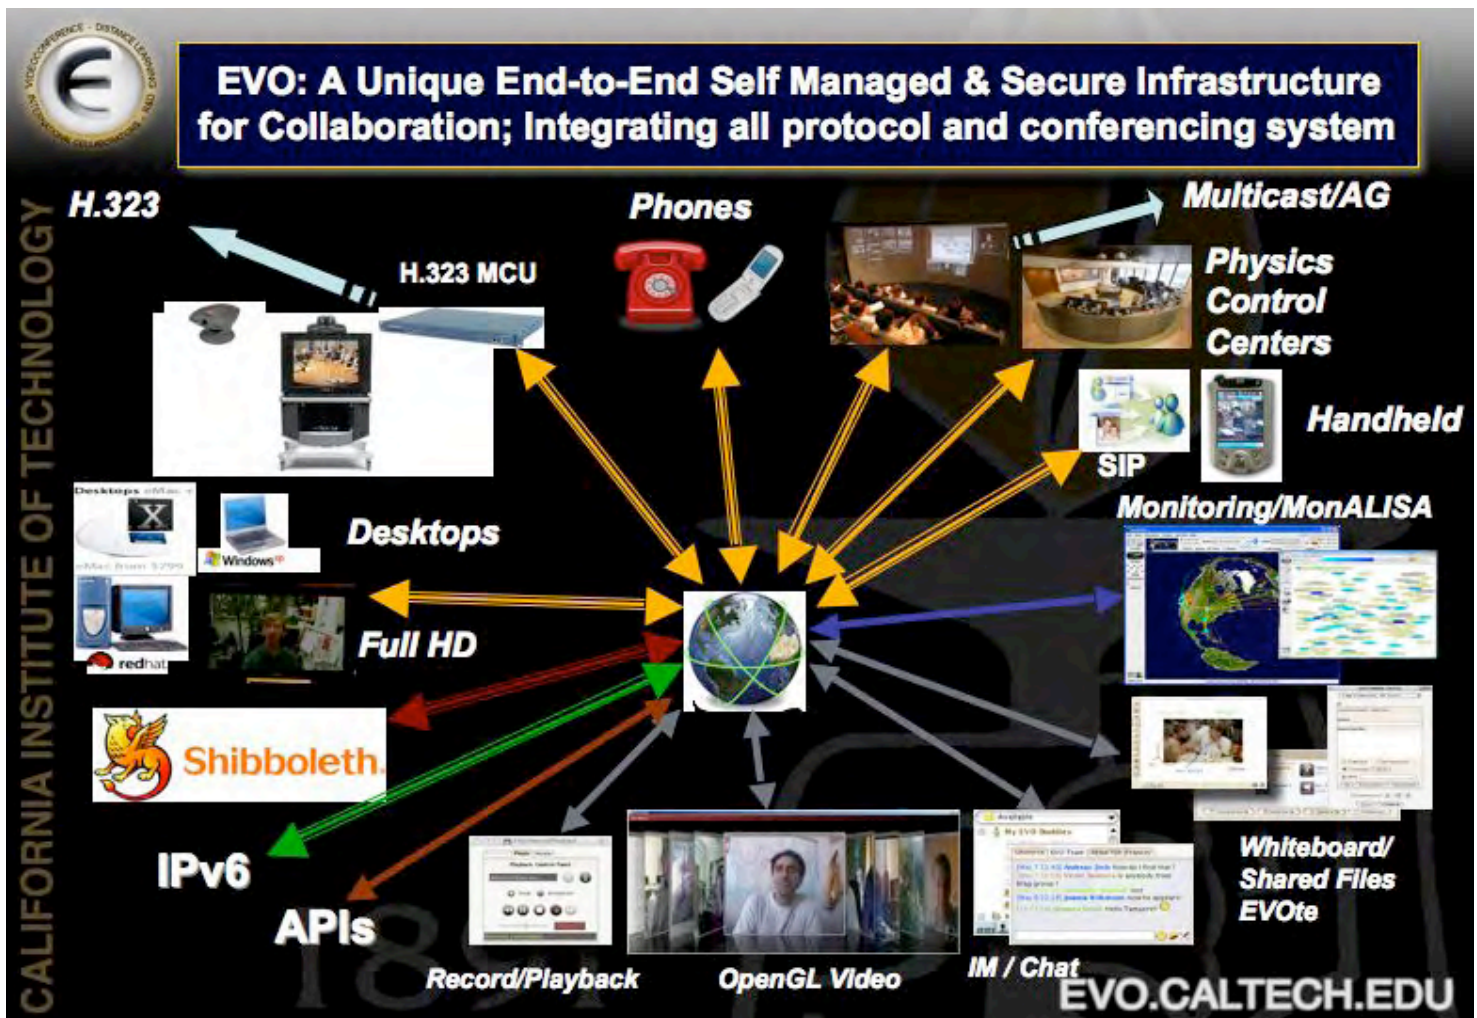
On the slide Mixing Vieve Applications page 35 above you see a mixture of the possible kinds of content including graphics displays of events, presentation material, live

videos, histograms, and plots. You can arrange this anyway you like. It is meant for collaborative work and for presenting lots of data in a limited space.

Slide 66 that we've already talked about above is our full HD prototype. There are now several of these using mainstream components. To give you a feeling the PC system I just got in my office has a low-end quadcore processor that costs approximately \$700. (A system with the latest Intel Core i7 920 processor is about \$1000). The capture card is \$ 240, the

graphics card \$ 100, camcorder \$ 700 dollars and the plasma monitor \$800 - 1400 depending on the size. The echo canceling mic is \$150. And that's it. You see behind the lettering in the slide Dave Adamczyk of our team, who puts these systems together in an exceptionally creative way.

Slide "Self Managed infrastructure" on the bottom of page 36 shows a few features of the MonALISA system that provides a self-monitoring, self-managing foundation for EVO.



EVO, which has operated around the clock for the last 18 months, effectively becomes one of the globally monitored parts of the MonALISA network.

The slide EVO: A Unique End-to-end Self managed infrastructure" on the bottom of the previous page summarizes many of the capabilities of EVO we been discussing. The slide highlights EVO's unique ability to provide a collaborative framework that integrates across the full range of working environments, using a diverse set of communication technologies.

To say it more simply - EVO is meant to connect anything to anything including different kinds of systems and different kinds of working environments.

COOK Report: Shibboleth is what?

Newman : Shibboleth is an identity management system. When you login you register yourself through Shibboleth. Folks in Aus-

tralia are going to help us to further our development with Shibboleth. It enables you to go from place to place having a single identity and enables you to be recognized anywhere. Shibboleth has been supported by Internet2. I understand it is some kind of de facto standard or close to it.

So EVO connects together desktops and laptops and small workgroup rooms, control centers and auditoria and we have had in the past handheld clients and will have them again in the future. One of my colleagues is looking into making a client

for a GPhone.

The slide below shows the 2008 Master Classes in physics. As we do every year now, EVO was used to bring together 4500 high school students in 22 countries from about 70 universities, who learn about particle physics both through lectures and hands-on exercises over a period of three weeks. Some of their exercises involve looking at event displays and analysis results and identifying different types of high energy physics events. And then they have a competition

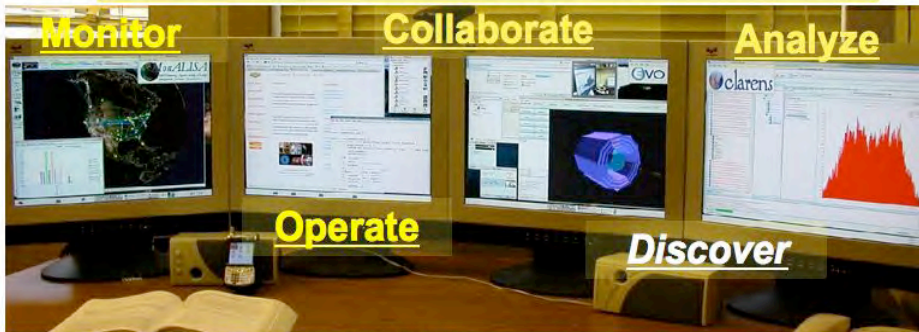
MasterClasses 08 Particle Physics

4500 high-school students from 70 institutes across Europe and the USA spent two weeks working at the frontier of physics in the third International MasterClasses for High-School Students.

EVO.CALTECH.EDU



Grid-Enabled Advanced Desktop: MonALISA, EVO, CMSSW, ROOT



- ★ **Automated Service Discovery**
- ★ **Remote Dataset Discovery, Movement, Processing**
- ★ **Monitoring and Control of Analysis, Data Motion**
- ★ **Smooth transition from small to large tasks**
- ★ **Integration of Collaborative Tools**

where teams collaborate to work out their answers and compare their results. We have had up to several dozen sites at a one time participating in this. This is just one of our many outreach programs. Another one is the US MasterClasses run by QuarkNet, which is aimed at high school teachers. And we already discussed our Digital Divide Workshops.

Now the grid enabled advanced desktop slide on the next page was put together a few years ago, when it was a bit harder to get many displays driven by a single PC. Michael Thomas, a software engineer who also manages our Tier2 center (and who was the 2008 U.S. Red Hat Certified Engineer of the Year One of our software developers has these four screens in his office. He tends to use his

laptop as a fifth screen.

On the screens in this somewhat old slide you can see event displays, analysis code and results in histograms, and MonALISA displays up and running, along with an early version of EVO.

COOK Report: So the 3-D technology shown a couple of slides earlier is to take these multiple screens and fold them down for use on a single screen?

Newman: You can just do more. Now there are also these little USB boxes by which you can extend one desktop across multiple screens and do so in a very convenient way. Each little box costs \$69. Or one can buy a four-fold full HD graphics card that costs \$ 400, which gives you better per-

formance. Screens are becoming tremendously cheaper. I think when we first bought them a few years ago these 19 inch screens cost more than \$1000 each, but now there are 22 inch screens at \$150, and a 28 inch screen is \$350.

So you can have in your office somewhere between one and four screens and have EVO in three dimensions to fold up some of the content You may also have a lot of things going on between the code and running applications and video images of collaborating people and also monitoring what's going on in the network if you are transferring a dataset. You can see how fast it's moving; and if it's not moving then you can look at the MonALISA displays to determine why not. This multiplicity of capabilities is what the environment is all about.

Some variation of this would be a good way to interact with other people. **Having enough computers or a wall full of streams is fine but people working in those environments should be able to interact with other people working all over the place - and in more modest environments as well. The point is that this is cheap and it is getting cheaper.** And it continues to get more capa-



View from Harvey Newman's command post toward the close of the 2008 Bandwidth Challenge on November 19, 2008. Various portions of five EVO session screens are visible in the large projected overlay of the MonALISA map. On Harvey's laptop is view from a slide set that he was explaining to me.

ble, through our ongoing developments.

Editor: here are some EVO related links:

- o From Australia:
<http://evo.arcs.org.au/>,
<http://www.arcs.org.au/>

- o From UK
<http://www.ja.net/services/video/agsc/AGSCHome/>,
<http://www.ja.net/services/video/agsc/services/services.html>

- o From Brazil,
<http://www.centraisite.com/>,

- o From France (the NREN),
<http://www.renater.fr/spip.php?article603>

- o From Slovakia,
<http://vk.upjs.sk/>

- o New Zealand,
http://www.bestgrid.org/index.php/EVO_in_NZ

- o Italy:
<http://server10.infn.it/video/index.php?page=vrvs>

- o South Korea,
http://www.ksc.re.kr/eng/project/project2_1.htm

Summation

We have covered a lot of things. Evolution of networks and how to use them well, as global dynamic systems. Both networks and all the applications that run on top of them. Both applications for managing the data at hand and for collaboration. As well as the digital divide issue and what we're doing about it. In general I think that our science collaborations are very nice examples for policy makers in the sense that we operate pretty much without borders,

and we try to bring along people in other countries to give them equal access and essentially equal rights, as partners in our worldwide scientific enterprise. We have a kind of world community paradigm which is very valuable, and which has also been applied in very tangible ways by our helping with the development of IT infrastructure that everyone can use.

I guess while I'm saying all this you are probably thinking that it's very ironic that I need to do some outreach within the United States.

COOK Report: Indeed the critical question would seem to be how to use what you've done in your experience to educate the new policy makers in the new administration?

Newman: I think that some of the high-level messages are that a lot of this conversation we've had has been to highlight some of the capabilities that are available. Our roadmaps are based upon our continued confidence that we will be able to use and further develop these technologies, using the networks effectively (up to their limits where needed) and bridging across networks of varying levels of capability, over long distances. The kind of Enablement we've shown is possible not only for scientific research, but also in

general, by looking at the learning process that we've undergone. As well as our demonstrated ability to change the way people interact across continental and transoceanic distances.

COOK Report: Well I am very interested in understanding how to make key administration people more aware of what you were doing and the trans-sectoral significance of the possibilities for enhanced education of investments in this area. How do you assess this situation?

Newman: I think the first item is to dispel the inaccurate perception that this is just one of several efforts of a similar character. People say well - there must be other things going on which are essentially the same? Well the answer is no there aren't.

The question to ask people who think they have seen something similar is whether they know how to build a fabric with this kind of messaging capability? Do you know how to build a large set of agents that are continuously aware of each other and really cooperate to carry out tasks over long distances, while running on small servers without a very heavy load? And just run all the time, around the clock and do this successfully for years,

without interruption ?

COOK Report: but you can do this in no small part because you have the huge network that runs MonALISA? Right?

Newman: Yes but the LHC network is not really a global private network. Our transatlantic links are indeed dedicated and mission-oriented but in the United States we use National Lambda Rail and Internet2, and in Europe we use the GEANT2 infrastructure. Elsewhere we use other national research and education networks, and sometimes we just use regular networks where people connect from home with a DSL or cable modem or perhaps with fiber to the home like FiOS.

COOK Report: The Panda servers all do have to reside on networks running MonALISA, correct?

Newman: Correct, but Koala clients can be activated by downloading the software and connecting to Panda servers by many different means.

COOK Report: Then MonALISA must be quite highly portable?

Newman: Oh yes there are many implementations of MonALISA that are run by

other organizations, without us. There are also some organizations that want to run it but don't want us to run it for them.

COOK Report: Is MonALISA open source?

Newman: No, but it's free. We don't want people to be able to tinker with a thing of this power. It is an intelligent global system that if corrupted by non-expert developers, could create havoc.

COOK Report: Somehow the huge overall unanswered question is how do you get to an open self-sustainable model based on the cost of the technology and cost of operation?

Newman: I think there is a huge potential for the research and education community but as people grow up within this, they can deal with information very differently. Children are getting into things that are immer-

sive. If not visually, then in ways in which they interact within their social networks; but they don't do the same for their learning experience and they don't do the same for interactive audiovisual things among many people.

Not yet.

To spread understanding further, I think intersections with policy makers are really critical.

Symposium Discussion February 18, 2009 - March 17 2009

Issues of Policy and Infrastructure

The Feb. 25 Supreme Court Decision on Anti-trust in Telco Pricing; Economic Thought Impacts Law Very Slowly; Why Fiber Always Superior to Copper

COOK Report: On February 25 the Supreme Court overturned the 9th Circuit and ruled out anti-trust as a ground for legal action in the claim of a price squeeze where the incumbent prices its wholesale service to a competitor high than it sells the same service to its own customers.

PACIFIC BELL TELEPHONE CO., dba AT&T CALIFORNIA, et al. v. LINKLINE COMMUNICATIONS, INC., et al. Certiorari to the United States Court of Appeals for the Ninth Circuit No. 07-512. Argued December 8, 2008--Decided February 25, 2009

Would someone translate this? It sounds ugly.

Brian Harris: Bottom line: Plaintiff gets another whack at recovering something or other from AT&T, if they can successfully amend their complaint. They get to amend (again) to try and show: "below-cost retail pricing and a " 'dangerous probability' " that the defendant will recoup any lost

profits." So they have to go back to the trial court level and try again.

More broadly, it appears to reflect an ongoing legal confusion wherein widgets (i.e., radios) are compared to services offered pursuant to a network. It is also a further narrowing of which antitrust laws may be applicable to the pricing behavior of telecommunications firms. The typical de facto price squeeze occurs when the rate for a retail service is lower than the rate for the same functionality at a wholesale level. Take the prosaic example of POTS. Say I can order POTS for \$13.50 + \$6.50 for the CALC. Yet the local CLEC must pay ~\$24 for the same functionality.

As long as both prices are arrived at in a lawful manner, no authority can find a de jure price squeeze. *Now trying to use American anti-trust law to prevent "upstream monopolists" from abusing "their power in the wholesale market to prevent rival firms from competing effectively in the retail market" is more of*

a sucker's game than ever. All nine justices were happy to see the case go back to trial court, but they used two different lines of reasoning.

This is a quick and dirty attempt to translate the legalese into more understandable language.

Cecil: Good summary Brian. Agreed. Trinko, just like the rest of the antitrust narrowing cases (Goldwasser etc.) should never have been brought. They were unfortunately fodder for a certain federal courts and in particular a Supreme Court (SOCUS) all too willing to mash it up and take a deep drag from their Chicago Economics Crack Pipe and blow the judicial smoke all over this. Unfortunately, telecom law, in particular is high quality regulatory economics crack. Taking a case like this up the appellate chain is about as risky as walking down addicts row with a big overflowing bag of the good stuff. You might make it to the other end, but your chances of being held up increase expo-

mentally with each step taken.

As to the immediate case, as Brian notes, SOCUS mostly punted, but not without expounding upon if not strengthening Trinko. At a high level, this points to a problem I've not only written about here before, but have been accused of dealing said mind altering substances and related paraphernalia. There is a natural tension in the law between a priori regulation (very broadly, the 1934 Act --> FCC ---> FCC Rules) and post hac (Sherman Act --> courts). Setting aside the SOCUS hooka, it is very difficult to maintain an antitrust case in an environment where the very entity you accuse of post hac antitrust violations are subject to a priori antitrust regulation.

If the rates are legal - e.g. approved or, in this day and age, even countenanced by a regulatory authority (typically state public utility commissions, noting however, there is no such thing as a "public" utility; these things are private subsidy harvesting machines but that's an entirely different adventure through the chemical jungle), you are going to have a damn hard time of convincing a judge that there has been an antitrust violation.

Now before anyone goes off half cocked and comes storming out of the houses with pitchforks, fire and brimstone, and holy hand grenades of righteous consumer indignation - just look at the logic of it - and stay, for a second at the conceptual level before we start carpet bombing each other with that case or this citation - on one hand they are regulated and on the other we claim a violation of the very antitrust principles the regulation is supposed to fix in the first place.

[Sure the 1996 Act has an antitrust savings clause but that's inapposite here for many reasons, particularly because I'm deliberately focusing on common sense before debating what the heck a savings clause enacted by a bunch of legislatures who'd just as soon cut off their PACs as see one scrap of law expire.]

My point is simple: the law is too damn complicated. As I said a long time ago on Cybertelecom and repeated here, the 1934 Act and all that goes with AND without it has become a massive Mobius strip. There is not a single position any reasonably competent telecom or regulatory attorney -- with regard to networks or the internet -- can take within this realm without the ability to build a case that

gets past motions. As a result, regulation and antitrust are basically deconstructing themselves. And the crack addicts are rolling on the floor laughing at all of us. Seriously, who here really expects this case to go anywhere? It's back to a federal district court for years and years of grinding, more appeals, and back up to SOCUS. Think AT&T is about to let up? And what do you think happens when plaintiffs get ahead? T offers a settlement. Plaintiff accepts. Who wouldn't? Litigation costs a fortune. There are 3 basic winners here: lawyers, AT&T, and judges/regulators.

Amazing thing is how much our species not only LOVES a good battle but consistently mistakes battles for progress.

Savage: UNE-based CLEC sued AT&T on a "price squeeze" theory: the retail price AT&T charged for DSL was so low, relative to what a CLEC has to pay for DSL piece-parts (loop, collocation space, etc.) that it was impossible for a CLEC to (a) cover the costs it would have to pay AT&T to offer DSL service while still (b) beating AT&T's retail DSL rates.

The Court said that this was not an antitrust problem. Since as an antitrust matter AT&T had no obligation to deal with the CLEC at all -- that is, no obligation to sell

the unbundled loop, collocation space, etc. -- that it can't be an antitrust violation for AT&T to sell those pieces-parts, but at rates that are too high to make a go of it.

This is totally distinct from whether or not AT&T has a regulatory obligation to sell loops, etc. at a sufficiently low price, or its retail DSL at a sufficiently high price, to allow a UNE-based competitor to function.

Under the Powell-Martin FCC the answer to the regulatory question was probably also "no," which means that it sucks to be a UNE-based CLEC. (Not a whole lot of news there.) Whether those regulatory positions will remain in a Jenachowski FCC remains to be seen.

Note: although I don't represent anyone in this case, it is probably worth noting that I have been publicly bear-ish on the UNE-CLEC business model starting in about 1996...

By the Time Economic Thought Impacts the Law it's Already Ten to 20 Years Behind Current Thinking

Savage in response to Cecil: A minor disagreement (Erik & I just talked about this...). The understanding of "economics" that becomes en-

shrined in Supreme Court antitrust cases, in the nature of things, lags literally 10-20 years behind what one might think of as "current" economic thinking. So if we use October 1, 2008 as a convenient marker for the Death of the Chicago School (financial meltdown edition), I would not expect Supreme Court antitrust cases to catch up until nearly 2030. Truly. On the other hand, what counts as a decent regulatory argument can change very quickly. In other words, *even if what AT&T was alleged to have done would not be an antitrust violation, that does not mean (or at least does not have to mean) that its conduct is not a form of "unjust" or "unreasonable" or "discriminatory" conduct, or some similar formulation based on regulatory, rather than antitrust, principles.*

So, while I am not thrilled with the Court's decision, (1) I'm not that surprised and (2) aggrieved folk have other avenues of relief.

Cole: Chris -- does this mean you are joining me as part of the [minority] OPTIMIST wing of this list?

Savage: I'm always an optimist. Plus, since I am in the business of making regulatory arguments, perhaps I see them as more viable in the long run.

Also, note that industry participants of a certain age grew up in a world in which the Terrible Sword of the Archangel Antitrust was wielded to smite down the Evil Monopolist AT&T. So deep in some folks' hearts, antitrust is sort of like the nuclear weapon in these deals -- maybe antitrust cases take a long time, but look what they can do! We can always be saved by antitrust! Goldwasser, Trinko, and now Linkline can be a bitter pill for someone who viewed antitrust as a soon-to-be-returning savior of pro-consumer competition.

In fact, with 30 years of hindsight, Judge Greene was slow on the uptake: the Chicago School was alive, well, and taking over the field even by 1984. The breakup of the Bell System was, arguably, the high water mark of the "old school" of antitrust thought. As noted in my earlier post, it will easily be another 20 years before antitrust catches up with current (non-Chicago) economic thinking.

McCullough: My problem is that **the courts are operating on the assumption that there is a "well greased" regulatory machinery that has been delegated the responsibility to take care of these things and is doing so consistent with the regu-**

latory statute; hence anti-trust need not step in. But there is not. We have complete regulatory failure. The regulators are captured by the regulated. And - as Erik says - the organic Act is a complete and utter mess which typically just gives rise to more Chevron deference and allowed regulator discretion because the words are so dissonant with the marketplace and the technology that nothing really makes sense. It is all ambiguous or full of interstices the regulator gets to play with.

So, antitrust does not work because there's regulation but there is not effective regulation. Sort of reminds me of what Judge Greene said in 1992 in his decision denying AT&T's motion for summary judgment in the case that led to divestiture. We need another Judge Greene.

I say that because I don't think regulation will ever work regardless of which regime is in power.

Harold Feld chimes in: So yesterday, the Court decided *Pac Bell v. Linkline*. <http://www.supremecourtus.gov/opinions/08pdf/07-512.pdf>

An interesting issue a reporter just asked me about is the Supreme Court's statement at Page 2 that the FCC deregulated in 2005 because

of "robust competition" between DSL and cable.

I don't think it has any impact on FCC regulation, although I expect the DC Circuit to make much of it. But I do think it pretty much nails any possibility of using anti-trust on any sort of network neutrality issue.

Paul Budde: I find it amazing that sophisticated people in a sophisticated country such as the US still hold on to such views, what on earth do we need to do to get the message through? The good thing with you guys sending us Sol Trujillo was that when he started to promote these 'American philosophies' in Australia, he got the whole country up in arms. Perhaps that was the best thing he did for the country. The good thing was that two successive governments didn't give in to his 'American way of thinking' and rejected the totally outdated 'laissez faire' policies that still seem to be accepted by courts and politicians around the USA.

Cooper: The courts will be the last to change. *The judiciary is filled with lawyers who grew up and entered legal practice during the ascendancy of Chicago School market fundamentalism.* This was a conscious effort to infect the judiciary. After they left law school and went into practice, the theory was

shredded in the academic literature, but there was not similar effort to re-educate the judiciary. Bork's bible on antitrust that all the judges memorized has been sliced, diced and flushed down the toilet, but the mind set of the judges remains the same.

I have pushed hard for both the DOJ and FTC under Obama to start a claw back campaign, reviewing all the bad decisions that have been so harmful, so that when the judge cites some lousy precedent, they can pull out the study and show why that precedent was wrong. Only the federal agencies can overcome stare decisis (the legal principle that once a court makes a bad decision it has to keep making it because precedent should not be overturned). If the Obama administration starts on day one and conducts a vigorous eight-year campaign (fingers crossed) it could have a big impact.

Of course, legislation that reverses some of the worst decisions to raise the stature of retail price maintenance) would help. The Trinko decision (that exempted local telephone companies from close antitrust scrutiny on the grounds that they are regulated, which they are not) and the Brand X decision, which abandoned Title II common carriage for broad-

band telecommunications, would also be on my short list.

Cecil: [to Paul Budde] I agree with you that the 'laissez faire' approach is a complete train wreck. I think that as applied to telecom it has been used even more perniciously. First, it was not as if all regulation was abandoned. Rather it was strategically abandoned. Wholesale regulation upon which competition depended was gutted. Retail regulation was relaxed in many respects but not altogether removed. Second, what resulted was continued forced subsidization of incumbent businesses by all, including their competitors. This occurred directly - via access charges - and indirectly - primarily by continued application of state jurisdiction and related regulations to competitor businesses drastically increasing their costs. Had all regulation been removed, the story would have played out much differently. I'm not saying it would have been better but I am saying that it would have been very very different.

My continuing concern with the U.S. is that there's a perception that everything is going to magically get better now that regulation is back. I think that's a simplistic view. It won't. This has nothing to do with the regulators, but there are issues

there. More deeply it is structural. As noted in previous discussions we, as a country, have not adequately solved the antitrust question. We have prophylactic regulation on one hand and post hoc on the other. They tend to cancel each other out. This occurs precisely because regulation is extremely complicated and sometimes contradictory both at the state and federal levels. It is no wonder, therefore, that the courts are either confused, just don't get it at all, or do get it but make decisions based upon ideological views. (That the courts in this country have been packed with right-leaning judges, particularly at the federal level is pretty much a given. It will take a long time to fix. And fix, in my view isn't necessarily "left-leaning" (though I am in many respects), rather, I think good judges, not good ideologues, are what's best for the country.)

Where that leaves us is not in the greatest place. There is vast confusion and disarray in the legal and business realms. But I still think there is incredible opportunity here. That opportunity, however, is realized by moving forward on terms other than those with which we got here. As applied, breaking up the bells again, say via loopco, is going to be a massive waste of time and effort. That's just rolling up the

seige engines to a fortified castle and pounding away. A lot of battle, but not significant progress. And it won't get done at the state level. Bell political power is just too deeply entrenched there. IMHO, we have to be more creative than that.

Tim Cowen: Perhaps is because I am writing this half way up a mountain in New Zealand, but I think we need to recall how interconnected the world's commercial system has become and how the effects in one jurisdiction have consequences in others. **To be blunt, the failure of US law to address these issues may be an opportunity for other jurisdictions to take up the challenge and to be the location for action.** The access issue in the US is a major problem both for US firms and supply chains but also for all of us and particularly for organisations and businesses that operate globally, (and I agree with those that are shocked by the politicisation of the US legal system, but that is a matter for another day).

I wanted to give an example and offer a little hope. The example is that if abusive and discriminatory access pricing in the USA is practised on the tail of a worldwide supply chain, that will have an effect on trade and distort competition on that world-

wide supply chain and between players in that relevant market. There are many supply chains that are vulnerable. Automotive production, I.T. production, banking and finance are all pretty global in their communications infrastructure needs.

The issue is then what possible legal system could apply as the place where a remedy may be sought. Different legal tests apply for jurisdiction, but the most obvious, EU law, is effects based. There have been a number of cases on the issue, perhaps the best known being the Woodpulp case. Much will depend on the facts but essentially if the distortion of competition has an effect on trade within the EU then EU law will apply to the issue. The test is widely interpreted by the courts and is essentially an issue of economic impact. Providing the effect is not de-minimis, then grounds for action under EU law would be possible. Given that communications infrastructure typically accounts for a substantial proportion of major Virtual Private Network deployments, or IT infrastructure outsourcing contracts, the thresholds may be met.

Of course much will depend on the facts, but the point I wanted to outline is that there are other laws that may well apply even if the domestic US

system is incapable of providing an adequate remedy.

The hopeful part: governments really don't like it when voters have to go abroad for justice. It is, and should be embarrassing for them that domestic justice is at a lower level than internationally recognised norms, and this type of thing can act as a powerful encouragement to action domestically. Of course the US is not quite at the level of a Zimbabwe but its position in the world rankings on the rule of law has slipped in recent years.

Customers that are likely to be particularly affected are those that operate in the financial sector given the high sensitivity to any and all costs on balance sheets at the moment. Anyone interested?

COOK Report: This is fascinating. Erik says we have to be more creative than divestiture 2.0 but doesn't say what that "getting creative" is. Tim Cowen drops a hint that it could be intercession by the EU legal authorities since the extortionate control of Verizon and ATT over the end loop enterprise customer hurts global supply chains.

The problem is that, while I am sure Susan Crawford "gets it," given the other fires

burning, what attention can she get in the Economic Council? If the FCC can find a way to address this great, but that will take a good number of months... of course EU action would take at least as long (?).

Can the affected large VPN supply chain oriented networks of enterprises get together and petition congress? Perhaps? Could BT Global Services team with Level 3 to bring last mile fiber to enterprise in the US? No idea but that would be interesting. Barack was going have a CTO. Wouldn't this be a worthwhile issue for him?

Eric Lee: This week I visited Capitol Hill and saw a senior staff person of a Member who is friendly to tech companies and CLECs. He confirmed that the ILECs were opposed to the NTIA and RUS broadband monies and plans but were helpless to stop the process, which originated with the Administration. This suggests strategically and tactically that we would do well to work with the Administration and avoid the legislative route. Instead by allowing the Administration's initiations to move both within the Executive Agencies and FCC and take the lead. This would mean, for example, to participate actively in the National Broadband plan envisioned and not allow the ILECs to capture the planning

process. Second, to find ways of using rural BB deployment to affect the policymaking process, e.g., to promote diversity of business models.

We need to celebrate the onset of this new Administration. In this, I am with Harold Feld. There is much to work with in the stimulus Act and we can look forward to more. This won't be Kevin Martin's FCC or George Bush's Congress. The ILECs may still occupy the pinnacle in Congress, but they no longer do so in the FCC and Administration.

Why Fiber is Superior to Copper -- The Importance of Technology Choice

Cole: From *Lightwave Direct*, and sponsored by an FTTH council, the "other side" weighs in:

Ovum: FTTH benefits European communities -- A study conducted by Ovum with the FTTH Council Europe concludes that FTTH positively contributes to improvements in socio-economic metrics such as quality of life, improved productivity, better

public services, and increased employment.

Robert Atkinson: I'm not surprised that a study would correlate various benefits with "high speed broadband" but why does it matter what technology is used? Wouldn't the same benefits come with DOCSIS 3.0 cable modem service or a high speed wireless service? At what point does the "unlimited" bandwidth of fiber, as a practical matter, allow applications that cannot be provided over any other technology? Or, is the report biased by the sponsor?

Cole: I was of course pointing out the sponsorship -- that does not make the study wrong, but should raise questions.

Your point about technology raises a long-standing discussion about when and if differences in degree (upload speed in particular) constitute differences in kind. I believe from about 10 Mbps upload on, the game really changes, but others disagree. Some think ubiquity is key, at least above a certain minimum (384K?). If we can count on "almost everyone" having a connection, the way we count on them to have electricity and/or a telephone, then the game changes.

Even fiber does not guarantee "game-changing" speeds;

most muni builds are rolling out at existing (top-level) cable speeds anyway -- say 10 Down, 2 Up, which I think is "too little" for true change. And, with refurbished technology (DOCSIS 3.0) or equipment, both cable and even ADSL have the potential to offer what I would deem "game-changing" speeds.

So I agree with your second point -- it is probably not the technology per se that creates the change.

Coluccio: The more direct reply would be, anything short of fiber is now considered a stop-gap measure only. Even DOCSIS 3.0 will strain to the point that, in order to make it work as advertised, extraordinary measures at the neighborhood and block levels will be required. It should be no great surprise, then, that the Cable industry itself is now shifting its focus for future builds to FTTH and PON technologies, even if the majority of the latter have been RF over Glass until now (although there's been some evidence of late that some smaller ops are beginning to look at point-to-point Ethernet and TDMA-based PONs as well. All of that said, sometimes in order to fill in the "gaps", especially in cash-strapped situations, lesser First Mile technologies are deployed (and here I'm not

trying to detract from wireless's other benefits), but these inevitably require a forklift before long.

Rood - to Robert Atkinson: It is not the unlimited bandwidth, but the near unlimited reach due to very low attenuation, in particular at low bit rates, that have provided the main gains for deploying fiber.

Fiber has been deployed first in the long distance, removing many repeaters per link compared to coaxial systems. Fiber also less expensive than radio relay links except for rough terrains and beats twisted pair transmission carrier systems hands down.

The benefits for fiber is that it allows for high-bandwidth with far less equipment and thus improved reliability, lower maintenance, and, **when leased out**, enabling a far more competitive market.

I think the entire idea that it is the high speed broadband, that brings benefits is misguided, it is the very different market structure that raises competitive levels with open access to fiber, that brings benefits.

For the matter of a transmission systems performance (e.g. a submarine cable, or a radio or satellite link), it is not bandwidth, but the

bandwidth-distance product that matters.

I think, if you want to promote FTTH, as the FTTH-council does, one should come up with new metrics: e.g.

- calculate the electric power consumption of an entire DOCSIS 3.0 based broadband USA vs a VDSL based broadband USA vs a FTTH based USA.

- calculate the required maintenance cost for the equipment mountains in the outside plant

- calculate the expected MTTF and repair cost for the cable plant of DOCSIS, VDSL and FTTH

- calculate for various fiber plants the OPEX and additional OPEX cost to change to a pro-competitive design

For an introduction to what such a calculation on OPEX looks like see:

http://www.itu.int/ITU-T/worksem/asna/presentations/session_2/asna_0604_s2_p4_jb.ppt

where the presentation of John Brouse, the Network Operations director of Charter Communications, contains that data. [**Editor's note** – it is 2004 data and the cost figures have gotten much better since then.]

The blunt point is this: **A nation, or any sovereign entity, should do an OPEX calculation comparable to**

the one discussed in the study above at least once, and then ask themselves the simple question: are they trading off the CAPEX required for rewiring against higher OPEX with DOCSIS and DSL technology paths, and is it this trade off that brings higher societal costs when compared to their capital outlay?

It is not that difficult to demonstrate that FTTH is a superior technology from an OPEX point of view, while bringing far more bandwidth. The issue was and is the CAPEX upfront, not the CAPEX deployed in several tech cycles over a typical 25-30 year outside plant economic life cycle.

Somehow I have the strong impression that all kind of advocates have been fooled by, again and again, discussing the benefits of FTTH in terms of bandwidth and new applications. I have always shied away from that "wonderful apps, wonderful services" approach and think that many other, unspoken benefits reside elsewhere.

The more hard-nosed argument seems to be:

1. A far cheaper fixed network in OPEX and annualised CAPEX depreciation

2. The opportunity to re-architect in a pro-competitive design
3. Lower energy consumption footprint
4. Higher reliability due to less active components (a dumber network) per communications link

And the dynamic enabling benefit is the opportunity to change gears from the electronics growth curve on a far steeper growth curve for optical technology.

It is this type of "making the economic case" that sometimes seems to have been forbidden in the policy fashion of the last few decades.

I have the impression that this will soon alter, when policy makers facing a severe economic turn down start to grasp that maintaining technological neutrality is a very expensive policy during an economic depression, in particular when you observe that other countries do make a switch and consolidate on a far more efficient new technological system.

So yes, you might get all that high speed broadband benefits by DOCSIS 3.0 you can imagine yourself, you cannot however get the operating cost advantages and dynamic benefits of FTTH.

Note: I have never been asked by the FTTH Council to

write a report for them. This view might not be fashionable and deployable for a lobbyist group. Touting new broadband applications and services is fashionable. Large opex cost reduction and improved competition as economic advantages tends to fly in the face of some of the FTTH Council's prospective customers.

Coluccio: Hi Hendrik. That was an excellent set of perspectives. My earlier reply in this sub thread was far more abbreviated, but assumed many of the same aspects of fiber's enabling characteristics that you presented. Sometimes being too close to something results in a tendency to trivialize or gloss over the fundamentals, but those fundamentals need to be aired nonetheless. They often go unspoken due to one's falsely assuming that they are already well understood, when, in fact, they are not. Incidentally, your post comes uncannily close to serving as a rationalization for my business plan. ;)

Earlier **Frank A. Coluccio:** Tim Nulty isn't pulling any punches here, echoing truths that have been aired here many times over the years (hat tip: Jim Baller):

Tim Nulty: **"The very widespread, duck the-issue viewpoint that, 'We need broad band but we**

shouldn't be an advocate of any particular technology, let a thousand flow-ers bloom and the market select which is best,' is a crock!.. Broadband equals fiber. Let's stop beating around the bush. Indeed, let's stop using the word 'broadband'. It has been abused to the point of uselessness,

Newman: I agree with this point of view. Claiming that one should be "technology agnostic" is at best a red herring. In the case of technologies that do, or do-not, enable a whole new generation of applications, one has to focus on the best technologies.

And there is the time-dimension to this process. If you delay the deployment and exploitation of the best technology at any stage, you delay the understanding of requirements and the orientation towards the next generation of "best" technology. So over the longer term, spanning a few generations, an attitude of claimed agnosticism towards technology leads to a delay in progress, and hence a penalty due to lost enablement of the population, that grows exponentially over time.

Kelly (BT): I need to declare an interest - terminal fatigue with this ongoing and increasingly sanctimonious debate - OK, I work for an in-

cumbent; typically organisations that invest shareholder funds against returns. If 'telcos' hadn't launched 'fraud-band' on copper in the early noughties (flawed thought it may have been at the time and since), we would still have niche web content enjoyed by a small number of insiders and kids sharing music illegally ... a movement strangled at adolescence ... not the mainstream market we enjoy today.

Goldstein: Maybe in the UK. Certainly that's not how it went down in the US.

Dial-up Internet became a popular public product starting the 1992-1993 time frame, with real high-volume growth by 1995. The ILECs here ignored it, occasionally making efforts to crush it. TA96 allowed CLECs to build the capacity for high-volume incoming calls; that pulled the ILECs' bacon out of the fire as dial-up boomed for the rest of the decade. (I was working on AOL net from 1995 to 1998 so I saw the growth rate, and how hard it was for the slow-moving ILECs to add modem port capacity to their network.)

Cable companies started experimenting with mass market-oriented cable modems in the 1994 time frame (noting earlier niche products going back to 1982), and as HFC rolled out (mostly done

1993-2001), they added cable modem support very quickly. Some enterprising folks figured out that telco alarm loops could be used for DSL, and they were the first to offer DSL Internet. (Some ILECs had tried and abandoned ADSL video in the early 1990s.) But ILECs put a stop to that, blocking the loops. With the 96 Telecom act, CLECs could get unbundled loops, and they rushed to put in DSL.

Then, gradually, dawn broke over Marblehead. The ILECs started rushing out their own CO-based DSL in the 1997-2000 time frame, largely following cable modem footprints, less likely to go where cable hadn't. And while the rules required them to make it available wholesale, they often priced the wholesale above their retail. So they forced out the independent ISPs. But only after the market had been proven.

Budde: I think the problem with our group is that we indeed are preaching amongst the converted, but one or two layers under us this debate is real and whether we like it or not we will have to repeat and repeat this message till we drop dead. Why because some of the counter forces have a vested interest to spread Fear Uncertainty and Doubt. While they publicly will support FttH at the same time they will do everything they can

to delay FttH for as long as possible as they want to maximise their returns on previous investments. We might not always like this but it is a reality and politicians are often sympathetic to such arguments, and I can see their viewpoint.

What I am arguing for is that these incumbents sit with the rest of the industry and the government around the table and discuss how to best implement FttH (make that a given). In that discussions the bottlenecks and stumbling blocks should be made clear and this is where the government can assist with policies and funding. In such an open debate I have no issue with taking on board the financial realities of the incumbents as well, but let them put their issues on the table so we can discuss them and find solutions.

A total lack of trust makes it impossible for us to have such a discussion. Some countries are now moving in the right direction and in relative terms there is more trust in the UK, New Zealand, Netherlands and some of the Scandinavian countries.

As for the US, perhaps we should look at how we can better engage with the industry to assist this country also moving towards

such an environment. We did so with the Obama Transition Team, how can we build on this and move forwards?

Cecil to Joe Kelley: Outside of Bruce Kushnick - whom I know well and won't mind me giving him a hard time <g> - I don't think anyone here truly believes there's anything inherently evil about ILEC or any other form of telecommunications business. I think both Fred and Paul make fair points which to my mind boil down to the fact that **any anger directed at any incumbent business per se on this list is misdirected to the extent it is really frustration with market and regulatory design. Outside of documented bad acts there are no "good" or "bad" actors here. All are rational economic actors given their business models and methods.**

I think the greater point, too often lost in the specifics, is that the overall regulatory design is simply out of date. From this lawyer's perspective, for example, the regulatory system in the U.S. is deeply flawed not necessarily because I think it does favor incumbents, but primarily because it is so imprecise and unpredictable to the point of having very little internal consistency. This is actually worse than always favoring

incumbents, but I think we can do a whole lot better than that. And I don't think we do any incumbent any favors by keeping them stuck in incumbency. Doing so only makes the inevitable changes in technology, business, cultural and monetary cycles more sudden and painful. I think what we all desire is a system that's relatively transparent, predictable and fair (or at least equally unfair enough to count as fair).

As applied to the policy question writ large, there is apparently a growing consensus amongst the 10-15% or so of folks on this list who routinely pipe up (as one can only guess with regard to the silent majority) that the regulatory world as it is no longer provides sufficient societal returns. Incumbency in any business tends to preserve whatever status quo was in place at the time that it was created. Corporations, even innovative ones, tend to lock into place certain business methods, and approaches to all sorts of things, including use and deployment of technology. Thus, 'drinking the kool-aid' but you can't have daily change either.

Still, if we live in a world where it makes more business sense to continue to deploy very old circuit switched services because the subsidies - both direct and indirect - are so great to make con-

version to DSL, for example, insane, then one might question the sanity of a system that makes it so. Were I running that business, I'd fight any requirement to open up that plant to anything other than that which maximized my shareholder's return, as would anyone. As insane as it may be, if it's profitable (and legal), then sure - rational business people will do it. This is one of many examples. And there are counter-examples, of course.

All in all, therefore, none of this is personal to any business, or individual as much as it a collective expression at frustration w/ policy, economic and regulatory assumptions that did not deliver as promised. Ultimately the frustrations are with political process b/c regulation of telecommunications, is, at its core, intensely political, which fact subjects telecommunications, technology, the Internet and its participants to the slings and arrows of outrageous fortune.

Speaking of such fortune, based upon your experience - both in- and out-of-region, what changes would you see? Are the access markets in the U.S. sufficiently competitive? State of the 911 system acceptable? Local franchising costs? Building and conduit access? Pole attachment rates? ROW and conduit fees?

Voice regulation vs data regulation? Maybe this system unchanged works perfectly from your perspective. I've stood before regulators who thought that the 1983 breakup of AT&T was a massive mistake and that all that has followed is proof. I didn't agree, but was informed by the dialogue. I don't think you'd advocate as much, but it doesn't matter. I simply don't know and would welcome your input.

So long as all minds remain open, we can have a profitable discussion.

Newman: I agree to this reasoned approach.

I do not think it is just repetition. Whatever we think is reasonable -- financially in the US, from the point of

view of innovation, strategy and competition, -- we are being bypassed. The FttH graph I sent is simple fact. And more countries are joining the fray.

I do recognize that repeating the same opinion, without new facts, does not add to the argument. But sending updates, with new facts as they emerge is something else. It is painting an evolving picture based on facts, that should change the opinions of more and more of us on the list, if they are willing to listen, and also motivated to respond to the emerging picture. Rather than to assume one "knows it all" already.

Some "know more" of the existing facts from their professional perspective and

daily experience (perhaps not in the global arena, really), others "know more" because of a honed ability (including long practice) to build roadmaps and to derive likely outcomes, in broad terms, a few years out [at least between the occasional revolutions]. Both sides have a lot to bring to the discussion...

It is so different in the [very high bandwidth] leased line market. We put out an RFP each year, and watch the competitive responses. The results are surprising - prices are evolving downward, and several non-selected vendors (and us too) are shocked by the competitive realities each year.

Municipalities Do Need to Offer Services Given the Obstacles that Vertically Integrated Duopolies Impose

St Arnaud: At one time I would have vehemently disagreed with Jim Baller on his comments below. But increasingly I am coming of the opinion shared by Jim Baller and Tim Nulty that municipalities need to provide services.

Goldstein: I'd like to think of it more as "public sector authorities", which could operate at the municipal or higher level. Bear in mind that municipalities vary widely in size. While Ottawa, for instance, has incorporated the whole county, and now includes forest and farmland within the city limits, some US states are divided into teensy-weensy municipalities with bizarre borders dating back to 19th century ethnic settlement patterns. (Look up "boroughitis".) Some are even discontinuous. The actual focus of my academic career was state and local government, and I got to know how a lot of these work. Public authorities and, in some states, special districts, can have more sensible borders. Come to think of it, how many urban areas still depend on privately-owned

mass transit? While Massachusetts has no county government and is fully incorporated into 351 Cities and Towns, there are regional Transit Authorities like the MBTA that covers Boston, the LRTA around Lowell and (I think) the PVRTA around Springfield.

But with local governments struggling and laying off police, fire, and teachers, there is no taste right now for new tax-funded obligations. That's why practical considerations may require a privately-owned entity under a new regulatory paradigm.

St Arnaud: As many of you know we have been struggling to deploy a customer owned fiber network in Ottawa where we only provide passive home run fiber and the customers can then connect at layer 0 to any service provider of their choice. We were hoping that this would get around the objections by the carriers in using someone else's infrastructure.

We have been in long series of discussions with an out of territory ILEC. The local team were very keen on our busi-

ness model as they knew that they only way they could compete with the local ILEC is by someone else underwriting the cost of the infrastructure. But when the local team took the business case back to head office in their home territory the project was killed. The senior management was terrified that other ILECs would invade their turf with a similar strategy. The ILECs are much more interested in protecting their local monopoly rather than competing in someone else's backyard.

So we are stuck looking to provide cable and Internet services ourselves. This actually requires a bigger investment than deploying the actual fiber (at least for a small scale project), because you have to purchase volume Internet transit to be competitive. But the biggest challenges is arranging for delivery of cable TV services. This is a Byzantine world if I ever saw one and much more difficult and costly than arranging for Internet service. I know that many municipal networks like Lafayette and others have run into the same problem.

What we desperately need is a North American organization than can handle all these negotiations with ISPs, TV content providers etc and provide a bundled package for municipalities to re-sell. Perhaps NATOA in partnership with NLR might be able to provide such a service?

Goldstein: That's an important point. TV content is, financially speaking, a big deal. The big cable operators negotiate better prices than little guys can get. Networks won't even talk directly to small operators, so they have to go through aggregators. Small incumbent cable operators (in the US) usually belong to the National Cable Television Cooperative, which negotiates on their behalf (something like 3M subscribers, which is a little bit of clout), but they're closed to most newcomers.

I've talked to some CLECs about the need for a similar body to aggregate demand and procure content on behalf of smaller new entrants. But it's a chicken-and-egg problem. They can't get started without content, so either they never get started, or they deal with an aggregator. And right now the volumes are pretty low, which limits their clout vs. what today's aggregators can do. I don't know of a good solution

that is practical in the short to medium term.

Baller: You're both right. These are serious challenges. Fortunately, however, obtaining content is not an insurmountable problem. As to NCTC, a recent article reports that NCTC's 1100 members -- which include most municipal cable systems -- collectively serve 16 million subscribers (<http://tinyurl.com/d5wua8>). That would give NCTC a TON of clout if it were true, but I believe that NCTC's members actually serve about half that number. Even so, that still gives NCTC a lot of clout. Also, NCTC recently ended its moratorium and began to consider new applications in January (<http://tinyurl.com/cg67xp>). If we find that it is discriminating against new municipal cable providers, it's going to regret it.

Budde: Interesting last paragraph Bill. You know the TransAct muni-system in Canberra (VDSL in operation for close to a decade now).

I think that this was for the first time that we faced a similar problem. We thought that if we could have some sort of an aggregation vehicle in the middle we could utilise a particular content and services bundle and use it not just for Canberra but also for several projects; this was essential as the Canberra mar-

ket on its own was not big enough for such a media approach. TransAct was looking at that time to either built new networks in other places or work together with other muni networks to replicate some of its model. Now a decade later we still don't have something like this. That is for a number of reasons.

(Traditional) content is another highly monopolised market and the content owners want to pick-off each 'sucker' (i.e. network operator) on its own, the game is played on their conditions only and this leaves hardly any margin for any network operator. Content aggregation is more of a media activity so its business history and skills is more linked to that market. From a business model perspective this is a very conservative and again monopolised market. Consequently, very little business model innovation and cooperation comes from that side.

With these two significant groups more or less out of their reach, the muni-network operators start trying building such content and services models themselves and are rapidly coming to the conclusion that they are not good at it. With dwindling business model opportunities they try to make money out of whatever they can and that sees them sometimes

taking on functions from ISPs which brings them in competition with these providers and that causes problems of its own.

So I would be extremely interested in the experiences from others as we have basically come to a dead end street without any clear vision on how to get ourselves out of this morass. I am not aware of any successful operation anywhere else.

Reorder the Process

It looks to me that the future has to be utilities based open fibre network. In order to built up extra higher margin revenues a value added infrastructure company/ or division should be attached to the utility operating data centres, content hosting, cloud computing, billing, network management, and offer that to individuals (UCG) and companies. In other words stay totally away from the content itself, ISP operation, etc.

Once you have this structure in place content owners and content aggregators will on their own terms start building their business models on that infrastructure and they might consider using the value added infrastructure services, but for the rest want to run this business totally separate and independent from the

network operator. Of course they will only become interested if such a network has sufficient mass, that's why very little activity is taking place on the current FttH networks as hardly any of them have reached any mass that warrant these media and services companies to start using them. They rely on the operators building a network based on the principle: 'built and they will come'. They don't want to make any commitment upfront but are happy to come onboard once you have got the network in place.

So, in most situations, only utilities based organisation can afford to built such networks and this again makes munies an ideal partner in building FttH networks, for all the 'common good' reasons that we have been discussing extensively on the LIST.

Baller: Bill, I had always hoped that you, Dirk, and others would prove me wrong, so that we could learn from your experiences and do it your way. Contrary to popular myth, municipalities in the US would generally prefer to provide only infrastructure or wholesale service. Those that provide retail service typically turned to that model only because they found that other models didn't work for them. As

Terry Huval of Lafayette puts it, retail service is the toll that municipalities must pay to get the fiber they want and need. Fortunately, in most cases, retail service has worked very well so far; in some cases, spectacularly so.

Still, I'm not convinced that there's no way to break through this barrier. Hopefully, intelligent multi-sector cooperation will provide the key that unlocks the door. Let's keep on working to make this happen.

In the meanwhile, I believe that those who condemn municipal retail service are doing a real disservice to the communities involved. After all, communities that are serious about fiber have to look reality in the face. They can't go with an infrastructure or wholesale-only model out of ignorance, ideology, or political expedience. If they want to acquire a fiber network to support economic development, educational opportunity, global competitiveness, modern health care, etc., they must do what works for them - not what may work under different conditions in Asia, Europe or anywhere else. In my biased opinion, municipalities that decide to pay the toll of providing retail service deserve applause, not condemnation -- and all the support we can provide them.

Ways of Thinking about Infrastructure from Technical and Legal Points of View

On March 10 - 11 a lively conversation mushroomed from a pointer to a presentation regarding the concept of a National Light Path Network. In the context of network architecture discussion began.

Kevin Barron: Fiber is public infrastructure; government builds public infrastructure. FULL STOP. [Editor - and later]

If we stop the petty squabbles over layers 0 and 1, we could actually see some of the 10-to-1 return on investment in BB infrastructure. **My fear though, is that even with the collapse of the market, the religious dogma of the "free market uber alles" will continue to obscure the fact that fiber is basic infrastructure (which *requires* public ROW's in most populated areas). We continue to be captives of that dogma, even while our economy melts down, our retirement plans go up in smoke and unemployment hits record highs as a result.**

The bogeyman of "big government" has been used for 20 years in this

country as a smokescreen to effectively capture it and primarily turn it into a hired gun for industry.

Government has a valuable role to play as witnessed by the "deregulation" of the financial markets that have left us in the mess we are in. Even Greenspan is calling for nationalizing the worst of the industry (although dumping toxic waste on the taxpayer is not my idea of fair play - I have some debt I would also like to share with my fellow taxpayers!).

Every sector of society has an important role to play: industry, government, individuals. None of these sectors can operate without the others, and require a balance to create a stable society. When we start to believe any one sector is more important or should dominate, we get into trouble. On the other hand, when each sector is productive within its sphere, we can perhaps reach Carlotta Perez's "Golden Age". Yes, the devils' in the details, but I don't think we can have any meaningful discussion (either here or in DC) until we acknowledge that requisite balance between the three realms.

Cole: Kevin et al -- the two assumptions (and/or "self-evident truths") in your statement are NOT fully accepted, even on this list of raving fans (although I do accept them).

Assumption/Truth #1. high-speed computer connectivity is public infrastructure. Unfortunately for all of us, this is NOT self-evident to many people including many working and/or leading provider organizations.

Assumption/Trust #2. only fiber counts, as anything else is just a stop-gap. We have people on this list who think universality at some "greater than dial-up" speed is much MORE important than high-speed, especially in remote and/or under served areas with only dial-up at present.

And even more unfortunately, we have people who are willing to accept #1, and perhaps #2, and still dream that somehow "competition" will save us. Or that "self-help" will save us. Or that "federal financial incentives" will save us. Or that "municipal/community ownership" will save us.

Note that electricity has been over the #1 hump for decades, and we still have debates about public versus private, large versus small, fossil versus renewable, centralized versus distributed, et al. at the generation, transmission, and distribution stages.

So -- even granting #1 and #2, we are NOT at "FULL STOP" but still have choices about ownership, operation, operating rules, etc. But if it were easy it would not be nearly so fun.

Barron: I wouldn't characterize it as "only fiber counts". Ubiquity is every bit as important; so in the first-mile space, wireless is a crucial player (and even in some difficult backhaul scenarios). But certainly even the most ardent wireless advocate will admit that we need to fiber up.

Goldstein: Let me add my support for that. Yes, the end game in most cases is optical fiber. But decent-speed wireless can reach a lot of places today where the same money spent stringing fiber would allow far fewer people to be reached. And there are still places where you just can't get fiber in practice (due to things like the terrain and CEQA- California Environmental Quality Act).

Investment in new copper, even coax, seems like an idea whose time has gone, though.

I'm always concerned about analogies. "Highway" is a useful one but imperfect -- we don't pay based on what we're carrying, but do pay a usage-related tax (on fuel) to pay for it as a commons. Most of the cost of road transport, though, is not in the highway itself, which is a shared, non reserved resource. It's in our own vehicles, which are autonomous. It's the ultimate "dumb pipe".

I don't see any glass having that nature. If it's shared, then it's got to either have reservation (my lambda, my TDM slot, etc.) or it has to be switched at a higher layer, which in turn implies a common protocol (like "Ethernet" datagramme frame relaying, IP, or perhaps a lambda-switching control protocol). So that requires network smarts. As such, a highway-style financing arrangement (public resource, just put there to use) seems unlikely.

Baller: Hold on. We're mixing apples, oranges, and bananas. Access to "public rights of way" and access to "support structures" are not the same thing. Welcome to the wonderful world of American law.

For example, if a service provider obtains a local franchise authorizing access to the public rights of way, that does not include a right of access to the telephone or electric utility poles in the locality. **For pole attachment rights, the provider would have to deal with the utility itself, meeting its engineering and other requirements, including its requirements for payment of make-ready charges and annual attachment fees.** This distinction has long been recognized in the federal Communications Act.

If a local government unreasonably refused to give a provider access to the public rights of way, the provider would have a claim under Section 253 of the Act. If the provider was a "cable system" or "telecom carrier" and a utility refused to give it access to its poles or sought to charge rates in excess of the maximum rates prescribed by the FCC's rules, the would-be attacher would have a claim under Section 224 of the Act.

(BTW, pure Internet service providers have no attachment or rate-protection rights under Section 224, at least until the FCC completes a pending rulemaking in which it is considering putting cable systems, telecom carriers, and ISPs on the same footing. In the meanwhile, municipali-

ties, cooperatives, and railroads that own poles are exempt from the federal pole attachments requirements under Section 224.)

But that's not all. Access to towers and other outdoor wireless support structures, is governed by Section 332 of the Act, which tries to balance the interests of wireless providers with the interests of homeowners in esthetics, property values, etc. -- which are covered by local zoning and other land use laws. Indoor facilities (e.g., wires, risers, etc.), are regulated under still other federal, state, and local rules. For example, the Marco Island litigation against Comcast that I mentioned a few weeks ago involved the interplay among the FCC's inside wiring rules, the Florida Deceptive and Unfair Trade Practices Act, Florida condominium law, and Florida antitrust law.

This may all sound very complicated, but it's really not. Those who do business in this space on a daily basis seem to understand the rules well enough. They may not like some of the rules, but that's a carol for another Christmas.

The bottom line is that one can't just loosely lump "access rights" together and deal rationally with the many competing interests involved.

Two Very Different Legal Approaches

Erik Cecil: Jim, with all due respect, welcome to American Democracy. American Law is a mess. We're here to reshape it, not to re-fit ourselves back into framing that no longer serves the highest possible goals for all involved. And I am one of those who play in all these realms all the time and understand them intuitively. And, since we're on the topic, for every Comcast case, there is a White Plains case. Governments, including municipalities, overreach just as much (and sometimes more) than corporations. But the point of this discussion is not to drag us back into the Act or the Act's definition and framing of competing interests.

The solution I specifically proposed was to eliminate the fee-generating dance - because when I'm running and managing all sorts of these kinds of litigation I get tired of the fees. And it's no secret that these disputes tend to be self-perpetuating. It is, therefore, the very basis of a framework that pits all sorts of common interests against each other for we know that:

1. ROW is dirt and space.

2. The costs of communications have gone way up but what the networks, in and of themselves do, isn't all that different. We get more, but it is more of the same.

3. **The market places premium values on applications and wants to drive commodity prices as low as possible.** Wood is a commodity. Poles are a commodity. Conduit is a commodity. Copper is a commodity. Coax is a commodity.

4. **Regulation places a premium on commodity. It drives commodity prices up. It pretends that telecom bits cost more than cable bits and has no idea of what to do with bit bits.**

This lends itself to the insanity of pole attachment battles where telco and electric utilities have taken down and laid live operational cable and telco plant on the ground or charged \$7 per pole one day and \$53 the next because of a change in law. Nothing on the poles changed one bit. Sure they had their justifications; there is never a shortage of justifications. And I'm certain that from their viewpoint the justifications were reasonable. But that's NOT the point. The point is in questioning a system that not only creates but energizes and rewards this endless internecine warfare.

Nothing could be better, honestly, than a world that worked well enough to put you and me both out of the business of fighting these battles. Why cut into the discussion and restart a thread with a refrain to American LAW? On an immediate level this misses the point, but the examples you cite actually make the point quite forcefully: there's no real progress to be had within American Law. It's broken. There is no antitrust and not much in the way of rational regulation either.

Rationally Fund Dark Fiber Utilities

Accordingly, I'd like to ask you whether you'd see immense value in a world where we could build and rationally fund dark fiber utilities, give strands to any qualified entity who wanted to provide service along the following lines:

- A.** Any player can provide service: telco, cable, wireless, ISP, muni-co whomever.
- B.** All players pay the same exact flat rate for access to fiber optic (conduit or pole - one flat rate - nationally), collocation space, power regardless of "service" provided.
- C.** All bits are "telecommunications" in the sense that common carriage applies to anything on any network anywhere at any time without regard to jurisdiction, service type, underlying provider, anything.
- D.** One flat fee to replace the insanity of dealing with seven or eight thousand local franchising authorities. Flat fee applies only to physical layer facilities. There's no such thing as channels, cable TV, telephone, or anything else because there's no need. Just light the stuff and provide what people want. **Be productive and make something new instead of endlessly fighting over 50 and 100 year old technology and business models. We have the means.**
- E.** National franchising. Not state, not local.
- F.** Access standards per Section 253 of the Act, and as codified in recent FCC cases, including the Section 621 Cases except we erase all of the regulatory silos because they are relevant only to the past and persist only in the present because we continue to tolerate them or are just too lazy to create anew.
- G.** Zero intercarrier compensation. It's all bill and keep. We got to pure peering for all networks. Mandatory. No charges whatsoever.
- H.** Building access is mandated - say under Section 201 - any provider can pull fiber in any building conduit at any reasonable time. Wireless - e.g. dish and fixed can already attach anywhere any time subject to reasonableness. There's no reason we can't do the same thing with physical facilities. The only allowable charges, therefore, are those associated with any retrofit necessary - say coring through floors to run conduit - but that work can be done by any qualified contractor. No charge for closet space, no charge for conduit space, and straight pass through on amps. Building benefits enormously by the presence of communications connectivity anyway. (There is no taking here - there is a giving - we know that because, were we to yank all communications connectivity out of a building, no landlord would be able to rent the space for anything other than maybe as a warehouse). Ditto for any city or community.

In other words we BAN any additional over the top charges of dirt and space vendors because they need to focus on the benefits of things creating rather than extracting money from things they did not create - DIRT and spaces inside of dirt or concrete - its just empty space through which one pulls fiber optic. (And let's not forget that the City of Dallas once tried to charge

Teligent - my old client; and AT&T a FEE for radio waves riding "through" city space. That's silly but had to go to federal court for a year to get resolved. But charging for photons moving through wave guides is no less silly than what Dallas pulled or taxing a nearby star for the light it casts across the landscape.)

Baller: Erik - you said "American Law is a mess. We're here to reshape it, not to re-fit ourselves back into framing that no longer serves the highest possible goals for all involved."

That may be why you're here, but it's not why I am here.

I agree with several of your ideas, and I find the rest interesting, even if I don't agree with them. *At bottom, however, I don't believe that your package of proposals has any realistic chance of being enacted in the foreseeable future - at least in the absence of an unthinkable catastrophe or a rapid acceleration of the collapse of our economy.*

For your package of changes to become a reality, federal, state, and local legislative bodies across America would have to decide that there was a pressing need for it. They would do that only if

an outraged public demanded it. Where's the rage? Polls indicate that the American public has many other, higher priorities, and it is generally satisfied with the crap that passes for broadband in this country. What will it take, and how long will it take, for the public to see the light? Three years? Five years?

Now let's suppose that we've arrived at the point at which the public was indeed ready for serious change - not the relatively minor changes that both parties called for during the recent national elections. Suppose further that this resulted in the introduction of your package of proposals in Congress and other relevant legislatures. Given the vast destruction of expectations, interests, and property rights that some of your proposed changes would produce, there's no assurance that Congress and the other legislatures would ultimately embrace your proposals. But let's suppose that they ultimately did. How long would that take? An additional two years? At best. More likely five years.

Then the litigation would begin. How long would that take? At least seven years. Maybe 15 years.

So, we're talking 12-25 years. I don't think that America has the luxury of 12-

25 years of self-immolation, while the Asian and European nations move rapidly into future. Personally, I'd prefer to devote my energies to pushing America to set ambitious goals and to implement them through bold but incremental changes to our legal structure.

Cecil: I'll concede up front that I'm swinging for the fences here. I will, therefore, concede that this is miles away from the way things are today and in conventional terms it might take a generation for all of this to happen. **But that's not really how change occurs. Rather, it bubbles along in the background for a really long time and then suddenly explodes into public consciousness.**

We are already at a crisis point. Present business models will not hold up under the strain. They can't. The money just isn't there and will never return in any sense that it was there in the first place. Moreover, your client set is deeply invested in putting conventional telecom out of business; that may not be their conscious aim - and I don't really care if it is or not; I've been around long enough to know that if anyone can do it, it's city hall. City hall is hell to fight. I'd rather litigate against bellco than take on a Section 253 case against a city; those are hell. They

tend to end up in places like the Supreme Court - and that's a long and expensive journey. Kind of have some, uh, empathy having been near that sort of thing.

Still, the cities or towns or municipalities either own or control the ROW most of the time or at least enough of it. They own the poles. They don't report to shareholders. they have a tax base, and they get to tax their competition. That's a sweet market position to have. From where I sit, therefore, I think that your guys will win. And you are the hero of that set. Meanwhile everyone hates bellco and they hate cableco too. CLECs are gone. All that leaves is wireless and the two biggest ones are owned by bellco, so you tell me how much more change is necessary before there's change. I'd really like to know b/c if that's not what's going on, I'm really missing something here.

Third, per above, I don't buy the "we can't change it will take too long, we don't have time" theme on two levels: (x) the plate tectonics have already shifted; (y) it is an artificial assumption that keeps us stuck. Apply that same reasoning to any societal change; the results, and the weakness of those sorts of justifications are readily evident.

Fourth, it is precisely these sorts of thought experiments that resulted in action at smaller levels, which changes led to much larger ones. Much of what ended up in the 1996 Act, for example, had already been tested at the state level. (And as much as I criticize the 1996 Act (and the entire 1934 Act for that matter), there's a whole lot of good stuff that came from all of it). Much of what will end up in the next rewrite will have to do with what you and your clients are doing right now. They and you do have a national agenda, are working hard at political levels to accomplish that, and have more than enough resources to effectuate significant change. I expect, for example, that legislation prohibiting states from restricting municipal entry will get through Congress within the next 4 years. I don't necessarily oppose that either, but I do think it fair to hold all parties - including governments of all sizes - just as accountable as we seek to hold corporations. [snip]

Enough people and enough smart, energetic, and committed people see the light that change will happen. It already has.

This time around the municipalities and communities will lead the way. And, again, I commend much of what you have done on behalf of com-

munity networking. I admire and celebrate it precisely because I am not wedded to any single business or regulatory design.

What I do not think is useful, however, is any perpetuation of failed regulatory models that result in incessant grinding of opposed interests that do not necessarily have to be opposed. If we look at this issue through the lens of game theory, it seems to fit the 'prisoner's dilemma' precisely because this is a non-zero sum game. We also know that cooperation is always the best strategy within such an environment. No field, it seems to me, better fits the prisoners dilemma than telecommunications precisely because all of the parts are so incredibly and intricately interrelated. It's why I posited the idea of "interstructure".

I did so to speak to the need to look at the system not as independent game pieces whose rival interests are best served by fighting those of other rivals, but rather to see it as a deeply intertwined and interrelated system whose value is maximized ONLY when all of the parts cooperate. If that is true then the only systems we should seek to create are

those which encourage incentive to align. Accordingly, for these reasons and the countless reasons cited on this list, economic literature, and the best thinking I hear coming from all directions, I cannot imagine better reasons for considering all of this more deeply right here and right now.

You (Jim Baler) said: So, we're talking 12-25 years. I don't think that America has the luxury of 12-25 years of self-immolation, while the Asian and European nations move rapidly into future. Personally, I'd prefer to devote my energies to pushing America to set ambitious goals and to implement them through bold but incremental changes to our legal structure.

I (Erik Cecil) comment: Bold but incremental? Bold but incremental is simply incremental in all caps and a really zippy brochure. Bold but incremental is what all the lobbyists used to say of any incremental rewrite of the Act - "it's a milker". Same assumptions, same process, same procedure, same results.

Bold but incremental is a field day for DC firms and consultants - nothing more than K street loves than change in legislation - damn, my inbox is nearly overflowing with newsletters from all over DC explaining why my clients

have to spend money on someone to tell them how to get this money or that, how to make the most of this change or that change, or how or why they have to get involved right now so they can "impact" this "important rule making" or this "critical legislation". There's a reason DC has more lawyers than any single state in the nation. It's not because Congress is there. It's precisely because there's money in the changes that Congress makes and the real or perceived necessity of influencing Congress, and then influencing every mini-congress and mini-judiciary - the Executive Agencies.

And let's be fair - I did that for a very long time, and am back into the consulting game myself - I will be doing it again. So I will readily accept responsibility for my past actions and willing if not enthusiastic participation in that game. I still play it and I still enjoy it; I plan to make some money doing it too. I, therefore, DO NOT think it is bad or evil or wrong. **I zero in on this precisely to beg the question of "to what end"? If we are to engage the machinery, and there's no avoiding that, what's the key change?**

BUT If you really think "bold but incremental change to the American legal system" is going to fix this, I AM ALL IN. I will support you and it

110%. Just kindly tip your cards if you could.

Defining an Agenda to Match Our Challenges

In closing, I will gently observe that you've avoided all of my questions, not responded to a single point, and pretty much told me that change takes forever so don't even try unless it's "bold but incremental change." So I'll take it up on your terms. Which of the following - or subset thereof - are principles and/or concepts where we can find agreement? Which are simply insane and will never work? Which are neither, but perhaps interesting? Where is it that cooperation would render the greatest benefits to your clients?

For example, they certainly have to hate the fact that they cannot get systems built without worrying about litigation from bellco and utility co - it is out of that very concern that I thought that where a city or municipality could give away fiber optic, you'd avoid all kinds of trouble. I can imagine standing up in court or hearing and saying, "Your Honor", or "Commissioners" "Company XYZ complains of a taking. Honestly, I'm astonished. What we have here is a

giving. Allow me to explain ..." It seems to me the incremental cost of additional strands will be minuscule in lost money, time and profits relative to litigation. Even if there is litigation, I'd rather open the case with a giving than one that looks like city is competing against private business. And honestly, this game is not really won on who controls the facilities - in the narrow sense yes, but in the bigger sense - in the 20 year horizon - it's won on providing people what they want; it's won on customer service; it's won on abundance.

1. Build conduit and dark fiber.
2. Provide fiber to any qualified entity on equal terms - whether muni, bell, CLEC, ICO, ILEC, Wireless, or ISP.
3. Level the fee field - one single rate on all poles, all ROW - make it national if possible. Rates are mostly arbitrage and mostly arbitrary; the only thing that gets proven at cost proceedings is who has most money to purchase the best experts and lawyers to put on the biggest case. No state or federal agency on the planet can swallow, much less digest the volume of information even a relatively modest operation can put into the mix. Com-

modity pricing for commodity inputs; let's get the incentives to disguise monopoly rents as something other than extortion "fair cost" "actual cost" "property interest" out of that game. Those things only make sense in a world where we agree to assume without question they make sense. We certainly can realign interests so that the incentives are all pointed toward lowering cost instead of justifying it.

4. Communications networks are a national priority. Give all municipalities rights of condemnation on zero or near zero cost access to locations inside of buildings necessary to run conduit for purposes of providing fiber optics. **The statutory offset - to prevent takings objections - is the value of having fiber optic available to their tenants.** I can tell you from direct experience most building owners are thrilled to have good fiber optic connectivity and they mostly welcome it. Maybe we don't need this part at all; perhaps there's another way to do it - some sort of offset to encourage the behavior so the 5th Amendment crowd's hair doesn't burst into flame.

5. Sooner or later we'll have to level the playing field with intercarrier comp but honestly that's not a priority in my mind. That playing field will level itself - it's con-

stantly at war; the participants on all sides have flipped positions so many times that sooner or later an insightful judge, jury or regulator is going to decide the case that just takes it all down. One good shot to the filed rate doctrine and we'll see the entire thing come down.

6. Level the siloes. They are mostly irrelevant to where technology, people, and markets want to go. Bits equal telecommunications. Single minimal rate. (This also means, that I think Fred's ideas around loop co have much to commend to them; so another way at this is to add to loopco poleco, ROWco, collocation-co under one single, non-retail, commodity provider entity - and perhaps we have meaningful change instead of change chasing meaning.)

Thanks in advance for your attention. You've been more than kind in tolerating my questions.

Paul Budde: Hmm tricky one Erik and Jim

Those who have followed some of my arguments will guess that I do lean towards Erik, if we are not bold and don't challenge the current system it is going to take also at least 15 years before we see changes, so there is in

my opinion not a lot to loose by being bold.

However, in a rational way Jim is correct unless we can either create massive uproar or unless something very dramatically happens, the chance of creating the change or even part of the change, as so passionately expressed by Erik, won't occur.

The way I would like to handle this is to maintain the more radical call for change, of course in the hope that we will be able to get at least some of this implemented at a minimum we will get some people thinking. (As a matter of fact we do already have some very senior people in the Administration at least interested in these more radical approaches.) This will create cracks in the system and than when that happens we need to be opportunistic on how to wiggle ourselves into these cracks to establish changes and that will be a matter of give and take, more along the lines Jim is talking about.

In one of our previous discussions we discussed the pros and cons of structural changes vs. grass root changes (e.g muni broadband). If, through guerrilla warfare, we could get FttH

muni penetration reaching 10% or so I bet we will start seeing cracks appearing in the incumbent fortresses and that would accelerate the change enormously. Consequently our bold activities when well executed don't have to see the system start to change overnight, we can be the lever that actually allows the system to implode on itself. **You can see in some of the European and Asian countries how quickly things can change once the incumbents decide that they have milked the old system for long enough and start moving toward a more open network environment. It can happen rather quickly (2-3 years)**

I think that its great that we have both sets of people with the LIST and in a combined effort I am sure that we can create change and that we will be able to get this sooner rather than later. If it comes to the crunch we will also have to sit down and negotiate the best possible outcome at any given time and than continue to create the next crack and the next victory. It certainly will not be easy, but hey that's why we are THE LIST!

Baller: Paul, let me be clear. I am not in favor of guerilla warfare. I do not want to destroy the incumbents. I am not eagerly awaiting, let

alone working toward, the implosion of America's telecom system. We are not in this together.

Furthermore, I have worked with a great many American municipalities, but I cannot name even one that thinks the way you and Erik apparently assume municipalities do. Municipalities are interested in advanced communications networks, not for their own sake, but primarily as vehicles to achieve critical community goals - including economic development and competitiveness, educational opportunity, public safety, affordable modern health care, energy efficiency, environmental sustainability, good government service, preserving free speech and democracy, digital inclusion, and many more. In my experience, municipalities would almost universally prefer to achieve these goals by working cooperatively with the private sector. They take matters into their own hands only if they have concluded, after careful study, that doing so is their best option.

You and Erik may have given up on the private sector and the American legal system, but I have not. You may ultimately turn out to be right. I respect that. If so, I will have failed, because I'm working to achieve more, not less, harmony between the public and private sectors,

including the incumbents. If I fail, I'll move in a different direction, but at least I'll know that I did my best to give cooperation a chance.

Cecil: [Jim:] None of what you say below comes even remotely close to points I've made. I'm not sure whom it is you are responding to or what it is you wish to clarify or rebut. We are, at this point, as two ships passing in the night.

I'm also really glad you haven't given up on the private sector. Considering your primary client base is government and mine is private, I'm thrilled to hear that. To be fair, I've not given up on government either.

And, for the record, private sector vastly prefers to work cooperatively with the public sector. Litigation is time consuming, uncertain and expensive. Unlike the public sector, every day in litigation is lost money - network not deployed, customers not served, customers leaving for another provider, investors scared off, shareholders upset, plaintiffs bar looking for a reason to sue your company, and so on. Like you and your clients, private sector only goes to litigation after deep and careful consideration and deep and extensive study. I can't tell you how hard it is to make those decisions. This is also be-

cause for the private sector, losing can mean losing a market and going out of business. And it always means losing money.

I'm also very happy and encouraged to hear that you've not given up on cooperation either.

The Overton Window

Goldstein: I'll stick up for Erik now. I don't agree with all of his positions, and I agree with Jim that Erik's positions can't, en bloc, be implemented. But maybe that's not what we need to be worrying about all the time.

There's a political concept called the Overton Window. This represents the scope of ideology that is considered acceptable in public discourse. Ideas outside of the window are considered "fringe" or "extreme". The Wikipedia description includes this passage:

The degrees of acceptance of public ideas can be described roughly as:

Unthinkable
Radical
Acceptable
Sensible
Popular
[Policy](#)

The Overton Window is a means of visualizing which ideas define that range of ac-

ceptance by where they fall in it, and adding new ideas that can push the old ideas towards acceptance merely by making the limits more extreme.

This is important in framing issues. The Overton Window moves. You start by pushing an idea that is outside of it, and it sounds nutty, but things that had previously been seen as radical suddenly look more acceptable by comparison. Newt Gingrich and Karl Rove got this, as did other Republican leaders and think tanks, and used it to push discourse far to the right. They knew that they couldn't get their proposals all implemented, but they moved the window. Clintonism, in contrast, was based on triangulation: Try to find a middle position acceptable to the broadest majority. It doesn't move the window, and when it runs up against a Gingrich, it loses.

So let's talk about radical and even unthinkable ideas. If we talk about the unthinkable long enough, our ideas that are merely considered sensible will have a chance at becoming policy. And it can keep moving.

But we also have to recognize when we are doing window-sliding, and when it is time to cut a deal, hopefully at a cen-

ter that has moved our way.

In telecom, let's face it, the extreme proprietarian view (Big Ed Whitacre's model) was not on the table in 1996, but by pushing it hard, stuff we liked, that was policy and popular in 1999, is now seen as radical.

COOK Report: So Erik your points 1-6 are major points for national clear the playing field legislation?? They would replace the regulatorium?

Sounds good to me. But it would horrify my wife who simply does not understand why her husband hates the phone company. And I think my wife is unfortunately like too many other Americans.

Now I have an idea. Explosive change is needed. You Erik see that change as being more likely than Jim. And who knows how close we might be? You exposition of the Overton Window concept I find personally much more attractive than triangulation. It would be terrific to find a Bernie Madoff in the telco fiefdom don't you think? R%aise the overall temperature.

After talking with Tom Alibone it is very plain that Verizon in New Jersey is racheting up its monopoly rent extraction in the midst of the meltdown in New Jersey. It actually is considering its

FiOS arm to be an independent competitor. Tom showed me further evidence. And it is running a guerrilla campaign to convince the league of municipalities to petition the state legislature to place the 7% sales tax on all video. Essentially unregulated it can pretty much do as it likes and what it likes is what ever it takes to keep its executives salaries flying. It can raise rates, charge what it chooses. Defy its customers to waste their time disputing bills. It stinks.... but where is Verizons Bernie Madoff when we need him?

I talked with Fred yesterday and Verizon has a different strategy in Massachusetts. But the basic theme... monopoly rent extraction is the same. ATT and Qwest are no different. Now I have done FOUR HOURS with Frank Colucio over the past 3 days. And Frank has described there in in excruciating detail the huge waste of ICT architectures in urban areas..... an accretion of the last 20 years of the consequences of copper in the last 100 meters rather than fiber. The economic opportunity here to reclaim lost real estate, less energy consumption less equipment cost and less staff cost is HUGE.

BUT getting people to see it rests on paradigm change on the order of Gallileos challenge to the church. There is

a multilevel orthodoxy that has accreted over the last 20 years. A faith founded in best practice standards libraries and firmly rooted in building design and construction.

<http://www.baselinemag.com/c/a/IT-Management/A-Day-of-IT-Reckoning/>

This article talks about all the techniques that the expert warns ICT managers they will have to get used to survive the meltdown. All the techniques save one. Namely begin to remove the copper LAN infrastructure from your buildings replacing it with fiber. This is the big kahuna.... but it goes unmentioned because it undercuts the assumptions of the basic faith. Levels the silos and forces the entire organization to begin a radical rethinking of its operational and financial structure.

It is I submit the kind of change you Erik are calling for. CLEAN slate. might it begin to take hold give the economic exigency that enterprises will face? Possibly... no sure thing.

I like Fred's Overton Window concept. A paradigm shift has a chance of being heard in a time of crisis where we are now. And what is fascinating is that while we have a lot of folk here beating on the issue of the regulatorium,

and deservedly so, there is another possibility for equally profound change, economic and environmental that is possible if one begins to lay out the technical feasibility of fiber to the desk designs in enterprises and begins to show people that the previously unthinkable.... IS POSSIBLE.

Erik is saying think and act on the previously unthinkable - that its a huge opportunity. Frank is saying that fiber in the last 100 meters is a huge opportunity - one facing the obstacle of a different entrenched paradigm but one that those challenging it can use for disruptive advantage. So the common thread that i see is that given a crisis there are opportunities for moving the window.

Where is the spark? ISKRA in Russian. Should we pray for a revival of the Hepting class action? or for a Madoff? Or some other tremor?

Cecil: Gordon, Jim is exactly right to point out how hard it is to change minds. Like Jim, I've got a career into that sort of thing; I empathize completely. But mind space is interesting. Once minds begin to change all kinds of serendipity is released, or so

I choose to believe.

My points 1-6, by the way, are vastly less ambitious than national legislation. They are more in the vein of whether we, as a group assembled here, can clear our minds and agree on anything at all. If this group agrees on that - or on some set of principles we can put on a single sheet of paper, I'd find that fact alone profoundly encouraging.

So the question is pretty simple. Does it make sense to aspire to a legal regulatory system that elevates cooperation over conflict? If so, do the points 1-6 below point toward that destination?

My guess is that Jim is right. We'll end up just going at it hammer and tongs all over again because so far, for this species, battle is progress. Maybe next time ... but meanwhile, I'm going to catch a late night rerun of Planet of the Apes b/c I want to feel better about all of this.

Coluccio: Jim, I can't speak for Paul or Erik, but I, for one, have not given up entirely on the private sector in the same ways you are ascribing, although the dominant players in the private sector require a great deal of collaring, and in some cases

they need to be defrocked outright. I say this because, with only some rare exceptions, the incumbents, like the unwitting enterprise industry IT participants that Gordon cited me as describing earlier today, spend far more time and energy delaying progress than they do promoting it.

And earlier -- Erik: where were you four or five years ago when I suggested setting a match to the regulatorium and calling it a day? I couldda used your support back then ;)

Cecil: I was inside a crucible called a hearing. It and hundreds like it taught me that if you really want to get away with something for a very long time make sure it is regulated. That, by the way, is not the American Legal system. It is a poor facsimile thereof whose time upon the stage is full of sound and fury signifying nothing. Let us hope we are wiser than to choose again those yesterdays leading to dusty death.

Behavioral Economics Versus the Chicago School

Chris Savage on March 12: Denizens of the list will recall that I have been ranting for some time about how

(a) the findings of behavioral economics undermine the intellectual foundations of "the Chicago School";

(b) the financial meltdown of fall 2008 represents the political death knell for the "Reagan Revolution" that established Chicago-School thinking as the dominant regulatory paradigm for the US; and, therefore

(c) we ought to be thinking about what regulatory paradigms and policies for telecom/Internet make sense in light of what we now know about economics.

Bob Atkinson and CITI have now put together a full-day seminar beginning to examine these issues!!! Mark is an early panelist. I am not sure that I personally will be able to go but if anyone is in New York and can, it looks like an awesome program. If I can't make it, I'd appreciate it if someone would ask Jerry Hausman about the difference between the mathematical predictions of equilibrium and/or optimum states, based on assumptions of how people will behave, versus welfare-based policy pre-

scriptions arising from the assumption that people's behavior in fact maximizes their own welfare. At a high level what behavioral economics and related empirical analyses of decision making show us is that people's decisions depart from self-welfare-maximization in systemic and predictable ways, which suggests that over time organized sellers (e.g., big, rich companies) can strive to structure decisions in ways that take advantage of those departures. Hence one cannot make welfare conclusions from people's decisions. Hence one cannot make affirmative, market-based policy prescriptions as a general matter.

Here are the details:

<http://www4.gsb.columbia.edu/citi/neweconomics>

Cole: Some of the differences were set out years ago in an article in *Science* magazine, later turned into the book: *Judgment Under Uncertainty: Heuristics and Biases* by Daniel Kahneman, Paul Slovic, Amos Tversky - 1982 - 574 pages.

In the political sphere, Graham Allison's book *Essence Of Decision*, based on an earlier article in the *American*

Political Science Review, identified 2 alternatives models to "rational decision-making" which he called Model II (based on organizational behavior) and Model III (based on political analysis pioneered by Richard Neustadt). (I was Allison's research assistant and Neustadt was chair of my dissertation committee.) All of these yield extremely interesting differences in both prediction and prescription vis-a-vis "economic/rational" decision-making. A very interesting concept for a conference!

Savage: Right. Picture economics progressing on one timeline. To the extent that economics generates results that are not totally obvious, there is a lag, measured at least in years and sometimes in decades, between the time economists figure something out, and the time that the "new" understanding of things gets implemented in public policy. So policy proceeds on a timeline that, as far as economic insight goes, is lagged by 15-25 years from the development of economics.

The "Chicago School" guys were doing great work starting in the 1950s and 1960s, and it got a number of them

Nobel prizes. 15-25 years later, their ideas had become the dominant economic paradigm underlying antitrust and regulatory law. Meanwhile, starting in the 1970s, and going on into the 1980s (and continuing), Kahneman, Tversky et al. started publishing their work, which conceptually undermines the strong(er) conclusions of the Chicago School. They are getting their own Nobel Prizes now.

So here we are, 15-25 years later, and what do we see happening in the policy sphere?

Cole: Tom Schelling just got his Nobel prize a few years ago -- maybe we have a bit more time?

Savage: Schelling was one of the readers for my undergraduate honors thesis in Economics and Sociology in 1977. My first key point was that when potential customers are ignorant about key features of a good or service they are considering buying, the determining factor will be the degree to which the customer trusts the salesperson with whom they are dealing - an unquantifiable social phenomenon, not a strictly economic one. My second key point was that people have at least three and probably more separate dimensions on which they evaluate potential courses of action, which

means that Kenneth Arrow's Paradox of Voting, in which a group of three or more voters can express intransitivity in their collective preferences (who wins depends on the order in which things are presented), meaning that one cannot derive conclusions about consumer welfare based on the consumer's observed behavior. (Arrow was my thesis advisor.)

Gregory Rose: Isn't the point of Arrow's Impossibility Result the fact that violation of any of a number of theorem's assumptions resolves the intransitivity (and several are violated routinely in group decision-making, particularly the unrestricted domain and independence of irrelevant alternatives assumptions).

Savage: His key point was that if you have at least three voters and at least three possible courses of action, it is possible for their preferences among the options to be such that they will collectively show intransitivity. Option A wins over B, B wins over C, but C wins over A.

The point of my paper is that if a consumer has three or more possible courses of action, as well as three or more separate internal valuation schemes, the consumer can internally reflect the paradox of voting.

From the perspective of a seller, that means you want a choice architecture (to use the current term from Nudge) in which the choice "buy my product" is what comes out on top.

From the perspective of a policymaker, it means that if utility is in this fundamental sense multidimensional, you can't make unambiguous welfare conclusions based on consumers' observed buying behavior. So much for "revealed preference theory"...

Rose: I'd be interested in seeing the specifics of the theorem.

Savage: My economic argument was not the generation of a new theorem: it was an explanation that Arrow's own existing observation regarding voting would apply to individual decisions once you move beyond the neoclassical formulation of people as unidimensional "consumers" whose decision processes and welfare experiences were really just a black box. For a hoot one time, in connection with a different economics project, I parsed Samuelson's introductory economics text, looking for references to "people" or "pleasure" or "enjoyment" - of which there were none. The entire discussion was phrased in terms of behavior and assumptions as to how that behavior would change in response to

price changes. Which is fair enough, I suppose, in a Skinnerian/Watsonian sort of way. But now that we aren't afraid to talk about brain processes, experiences, etc. any more. It seems fairly sterile.

My other argument was in the form of a game (in the Schelling sense) which I devised, in which one's optimal strategy depends entirely not on any information available in the game set up, but rather on whether you trust someone who you know to be a confederate of the experimenter and who you know to have some economic incentives possibly, but not necessarily, adverse to yours. I'll dig up the paper to recreate that and post it later on.

The Real Culprit - Clinton Captured by Wall Street

Earlier **Cole**: Actually, although it started with Reagan in 1980, the biggest steps forward were with Gingrich in 1994, so by that reckoning, Obama is the "new economy Reagan," and we are yet to have the "new economy Gingrich" who on that timetable will not arrive until 2022 (14 years after the related President).

Cooper: Reagan and Gingrich are villains, but we could expect nothing less.

The real culprit was Clinton. Throughout the 1980s, Reagan could not do his worst because he faced a liberal Democratic Congress. After campaigning as a populist, he was captured by the Wall Street crowd. He signed the legislation that ripped the heart out of the New Deal gleefully. His thugs, Ruben, Summers, Geitner and Gensler, backed up by Greenspan, opposed any effort to preserve prudential regulation in financial markets.

Joe Stiglitz, former chief economist at the World Bank, Chairman of the Council of Economic Advisors under Clinton, and later Nobel Laureate, fought Ruben tooth and nail, but Clinton always made the wrong choice, Stiglitz wrote a devastating critique of Rubinomics ("the Roaring Nineties") in 2002.

I have been giving a speech generically entitled "The Collapse of Market Fundamentalism" with applications to various sectors (finance, commodities, electricity) including telecom, which I gave at Bruce Kushnik's event last week. I will give a version of it at the CITI event.

Savage: I'm not sure I'd be as hard on Clinton as you seem to be; he never won a majority of the electorate, and his first attempt at doing

something progressive - health care - was, yes, screwed up, but also totally slapped down. Then came Gingrich, and he was fighting Congress every step of the way. So Clinton was kind of hosed from the get-go.

Cooper: On these issues he did not fight with Congress.

Savage: Fair enough. I'm not going to mount any spirited defense of Clinton; I just think that anyone trying to do anything really progressive in the mid- to late-1990s was still fighting the tide of intellectual and political history, so I cut him more slack than you do.

Cooper: Agreed.

Rose: Bounded rationality theorists have been working on this problem [welfare-based policy prescriptions] for over fifty years. While there is certainly room for disagreement, I think it is intellectually cavalier to dismiss a complex set of ideas because what amounts to a *reductio ad absurdum* "toy" model turns out to be absurd. I don't want to sidetrack the discussion into technical issues outside the general interest of the list, but the two assumptions you cite don't reflect the depth of current thought on welfare theory in behavioral economics. There are real and complicated problems here.

Savage: I think you are slightly misunderstanding me. I'm not saying that neo-classical welfare theory is wrong because I can create a reductio based on a simplistic model; I'm saying that **the basic premises of neoclassical welfare theory do not seem to me to correspond with what we are coming to understand about how people actually make decisions and experience happiness, unhappiness, etc.** The toy model illustrates how a simple step in the direction of realism (multiple, inconsistent rankings of outcomes in the same individual) can cause the theory to break down. That doesn't mean it can't be patched up again in a more sophisticated version.

To put my two assumptions affirmatively: (1) I think that it is very common for people not to know, and not really to be able to discover, what will make them happy; and (2) I think that people have a hard time making decisions that rationally use what (sometimes, little) they do know. As a result I question the normative side of economics that makes the leap from examining and modeling how people actually make decisions to policy prescriptions that proceed from the view that those decisions are the best indicators of welfare even for the individuals making them, much less for soci-

ety as a whole.

My deep reservations about normative economics do not at all lead me to fail to recognize that there are, as you say, real and complicated problems in this area and that there are some very smart, dedicated people working on them. Partly I'm just tired, after 30 years, of having "market" outcomes played as a trump card in complicated policy arguments.

March 13 **Cowen:** I have argued with economists for 20 years and think it is a worthwhile sport. A few thoughts about the direction that policy is going in the EU and how that is relevant to this list below.

In the EU there are an increasing number of references to "behavioral economics" being made in speeches by policy makers and indeed EU Commissioners. For example:

europa.eu/rapid/pressRelease_sAction.do?reference=SPEECH/08/660&format=HTML&aged=0&language=EN

I attended a presentation by Commissioner Kuneva at Kings College in London University last month where she was speaking about her approach to competition policy; she explicitly referred to behavioral economics in her re-

view of the importance of the market as part of the mechanism that drives the EU. Interestingly she also made a point that is critical for all to understand: the market mechanism is the cornerstone of democracy and given her personal history of growing up in Eastern Europe, she felt strongly that the market is vital to ensure personal freedom, particularly from the perils of state control.

The basis of US anti-trust came from similar thinking: anti-trust was created to bust the trusts that developed as a consequence of free market capitalism in the late 1800s and strip power from the robber barons. The trusts were seen at the time as a threat to democracy as too much power was concentrated in too small a number of hands. The market wasn't working to provide opportunity innovation, growth, and personal freedom. It had become controlled and there was a threat to democracy.

I have provided the link to the behavioral economic conference held by the Commission on the 28th November last year above. In the comments that Commissioner Kuneva made, she drew out the thought that consumers do not always act in their best economic self-interest. This is a challenge to much of

classical economic thinking or at least to those people who read Adam Smith's *Wealth of Nations* without reading the *Theory of Moral Sentiments*.

Approaches to the “How To’s” of Regulation

Consumers have been shown by behavioralists to value many different things and their value systems drive their choices. As the Commissioner pointed out, this is important in formulating the extent and degree and the way in which regulation should be implemented. She drew out four different issues:

1. **Default Bias** (in which when making decisions we default to a previously successful behavior or rule),
2. **Framing**, (weighing losses above potential gains leading to risk aversion),
3. **Present Bias** (or one in the hand is worth two in the bush),and
4. **Choice Overload**.

In particular she pointed out that these issues are critical for the regulation of industries such as telecommunications and energy. The thinking and references are useful in determining predictably irrational

decisions. They inform policy makers when thinking about entrenched monopoly. For example understanding of the default bias idea is important when thinking about the extent of entrenched monopoly. It is also a bit more accessible than talking about switching costs or loosely covering different motivations with the redefinition of common phrases and words.

(I have always had a problem with talking about people's motivations in terms of 'utility' and other such expressions that are used to redefine commonly used language to mean the exact opposite of normal usage. How can happiness be encompassed by the word utility? To say that a social worker, often motivated by caring and feeling for common humanity is motivated by personal utility indicates more about the cynical mentality of the analyst/economist that has to see everything in terms of personal benefit, than the reality of people's motivations. We can define a pot as a pan handle but its still a pot to most people).

Social good and public goods are at the moment being redefined and the role of the market is under intense scrutiny. I am not one that says Greenspan got it all wrong, but I don't think that Keynes was all wrong either. These issues are central to the

regulation of telecommunications as those policy makers that are concerned with outcomes need to understand all aspects of market failure, not just the ones we have seen before.

March 17: **Wedeman:** Wow. What a great conversation. I wish I had been paying attention last week when it was going on.

One thing that nobody seems to have mentioned is that Kahneman is and Tversky was (he passed away in 1996) psychologists, not economists. The fact is that although most associate psychology with personality theory and/or psychotherapy, these together constitute only one (of many) aspects of a vast field, with a rich tradition of rigorous empirical research -- research on the way people actually behave rather than on the way someone expects they should, in theory, behave.

Even though the dialogue happened last week, I may feel moved to offer a few more comments, even to a conversation in which I am the only participant. This stuff is just too good to pass up.

Cowen: I agree that this is an important issue since it gets at the 'how to's' of regulation.

One issue that has arisen in relation to economic assumptions is in relation to the basis for regulation. I am in Brussels today discussing these issues with the European Commission and others. One way in which the issue has been described is that the issue is not about economics but is about the extent of the discretion that is afforded to regulators in regulating regulatees. In sum, even if classical economics is not perfect that is OK since it provides a predictable basis for regulatory intervention. The issue with behavioral economics is that it demands information and actual evidence. That requires a detailed enquiry and rigorous testing of evidence. That takes time. The argument goes that time is in short supply and its better to get on with decisions even if on a flawed basis if it is predictably flawed, because markets need certainty.

My personal perspective is that any justice system based on demonstrative nonsense is likely to fail, and the thing that needs fixing is the certainty/evidence/time paradox. This is not new since any justice system makes trade offs between the amount of time and the availability of evidence which has an effect on the quality of decisions. Part of the fix would include the anticipated accumulation and disclosure

of accurate information such that regulators can make good decisions. This requires regulators to look forward and to anticipate the sorts of information that will be needed. That would allow speedier decision making and solve the certainty/time/quality of decisions paradox issue.

What Behavioral Economics Tells Us

Savage: What behavioral economics tells us in the short run is several things:

(1) Intelligent regulation is not easy. You actually have to know the devil is going on in the markets and industries you want to regulate.

(2) Actual people, in their role as consumers, make systemic and predictable errors in making certain types of decisions. (Someone listed some of the key ones, e.g.: status quo bias, option overload, fear of loss of current benefits more than outweighs equivalent potential of gain). It follows that people who would sell to them have systemic and predictable incentives to exploit those errors. These are not problems that are fixed or avoided by nominal "competition." Therefore, from a consumer protection

standpoint, one should focus on the details of consumer transactions more than we have.

(3) If, as noted above, firms can maximize their profits by exploiting flaws in the way people are wired up to make decisions, then assuming the firms themselves are rational profit-maximizers (but see below) they will misallocate their own resources towards those types of activities (e.g., new marketing plans, long-term contracts, etc.) and less towards what neoclassical theory would say they would do (product improvements or price decreases to attract customers). This suggests that some greater degree of regulatory inquiry/oversight into product development processes might be appropriate.

(4) The flip side of people making bad decisions in the context of their individual choices as consumers is people making bad decisions in their role as managers of businesses. It's not like the bugs in our decision-making programs only come into play when we are deciding whether to buy a triple-play option from Verizon. The people who run Verizon are people too, subject to the same flaws. Perhaps someone more familiar with the economic literature than I

am can point out whether there have been systematic studies of decision making errors by businesses as well as consumers.

Feld: There is a further element at work, which is what I often call the "morality" argument in the Chicago School. Despite its claim to being a school of economics, the University of Chicago school rests heavily on an assumption that as long as a theoretical opportunity exists for consumers to learn something and a theoretical possibility for them to take some sort of action which might be effective, that is the only thing that matters. Failure to act is their fault.

The classic example for me is the presence of switching costs and network effects on consumer behavior in broadband. The argument I run into is not even that these costs are minimal (although some make this argument) but that they are *irrelevant*. Worse, any effort by government to minimize switching costs or network effects is "picking winners" and interfering with the "level playing field."

Mind you, it is obvious why industry incumbents push this line. **But it finds a ready home in the Chicago School, which long ago gave up any pretense of being about consumer**

welfare and became instead about worship of the gods of the marketplace.

Problems of Free Peering

March 12: **Tim Cowen:** One issue that I have found to be a problem is peering on a 'free' basis. We saw this in the Worldcom-MCI-Sprint case and the issue was the same as the earlier correspondent networks case. Essentially the issue is that free peering is only feasible when the balance of traffic is balanced within certain boundaries. We saw Worldcom get to a certain point in the late 1990s where it had the opportunity and incentive to charge others for peering, and indeed started to do so. The reason that it started to charge was that it could. (Incidentally it also took the position that it could degrade links and services to its competitors over its own downstream offerings, which is a form of discrimination that can be expected from monopolists).

Rood: As far as I know, this was the "Theory of Harm", put forward by Cremer, Rye and Tirole in their paper for EU DG Competition in the MCI - Worldcom merger case. Just like the "Theory of Harm" by David Farber and Gerald Faulhaber on Instant Messaging dominance from AOL - Time Warner merger case, it was a thought experiment on what could happen. As such it remained

theory.

While I deliberately named the academics above, I could also have pointed at Hal Varian and Jeffrey McKie-Mason in 1994, who thought Internet congestion was a real danger at NAPs and then proposed to start charging for bandwidth. The entire concept of peak-load measurement was not in their paper, and they overlooked that the incremental cost for capacity expansion was considerable lower than the incremental cost for establishing and operating a complex billing arrangement for traffic.

The ability to serve peak-load is the primary source of cost causation in every network, whether it is roads, electricity, water distribution or the number of TV-channels. This is not unique. Serving peak-load is also the genuine cost causation in any capital intensive industrial plant or computing center and even in most labour intensive service operations.

The issue is that there are many ways to make the cost of peak-load divisible over product/service units in industries. You then make a switch from marginal cost and incremental investment to average

cost per unit and call that average cost the incremental cost. These two concepts are not equivalent and it is this step in any "engineering cost model" that determines where the costs fall and who receives the benefits.

The incremental cost, as the average unit cost of an arbitrarily chosen unit of accounting is not a marginal cost. Therefore it is why wholesale operators enjoy high margins when expanding traffic load under regulatory approved incremental costs. Attempts to fix this problem were made by determining Forward Looking, Long-Run incremental cost. This introduced rate-of-return regulation via the back door.

Tom Vest: Bingo

Rood: What the early internet engineers did was choose a different set than long-distance telegraph and telephone engineers and Prussian politicians did in the 2nd half of the 19th century. The early internet engineers came from the academic environment and peering was a natural approach to exchange for them.

The telegraph and telephone accounting between different network operators has been a result of ITU agreements. The ITU was established in 1865 under strong Prussian and Bismarckian influences, among the most relevant was Prussia requiring only to do cross border correspondent business with government owned Telegraph operations. Except for the Victorian British Empire (too powerful) and the USA (too remote, first successful submarine cable just installed) all countries around Prussia had to follow them as Prussia was already a too important trading partner.

The result was the rise of the ITU as a collection of arrangements based on medieval royal privileges and authoritarian economic thought. The was also combined with the rise of 19th century nationalism and militarism, as well as the prime recognition of telegraphs as a military tool. The result was governmental monopoly provisioning.

Napoleon's France has had great advantage by the optical Telegraph of Chappe as it allowed Napoleon to keep in touch with Paris, while engaging in military warfare around Europe. After Waterloo, the French did keep business traffic off the optical telegraph grid and thus obstructing it's useful deploy-

ment in peace time. A social costly decision that was only broken by the electric telegraph (which was introduced rather late in France). That social costly habit of military strategists may also be seen when looking at the US military's privileges in spectrum allocation.

The ITU telegraph regime was later extended to cross-border telephone exchange and radio, after governments started their power grab in this new communications industry around 1900 and then observing the UK and USA siding with them by creating national monopolies too.

Modern technology has opened tremendous possibilities to game (arbitrage) traffic payment systems. The capacity based arrangements are however closer to real cost causation and CAPEX investment decisions and therefore more difficult to game and profit from than the refilling and call back mechanisms and means of arbitrage introduced in the telephone system.

A call for a relaunch of traffic based accounting should also provide a proper assessment for the cost of accounting as well as the extent of the different arbitrage it provokes. Only then one can declare that it delivers more benefits than costs.

Vest: This is the best synthesis of historical, theoretical, and real-world operational insights on this particular subject that I have ever seen. Kudos Hendrik!

Of course, many advocates of traffic-based accounting privately reject the very idea of cost-referenced pricing on philosophical grounds. Producers should/may set prices at whatever level they see fit, based on their own judgments about their own near-term and long-term interests. Proponents of this view tend to assert that overpricing will naturally lead to the emergence of new competitors -- which is exactly why this view is so often coupled with commercial strategies that focus on establishing and defending non-by passable input bottlenecks.

Needless to say, this is perfectly sensible and internally consistent strategic behavior on this particular view of the world.

Bottleneck & Bypass® -- that's still the name of the game in (too) many places. Unfortunately, if either side wins decisively, everybody loses.

Cowen: I have heard it said that 'bill and keep' is the answer to many a maiden's prayer. My concern is that bill and keep has essentially the same economic characteris-

tics as free peering and the same issues arose in relation to correspondent networks; so will only be likely to operate when there is a balance in the relative market power of players in either side of the relationship.

As a purely commercial and accounting matter I have also been concerned about bill and keep etc on the grounds that each activity should carry a revenue stream, and be accounted for as an activity that carries risk and reward. Commercially the trade occurs within a value chain and tax authorities are keen that turnover taxes apply to the different steps in the chain, typically as a % of the value added through the supply chain. I don't understand how a peering or bill and keep model addresses this issue.

Concepts of value are traditionally limited when the data transferred is thought of as a commodity; in reality different data has differential value depending, often, on the use to which it can be put. This is something else that should be thought through in inter-connection relationships between players so that value is properly compensated. Perhaps the best example is the need to have different classes of service for different types of end user application; and the different classes of service and service quality need

to be assured at each inter-connection point for the end to end service to be truly assured. I don't understand how the bill and keep model achieves this, nor do I see how service levels are incentivised in such a model. I guess I just don't get the bill and keep thing, and that is perhaps because I haven't been paying enough attention. Education welcome.

On the other points in the response I tend to agree, but I don't fully follow what you mean by loop co. If you mean a shared local access operator that operates in a particular locations where there is no economically feasible substitute then it has merit but, there are issues. If set up as a separate corporate entity and a profit center it will seek to expand beyond its area and into competitive areas. That would require close regulatory oversight to control. (Expansion happened with the BOCs and the US legal system was inadequate to the task)

Incidentally the comments in the earlier thread about the US legal system are beside the point. I see force in the argument that in certain sectors of the economy; and telecoms is one of them, **the US" legal system" lacks the essential characteristics to be called a 'legal' system at all: it is unequivocally a political sys-**

tem. After all, the FCC is a creature of Congress, passing rules is an executive branch function, weak or ineffective judicial oversight in a system that allows the political appointment of enforcement agencies and judges is a questionable legal system at best: if the legal system is not separate and independent from the political system it fails one fundamental test for the rule of law.

Meanwhile back to Loop Co. Loop Co would also have to have a load of contracts between itself and all service providers. To what extent the loop co is truly a provide of a utility facility and not a service depends on its commodity characteristics. One issue is configuration; both configuration and network deployment affect speed and quality of service so the product lacks the characteristics of a utility and again that requires regulatory oversight. Again the boundary will shift over time and the drivers of a corporation will create conflict. That is why we adopted the functional separation model with regulators on the board of the functionally separate entity; and while respecting the profit motive and ensuring physical separation of staff; targeted incentives for the functionally separate entity etc, we avoided the trap of a separate corporate entity.

Rood: Tim, Three responses:

1. In 1996 the Dutch government effectively forced KPN to withdraw from its Casema cable subsidiary. As a consequence the firm also sold its holdings in many other European cable networks, which totaled somewhere up to 7 million amongst them quite some holdings in the UK.

It was then expected that also France Telecom and Deutsche Telekom would have to sell. That ultimately happened but in quite protracted stages.

Telecom Danmark (TDC) is the last major incumbent that also controls a substantial set of CATV networks in its home market. As TDC is also extremely leveraged (they have been acquired by Private Equity firms who then gave them a major dividend payout that reduced capital with 98%) there is a very big risk of monopoly rent extraction. Hence opening up cable.

In 2008 the European Commission has agreed with Dutch regulator OPTA to open up CATV, however they decided that KPN is not allowed access on their networks as well as CATV operators are not allowed to access KPN's networks.

The Dutch opening however does not apply for broadband, as our regulator tends to think in a very restricted mold about unbundling. They only opened wholesale of an RF downstream channels (for Digital TV), as well as reselling analogue CATV services and not for 5-65 MHz upstream channels.

As far as I was told a major reason the EC agreed was the decision of several member states to withhold support for some parts of the proposed measures of the 2006 Telecom Review, the argument was that as long as the EC showed very limited regard for specific national circumstances, they did not want to shift more regulatory powers to a pan-European level.

In the 2005 round of Market Analyses the EU has taken a rather tough position not allowing any definition of additional markets and rejecting some specific national regulatory arrangements. The commission now has weakened a bit on their attempt to get a near uniform regulatory approach operational.

P.S. I reviewed a brief for the EC made by a consultant to the Finnish government, who attempted to explain their very different approach to regulation of their hundreds

of local telephone companies, an approach which resulted in a lack of wholesale fees. In short, when one places a long-distance telephone call in Finland a consumer gets a bill listing a per minute charge from his originating local telco, the long distance telco and the terminating local or mobile operator for this single call. There is no money flowing between those networks, just interconnection interfaces, they sent their CDRs from their own network leg used via a clearinghouse to each other. Although long distance is competitive in Finland for carrier select, this approach strongly differed from the EU and US model where originating and terminating fees are involved, while the customer is billed by a single entity.

I wonder what they do with Finland this time, as they enforced in 2005 the wholesale model with network owners charging fees to service providers.

Executive Summary

Harvey Newman - Part 2

Last month we began to explore Harvey Newman's discovery of the necessary role for data networks in the world of research in high energy physics. We traced its development in the global network that enables the use of the Large Hadron Collider and examined the Collider itself and the research that it enables.

The concluding portion of our interview with Harvey in this month's issue explores the nature of the global optical virtual private network he created and the technology innovations that his group have made as part of their effort to make an unprecedented amount of data available to the global community of high energy physicists and do so in such a way that they gained the most cost-effective use of the technology on which they depend.

Harvey explains how his MonALISA application that is used in monitoring the LHC global VPN enables operational costs to be kept at a reasonable level and describes why, with the rapid escalation of the amount of data needed for delivery into the hands of a distributed

global network of researchers, capital costs of hardware and optical wavelengths are a more significant issue than operational expenditures. The interview covers via funding strategies chosen for the development of the multi tiered LHC Network. A key insight is the funding of local universities to set up regional centers as part of a network in a way that enables them to become global players. To make his monetary resources grow as fast as possible he also explains how all the development of FDT (the fast data transfer application) has enabled him to, by means of optical circuit switching, run his wavelengths at about 85% of capacity or even higher rather than the more normal roughly 30% common to the best effort Internet.

He explains in some detail how over the last seven or eight years he has participated in the development of the global new e-science community where as the representative of high-energy physics he has gone into countries in Latin America, Eastern Europe, India and places far outside the general purview of European or American science to get local leaders to work with their governments to establish

fiber-optic research networks that expand the reach of e-science in ways where local governments find the necessary investment in the enabling optical network infrastructure to be a cost-effective means of increasing local intellectual capital and competitiveness.

We close with a detailed look at EVO (Enabling Virtual Organizations), the peer-to-peer client/server based telepresence collaboration tool that he has developed in parallel with the network. Its capabilities and uses are extremely impressive. Its costs are also low. But it is not open source and is dependent on licensing through Caltech.

Finally, the tools that his group have developed have the potential to be used to solve some of our infrastructural problems. Unfortunately, predisposition is that federal funding should not be used to give away something that a private corporation could sell, is the primary shaper of policy. As long as this is the case, the benefits of these publicly funded efforts are less likely to pay social dividends in some of the areas of education, public administration and local gov-

ernment where they could otherwise become extremely useful.

Symposium: Issues of Policy and Infrastructure

Summary – As the new administration looks at policy and the economic meltdown continues, regulation is up in the air. We start with the February 25 Supreme Court Decision that essentially removed anti-trust remedies to telco pricing practices and move onto a discussion that says muni-networks that survive in the face of hostile duopoly may have no choice but to offer their own services to citizens. From there the discussion moves into support structures and control there of where we trace a spirited argument between Jim Baler and Erik Cecil over if or how fundamental regulatory premises should be swept aside. From that point on we conclude with a discussion of the impact of the Chicago School in regulatory economics and a behavioral economic point of view as being something that merits a new approach in creating an alternative approach.

Feb 25 Decision voids Antitrust as Remedy for carrier abuse. p. 43

Cecil: As Brian notes, if the rates are legal - e.g.

approved or, in this day and age, even countenanced by a regulatory authority, . . . you are going to have a damn hard time of convincing a judge that there has been an antitrust violation.

Savage: *But even if what AT&T was alleged to have done would not be an anti-trust violation, that does not mean (or at least does not have to mean) that its conduct is not a form of "unjust" or "unreasonable" or "discriminatory" conduct, or some similar formulation based on regulatory, rather than antitrust, principles.*

Cooper: The courts will be the last to change. *The judiciary is filled with lawyers who grew up and entered legal practice during the ascendancy of Chicago School market fundamentalism. . . . I have pushed hard for both the DOJ and FTC under Obama to start a claw back campaign, reviewing all the bad decisions that have been so harmful, so that when the judge cites some lousy precedent, they can pull out the study and show why that precedent was wrong.*

Why Fiber is Always Superior to Copper p. 49

Rood: I think the entire idea that it is the high speed broadband, that brings benefits is misguided, it is the

very different market structure that raises competitive levels with open access to fiber, that brings benefits. . . **A nation, or any sovereign entity, should do an OPEX calculation comparable to the one discussed in the study above at least once, and then ask themselves the simple question: are they trading off the CAPEX required for rewiring against higher OPEX with DOCSIS and DSL technology paths, and is it this trade off that brings higher societal costs when compared to their capital outlay?**

Cecil: p. 53 **any anger directed at any incumbent business per se on this list is misdirected to the extent it is really frustration with market and regulatory design. Outside of documented bad acts there are no "good" or "bad" actors here. All are rational economic actors given their business models and methods.**

I think the greater point, too often lost in the specifics, is that the overall regulatory design is simply out of date. From this lawyer's perspective, for example, the regulatory system in the U.S. is deeply flawed not necessarily because I think it does favor incumbents, but primarily because it is so imprecise and unpredictable to the point of

having very little internal consistency. This is actually worse than always favoring incumbents, but I think we can do a whole lot better than that.

Municipalities Do Need to Offer Services p. 55

St Arnaud: We have been in long series of discussions with an out of territory ILEC. The local team were very keen on our business model as they knew that they only way they could compete with the local ILEC is by someone else underwriting the cost of the infrastructure. But when the local team took the business case back to head office in their home territory the project was killed. The senior management was terrified that other ILECs would invade their turf with a similar strategy. The ILECs are much more interested in protecting their local monopoly rather than competing in someone else's backyard.

So we are stuck looking to provide cable and Internet services ourselves. This actually requires a bigger investment than deploying the actual fiber (at least for a small scale project), because you have to purchase volume Internet transit to be competitive. But the biggest challenges is arranging for delivery of cable TV services. This is a Byzantine world if I ever saw one and much more dif-

ficult and costly than arranging for Internet service. I know that many municipal networks like Lafayette and others have run into the same problem.

Budde: It looks to me that the future has to be in a utilities based open fibre network. In order to built up extra higher margin revenues a value added infrastructure company/ or division should be attached to the utility operating data centres, content hosting, cloud computing, billing, network management, and offer that to individuals (UCG) and companies. In other words stay totally away from the content itself, ISP operation, etc.

Once you have this structure in place content owners and content aggregators will on their own terms start building their business models on that infrastructure and they might consider using the value added infrastructure services, but for the rest want to run this business totally separate and independent from the network operator. Of course they will only become interested if such a network has sufficient mass, that's why very little activity is taking place on the current FttH networks as hardly any of them have reached any mass that warrant these media and services companies to start using them. They rely on the operators building a network

based on the principle: 'built and they will come'. They don't want to make any commitment upfront but are happy to come onboard once you have got the network in place.

So, in most situations, only utilities based organisation can afford to built such networks and this again makes munies an ideal partner in building FttH networks, for all the 'common good' reasons that we have been discussing extensively

New Ways to think About Infrastructure p. 58

Kevin Barron: Fiber is public infrastructure; government builds public infrastructure. . . . If we stop the petty squabbles over layers 0 and 1, we could actually see some of the 10-to-1 return on investment in BB infrastructure. **My fear though, is that even with the collapse of the market, the religious dogma of the "free market uber alles" will continue to obscure the fact that fiber is basic infrastructure (which *requires* public ROW's in most populated areas). We continue to be captives of that dogma, even while our economy melts down, our retirement plans go up in smoke and unemployment hits record**

highs as a result.

The bogeyman of "big government" has been used for 20 years in this country as a smokescreen to effectively capture it and primarily turn it into a hired gun for industry.

P. 60 Cecil The solution I specifically proposed was to eliminate the fee-generating dance - because when I'm running and managing all sorts of these kinds of litigation I get tired of the fees. And it's no secret that these disputes tend to be self-perpetuating. It is, therefore, the very basis of a framework that pits all sorts of common interests against each other. . . . [On p. 61 Cecil proposes a way to "rationally fund" dark fiber utilities.

What I do not think is useful, however, is any perpetuation of failed regulatory models that result in incessant grinding of opposed interests that do not necessarily have to be opposed.

Policy Strategy

Budde: The way I would like to handle this is to maintain the more radical call for change, of course in the hope that we will be able to get at least some

of this implemented at a minimum we will get some people thinking. (As a matter of fact we do already have some very senior people in the Administration at least interested in these more radical approaches.) This will create cracks in the system and then when that happens we need to be opportunistic on how to wiggle ourselves into these cracks to establish changes and that will be a matter of give and take, more along the lines Jim is talking about.

In one of our previous discussions we discussed the pros and cons of structural changes vs. grass root changes (e.g muni broadband). If, through guerrilla warfare, we could get FttH muni penetration reaching 10% or so I bet we will start seeing cracks appearing in the incumbent fortresses and that would accelerate the change enormously. Consequently our bold activities when well executed don't have to see the system start to change overnight, we can be the lever that actually allows the system to implode on itself. **You can see in some of the European and Asian countries how quickly things can change once the incumbents decide that they have milked the old system for long enough and start moving**

toward a more open network environment. It can happen rather quickly (2-3 years)

This is important in framing issues. The Overton Window moves. You start by pushing an idea that is outside of it, and it sounds nutty, but things that had previously been seen as radical suddenly look more acceptable by comparison. . . .

So let's talk about radical and even unthinkable ideas. If we talk about the unthinkable long enough, our ideas that are merely considered sensible will have a chance at becoming policy. And it can keep moving.

What Behavioral Economics Tells Us p. 75

Savage: What behavioral economics tells us in the short run is several things:

(1) Intelligent regulation is not easy. You actually have to know the devil is going on in the markets and industries you want to regulate.

(2) Actual people, in their role as consumers, make systemic and predictable errors in making certain types of decisions. (Someone listed some of the key ones, e.g.: status

quo bias, option overload, fear of loss of current benefits more than outweighs equivalent potential of gain). It follows that people who would sell to them have systemic and predictable incentives to exploit those errors. These are not problems that are fixed or avoided by nominal "competition." Therefore, from a consumer protection standpoint, one should focus on the details of consumer transactions more than we have.

(3) If, as noted above, firms can maximize their profits by exploiting flaws in the way people are wired up to make decisions, then assuming the firms themselves are rational profit-maximizers

(but see below) they will misallocate their own resources towards those types of activities (e.g., new marketing plans, long-term contracts, etc.) and less towards what neoclassical theory would say they would do (product improvements or price decreases to attract customers). This suggests that some greater degree of regulatory inquiry/oversight into product development processes might be appropriate.

(4)The flip side of people making bad decisions in the context of their individual choices as consumers is people making bad decisions in their role as managers of businesses.

Peering p. 77

I see force in the argument that in certain sectors of the economy; and telecoms is one of them, the US" legal system" lacks the essential characteristics to be called a 'legal' system at all: it is unequivocally a political system. After all, the FCC is a creature of Congress, passing rules is an executive branch function, weak or ineffective judicial oversight in a system that allows the political appointment of enforcement agencies and judges is a questionable legal system at best: if the legal system is not separate and independent from the political system it fails one fundamental test for the rule of law.

A Note from the Editor on the May 2009 Format and Presentation

This issue leads off with the final part of a two part interview with Caltech physicist Harvey Newman and four weeks of symposium discussion - including an emphasis on regulatory policy issues

Text, URLs and Executive Summary: I have attempted to identify especially noteworthy text by means of boldface for REALLY good "stuff" . Also the proper Executive Summary in this issue continues. I hope you find it useful. Feedback welcomed. You will also find live URL links and page links in this issue.. (I am also no longer changing British spellings of things like fibre to the American fiber.) Thanks to Sara Wedeman - see sarasworld.blogspot.com/behavioraleconomics/ for assistance with the masthead logo. Captain Cook now charts direction by looking at a compass rosette.

Coming in the June 2009 issue - out about April 30 the promised interview with Frank Coluccio. This treats in very detailed fashion the complex story of the copper strangle hold on our networks. It is paradigm shifting in its importance and is hinted at <http://gordoncook.net/wp/?p=298> Why Fiber is in Every Respect Superior to Copper. See also page 49 of this issue.

Correction: page 51 of April 2009 issue the third complete paragraph beginning: It may only shift demand in time, but that's GREAT. -- should have been attributed to Alexander Harrowell -- ending 4 paragraphs later with " whatever push-messaging protocol you like."

I am omitting the contributors' page since a cumulative list may now be found at http://www.cookreport.com/index.php?option=com_content&view=article&id=121&Itemid=74

Contents

Shaping a New Science Infrastructure

Global Collaboration Mandates Inventive

Solutions to Resource Constraint Problems

Editor's Introduction	p. 1
Building the Global VPN	p. 1
Capital Constraints Meet Network Management and Growth Needs	p. 4
Lambda Switching Advancements – VCAT and LCAS – FDT Gets More out of Each Lambda	p. 7
These Networks Need Dynamically Provisioned Circuit Overlays	p. 13
Digital Divide Issues	p. 15
Monitoring Regional Connectivity	p. 20
Digital Divide Workshops – Putting Brazil on the GLIF	p. 21
Other Workshops	p. 24

Enabling Virtual Organizations

A Low Cost but Highly Efficient

Global Telepresence System

The Economics of EVO	p. 27
More about the Architecture and Use	p. 28
OpenGL's Use in EVO	p. 30
A Multiplicity of Capabilities	p. 34
Summation	p. 36
	p. 40

Symposium Discussion February 18, 2009 - March 17 2009

Issues of Policy and Infrastructure

The Feb. 25 Supreme Court Decision on

Anti-trust in Telco Pricing; Economic

Thought impacts Law Very Slowly; Why

Fiber Always Superior to Copper

By the Time Economic Thought Impacts the Law it's Already

Ten to 20 Years Behind Current Thinking

	p. 43
	p. 45

Why Fiber is Superior to Copper -- The Importance of Technology Choice	p. 49
Municipalities Do Need to Offer Services Given the Obstacles that Vertically Integrated Duopolies Impose	p. 55
Reorder the Process	p. 57
Ways of Thinking about Infrastructure from Technical and Legal Points of View	p. 58
Two Very Different Legal Approaches	p. 60
Rationally Fund Dark Fiber Utilities	p. 61
“Interstructure”	p. 63
Defining an Agenda to Match Our Challenges	p. 64
The Overton Window	p. 67
Behavioral Economics Versus the Chicago School	p. 70
The Real Culprit - Clinton Captured by Wall Street	p. 72
Approaches to the “How To’s” of Regulation	p. 74
What Behavioral Economics Tells Us	p. 74
Problems of Free Peering	p. 77
Executive Summary	p. 81